



**Reference: Building Blocks, Tables, and
Procedures**

Sybase IQ 15.4

DOCUMENT ID: DC38151-01-1540-02

LAST REVISED: May 2012

Copyright © 2012 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Java and all Java-based marks are trademarks or registered trademarks of Oracle and/or its affiliates in the U.S. and other countries.

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Contents

Audience	1
File Locations and Installation Settings	3
Installation Directory Structure	3
How Sybase IQ Locates Files	4
Simple File Searching	5
Extensive File Searching	5
Environment Variables	6
Setting Environment Variables on Windows	6
Running UNIX Environment Source Files	6
Setting Environment Variables on UNIX	7
IQCHARSET Environment Variable	7
IQDIR15 Environment Variable	7
IQPORT Environment Variable	8
IQLANG Environment Variable	8
IQLOGDIR15 Environment Variable	10
IQTIMEOUT Environment Variable	10
IQTMP15 Environment Variable	11
JAVA_HOME Environment Variable	12
LIBPATH or LD_LIBRARY_PATH Environment Variable	12
PATH Environment Variable	12
SQLCONNECT Environment Variable	13
SYBASE Environment Variable	13
\$SYBASE_JRE6_32, \$SYBASE_JRE6_64, \$SYBASE_JRE5_64 Environment Variables	14
SYBASE_OCS Environment Variable	14
Registry Entries	14
Current User and Local Machine Settings	15
Registry Structure	15
Registry Settings on Installation	16

SQL Language Elements	17
Keywords	17
Reserved Words	17
Identifiers	20
Strings	21
Expressions	22
Constants in Expressions	23
Column Names in Expressions	24
Subqueries in Expressions	24
SQL Operators	24
IF Expressions	28
CASE Expressions	28
Compatibility of Expressions and Constants	29
Search Conditions	32
Comparison Conditions	33
Subqueries in Search Conditions	35
ALL or ANY Conditions	38
BETWEEN Conditions	39
LIKE Conditions	39
IN conditions	43
CONTAINS Conditions	44
EXISTS Conditions	44
IS DISTINCT FROM and IS NOT DISTINCT FROM Search Conditions	45
IS NULL Conditions	45
Conditions with Logical Operators	45
NOT Conditions	46
Truth Value Conditions	46
Three-Valued Logic	46
User-Supplied Condition Hints	48
Special Values	54
CURRENT DATABASE Special Value	54
CURRENT DATE Special Value	54
CURRENT PUBLISHER Special Value	55
CURRENT TIME Special Value	55

CURRENT TIMESTAMP Special Value	55
CURRENT USER Special Value	56
LAST USER Special Value	56
SQLCODE Special Value	57
SQLSTATE Special Value	57
TIMESTAMP Special Value	57
USER Special Value	58
Variables	58
Local Variables	58
Connection-Level Variables	59
Global Variables	60
Comments	65
NULL Value	66
SQL Data Types	69
Character Data Types	69
Storage Sizes	70
Character Sets and Code Pages	71
Indexes	71
VARCHAR Data and Trailing Blanks	71
Restriction on CHAR and VARCHAR Data Over 255 Bytes	72
Character Data Compatibility	72
Long strings	73
Numeric Data Types	73
Usage for Numeric Data Types	73
Binary Data Types	78
Usage for Binary Data Types	78
Bit Data Type	83
Bit Data Compatibility	84
Date and Time Data Types	84
Usage for Date and Time Data Types	84
Domains	89
Simple Domains	89
CREATE DOMAIN Statement	90
Domain Compatibility	92

Data Type Conversions	93
Compatibility of String to Datetime Conversions	93
Compatibility of Exported Dates	94
Conversion of BIT to BINARY Data Type	94
Conversion Between BIT and CHAR/VARCHAR Data Types	96
SQL Functions	99
Aggregate Functions	99
Analytical Functions	101
Windowing Aggregate Function Usage	103
Ranking Functions Usage	104
Statistical Aggregate Analytic Function Usage ..	105
Distribution Functions Usage	105
Inter-Row Functions Usage	106
Data Type Conversion Functions	106
Date and Time Functions	107
Date Parts	110
HTTP Functions	112
Numeric Functions	112
String Functions	114
System Functions	116
Connection Properties	119
Properties Available for the Server	120
Properties Available for each Database	120
SQL and Java User-Defined Functions	121
User-Defined Functions in SQL	121
User-Defined Functions in Java	122
Time Series and Forecasting Functions	122
Miscellaneous Functions	122
Alphabetical List of Functions	123
ABS Function [Numeric]	123
ACOS Function [Numeric]	124
ARGN Function [Miscellaneous]	125
ASCII Function [String]	125
ASIN Function [Numeric]	126

ATAN Function [Numeric]	127
ATAN2 Function [Numeric]	128
AVG Function [Aggregate]	129
BFILE Function [Data extraction]	130
BIGINTTOHEX Function [Data Type Conversion]	130
BIT_LENGTH Function [String]	131
BYTE_LENGTH Function [String]	132
BYTE_LENGTH64 Function	133
BYTE_SUBSTR64 and BYTE_SUBSTR Functions	133
CAST Function [Data Type Conversion]	133
CEIL Function [Numeric]	135
CEILING Function [Numeric]	136
CHAR function [String]	136
CHAR_LENGTH Function [String]	137
CHAR_LENGTH64 Function	138
CHARINDEX Function [String]	139
COALESCE Function [Miscellaneous]	140
COL_LENGTH Function [System]	141
COL_NAME Function [System]	141
CONNECTION_PROPERTY Function [System]	142
CONVERT Function [Data Type Conversion]	144
CORR Function [Aggregate]	148
COS Function [Numeric]	149
COT Function [Numeric]	150
COVAR_POP Function [Aggregate]	151
COVAR_SAMP Function [Aggregate]	152
COUNT Function [Aggregate]	153
CUME_DIST Function [Analytical]	154
DATALength Function [System]	155
DATE Function [Date and Time]	156
DATEADD Function [Date and Time]	157
DATECEILING Function [Date and Time]	158

DATEDIFF Function [Date and Time]	161
DATEFLOOR Function [Date and Time]	164
DATEFORMAT Function [Date and Time]	166
DATENAME Function [Date and Time]	168
DATEPART Function [Date and Time]	169
DATEROUND Function [Date and Time]	170
DATETIME Function [Date and Time]	173
DAY Function [Date and Time]	174
DAYNAME Function [Date and Time]	174
DAYS Function [Date and Time]	175
DB_ID Function [System]	176
DB_NAME Function [System]	177
DB_PROPERTY Function [System]	178
DEGREES Function [Numeric]	179
DENSE_RANK Function [Analytical]	180
DIFFERENCE Function [String]	181
DOW Function [Date and Time]	182
ERRORMSG Function [Miscellaneous]	183
EVENT_CONDITION Function [System]	184
EVENT_CONDITION_NAME Function [System]	186
EVENT_PARAMETER Function [System]	186
EXP Function [Numeric]	188
EXP_WEIGHTED_AVG Function [Aggregate] ...	188
FIRST_VALUE Function [Aggregate]	190
FLOOR Function [Numeric]	192
GETDATE Function [Date and Time]	193
GRAPHICAL_PLAN Function [String]	193
GROUPING Function [Aggregate]	196
GROUP_MEMBER Function [System]	197
HEXTOBIGINT Function [Data Type Conversion]	198
HEXTOINT Function [Data Type Conversion]	199
HOUR Function [Date and Time]	201
HOURS Function [Date and Time]	201

HTML_DECODE Function [HTTP]	203
HTML_ENCODE Function [HTTP]	203
HTML_PLAN Function [String]	204
HTTP_DECODE Function [HTTP]	205
HTTP_ENCODE Function [HTTP]	206
HTTP_HEADER Function [HTTP]	206
HTTP_VARIABLE function [HTTP]	206
IFNULL Function [Miscellaneous]	207
INDEX_COL Function [System]	208
INSERTSTR Function [String]	208
INTTOHEX Function [Data Type Conversion] ...	209
ISDATE Function [Date and Time]	211
ISNULL Function [Miscellaneous]	212
ISNUMERIC Function [Miscellaneous]	213
LAG Function [Analytical]	214
LAST_VALUE Function [Aggregate]	215
LCASE Function [String]	217
LEAD Function [Analytical]	219
LEFT Function [String]	220
LEN Function [String]	222
LENGTH Function [String]	223
LIST function [Aggregate]	224
LN Function [Numeric]	224
LOCATE Function [String]	225
LOG Function [Numeric]	226
LOG10 Function [Numeric]	227
LOWER Function [String]	228
LTRIM Function [String]	229
MAX Function [Aggregate]	230
MEDIAN Function [Aggregate]	231
MIN Function [Aggregate]	232
MINUTE Function [Date and Time]	233
MINUTES Function [Date and Time]	234
MOD Function [Numeric]	235
MONTH Function [Date and Time]	236

MONTHNAME Function [Date and Time]	236
MONTHS Function [Date and Time]	237
NEWID Function [Miscellaneous]	239
NEXT_CONNECTION Function [System]	240
NEXT_DATABASE Function [System]	242
NEXT_HTTP_HEADER Function [HTTP]	243
NEXT_HTTP_VARIABLE Function [HTTP]	243
NOW Function [Date and Time]	244
NTILE Function [Analytical]	244
NULLIF Function [Miscellaneous]	246
NUMBER Function [Miscellaneous]	247
OBJECT_ID Function [System]	248
OBJECT_NAME Function [System]	249
OCTET_LENGTH Function [String]	250
PATINDEX Function [String]	250
PERCENT_RANK Function [Analytical]	252
PERCENTILE_CONT Function [Analytical]	254
PERCENTILE_DISC Function [Analytical]	256
PI Function [Numeric]	258
POWER Function [Numeric]	258
PROPERTY Function [System]	259
PROPERTY_DESCRIPTION Function [System]	260
PROPERTY_NAME Function [System]	261
PROPERTY_NUMBER Function [System]	261
QUARTER Function [Date and Time]	262
RADIANS Function [Numeric]	263
RAND Function [Numeric]	264
RANK Function [Analytical]	265
REGR_AVGX Function [Aggregate]	266
REGR_AVGY Function [Aggregate]	267
REGR_COUNT Function [Aggregate]	269
REGR_INTERCEPT Function [Aggregate]	270
REGR_R2 Function [Aggregate]	271
REGR_SLOPE Function [Aggregate]	272

REGR_SXX Function [Aggregate]	274
REGR_SXY Function [Aggregate]	275
REGR_SYY Function [Aggregate]	276
REMAINDER Function [Numeric]	278
REPEAT Function [String]	278
REPLACE Function [String]	279
REPLICATE Function [String]	282
REVERSE Function [String]	283
RIGHT Function [String]	284
ROUND Function [Numeric]	285
ROW_NUMBER Function [Analytical]	286
ROWID Function [Miscellaneous]	287
RTRIM Function [String]	289
SECOND Function [Date and Time]	290
SECONDS Function [Date and Time]	291
SIGN Function [Numeric]	292
SIMILAR Function [String]	292
SIN Function [Numeric]	293
SORTKEY Function [String]	294
SOUNDEX Function [String]	299
SPACE Function [String]	300
SQLFLAGGER Function [Miscellaneous]	301
SQRT Function [Numeric]	302
SQUARE Function [Numeric]	303
STDDEV Function [Aggregate]	303
STDDEV_POP Function [Aggregate]	305
STDDEV_SAMP Function [Aggregate]	306
STR Function [String]	307
STR_REPLACE Function [String]	308
STRING Function [String]	310
STRTOUUID Function [String]	311
STUFF Function [String]	312
SUBSTRING Function [String]	313
SUBSTRING64 Function [String]	315
SUM Function [Aggregate]	315

SUSER_ID Function [System]	316
SUSER_NAME Function [System]	317
TAN Function [Numeric]	318
TODAY Function [Date and time]	319
TRIM Function [String]	319
TRUNCNUM Function [Numeric]	320
TS_ARMA_AR Function [Time Series]	321
TS_ARMA_CONST Function [Time Series]	321
TS_ARMA_MA Function [Time Series]	321
TS_AUTOCORRELATION function [Time Series]	322
TS_AUTO_ARIMA Function [Time Series]	322
TS_AUTO_ARIMA_OUTLIER Function [Time Series]	322
TS_AUTO_ARIMA_RESULT_AIC Function [Time Series]	323
TS_AUTO_ARIMA_RESULT_AICC Function [Time Series]	323
TS_AUTO_ARIMA_RESULT_BIC Function [Time Series]	323
TS_AUTO_ARIMA_RESULT_FORECAST_VA LUE Function [Time Series]	323
TS_AUTO_ARIMA_RESULT_FORECAST_ER ROR Function [Time Series]	324
TS_AUTO_ARIMA_RESULT_MODEL_D Function [Time Series]	324
TS_AUTO_ARIMA_RESULT_MODEL_P Function [Time Series]	324
TS_AUTO_ARIMA_RESULT_MODEL_Q [Time Series]	324
TS_AUTO_ARIMA_RESULT_MODEL_S function [Time Series]	325
TS_AUTO_ARIMA_RESULT_RESIDUAL_SIG MA [Time Series]	325
TS_AUTO_UNI_AR Function [Time Series]	325

TS_BOX_COX_XFORM Function [Time Series]	325
TS_DIFFERENCE Function [Time Series]	326
TS_DOUBLE_ARRAY [Time Series]	326
TS_ESTIMATE_MISSING Function [Time Series]	326
TS_GARCH Function [Time Series]	326
TS_GARCH_RESULT_A Function [Time Series]	327
TS_GARCH_RESULT_AIC Function [Time Series]	327
TS_GARCH_RESULT_USER [Time Series]	327
TS_INT_ARRAY [Time Series]	327
TS_LACK_OF_FIT Function [Time Series]	328
TS_LACK_OF_FIT_P Function [Time Series]	328
TS_MAX_ARMA_AR Function [Time Series]	328
TS_MAX_ARMA_CONST Function [Time Series]	329
TS_MAX_ARMA_LIKELIHOOD Function [Time Series]	329
TS_MAX_ARMA_MA Function [Time Series]	329
TS_OUTLIER_IDENTIFICATION Function [Time Series]	330
TS_PARTIAL_AUTOCORRELATION Function [Time Series]	330
TS_VWAP Function [Time Series]	330
UCASE Function [String]	331
UPPER Function [String]	332
USER_ID Function [System]	333
USER_NAME Function [System]	333
UIDTOSTR Function [String]	334
VAR_POP Function [Aggregate]	335
VAR_SAMP Function [Aggregate]	336
VARIANCE Function [Aggregate]	338
WEEKS Function [Date and Time]	340

WEIGHTED_AVG Function [Aggregate]	341
WIDTH_BUCKET Function [Numerical]	343
YEAR Function [Date and Time]	345
YEARS Function [Date and Time]	346
YMD Function [Date and Time]	347
Differences from Other SQL Dialects	349
Dates	349
Integrity	349
Joins	350
Updates	350
Altering Tables	350
Subqueries Not Always Allowed	351
Additional Functions	351
Cursors	351
Physical Limitations	353
System Procedures	357
Syntax Rules for Stored Procedures	357
Understanding Statistics Reported by Stored Procedures	358
System Stored Procedures	358
sa_char_terms System Procedure	359
sa_dependent_views system procedure	359
sa_external_library_unload Procedure	360
sa_list_external_library Procedure	360
sa_nchar_terms system Procedure	360
sa_text_index_vocab Procedure	361
sa_get_user_status system procedure	361
sp_expireallpasswords Procedure	362
sp_iqaddlogin Procedure	363
sp_iqbackupdetails Procedure	364
sp_iqbackupsummary Procedure	366
sp_iqcardinality_analysis Procedure	368
sp_iqcheckdb Procedure	370
sp_iqcheckoptions Procedure	379
sp_iqclient_lookup Procedure	380

sp_iqcolumn Procedure	382
sp_iqcolumnuse Procedure	385
sp_iqconnection Procedure	386
sp_iqconstraint Procedure	390
sp_iqcontext Procedure	391
sp_iqcopyloginpolicy Procedure	394
sp_iqcursorinfo Procedure	395
sp_iqdatatype Procedure	398
sp_iqdbsize Procedure	401
sp_iqdbspace Procedure	402
sp_iqdbspaceinfo Procedure	405
sp_iqdbspaceobjectinfo Procedure	409
sp_iqdbstatistics Procedure	413
sp_iqdroplogin Procedure	414
sp_iqemptyfile Procedure	415
sp_iqestjoin Procedure	416
sp_iqestdbspaces Procedure	417
sp_iqestspace Procedure	420
sp_iqevent Procedure	420
sp_iqfile Procedure	424
sp_iqhelp Procedure	426
sp_iqindex and sp_iqindex_alt Procedures	433
sp_iqindexadvice Procedure	437
sp_iqindexfragmentation Procedure	438
sp_iqindexinfo Procedure	440
sp_iqindexmetadata Procedure	442
sp_iqindexsize Procedure	443
sp_iqindexuse Procedure	446
sp_iqjoinindex Procedure	447
sp_iqjoinindexsize Procedure	451
sp_iqlmconfig Procedure	452
sp_iqlocks Procedure	454
sp_iqmodifyadmin Procedure	457
sp_iqmodifylogin Procedure	458
sp_iqmpxcheckdqpconfig Procedure	459

sp_iqmpxfilestatus procedure	461
sp_iqmpxinconnpoolinfo Procedure	461
sp_iqmpxinheartbeatinfo Procedure	462
sp_iqmpxinfo Procedure	462
sp_iqmpxvalidate Procedure	462
sp_iqmpxversioninfo Procedure	462
sp_iqobjectinfo Procedure	462
sp_iqpassword Procedure	465
sp_iqkeys Procedure	467
sp_iqprocedure Procedure	469
sp_iqprocparm Procedure	472
sp_iqrebuildindex Procedure	475
sp_iqrename Procedure	477
sp_iq_reset_identity Procedure	480
sp_iqrestoreaction Procedure	481
sp_iqrowdensity Procedure	482
sp_iqsharedtempdistrib procedure	483
sp_iqshowpsexo Procedure	484
sp_iqspaceinfo Procedure	486
sp_iqspaceused Procedure	487
sp_iqstatistics Procedure	489
sp_iqstatus Procedure	491
sp_iqsysmon Procedure	494
sp_iqtable Procedure	501
sp_iqtablesize Procedure	504
sp_iqtableuse Procedure	506
sp_iqtransaction Procedure	506
sp_iqunusedcolumn Procedure	510
sp_iqunusedindex Procedure	511
sp_iqunusedtable Procedure	513
sp_iqversionuse Procedure	514
sp_iqview Procedure	515
sp_iqwho Procedure	517
sp_iqworkmon Procedure	521
Catalog Stored Procedures	523

sa_ansi_standard_packages system procedure	523
sa_audit_string system procedure	524
sa_checkpoint_execute System Procedure	524
sa_conn_activity system procedure	526
sa_conn_info system procedure	527
sa_conn_list System Procedure	531
sa_conn_properties System Procedure	531
sa_db_info system procedure	532
sa_db_properties system procedure	533
sa_disable_auditing_type system procedure	534
sa_disk_free_space system procedure	535
sa_enable_auditing_type system procedure	537
sa_eng_properties system procedure	538
sa_flush_cache system procedure	539
sa_make_object system procedure	539
sa_rowgenerator system procedure	540
sa_server_option system procedure	543
sa_set_http_header system procedure	552
sa_set_http_option system procedure	553
sa_table_page_usage system procedure	557
sa_validate system procedure	558
sa_verify_password system procedure	559
sp_login_environment system procedure	559
sp_remote_columns System Procedure	560
sp_remote_exported_keys system procedure	560
sp_remote_imported_keys system procedure	562
sp_remote_primary_keys System Procedure	564
sp_remote_tables system procedure	564
sp_servercaps system procedure	566
sp_tsql_environment system procedure	566
Adaptive Server Enterprise System and Catalog Procedures	567

Adaptive Server Enterprise System Procedures	568
Adaptive Server Enterprise Catalog Procedures	570
SQL Anywhere Supported Procedures	570
System Tables and Views	573
System Tables	573
SYS.DUMMY Table Versus IQ_DUMMY Table	576
System Views	577
Consolidated Views	577
Compatibility Views	578
ASE T-SQL Compatibility Views	578
SYSARTICLE system view	578
SYSARTICLECOL system view	579
SYSARTICLECOLS consolidated view	579
SYSARTICLES consolidated view	580
SYSCAPABILITIES consolidated view	580
SYSCAPABILITY system view	580
SYSCAPABILITYNAME system view	581
SYSCATALOG consolidated view	581
SYSCHECK system view	582
SYSCOLAUTH consolidated view	582
SYSCOLLATION compatibility view (deprecated)	582
SYSCOLLATIONMAPPINGS compatibility view (deprecated)	583
SYSCOLPERM system view	583
SYSCOLSTAT system view	584
SYSCOLSTATS consolidated view	585
SYSCOLUMN compatibility view (deprecated)	585
SYSCOLUMNS consolidated view	586
SYSCOLUMNS ASE Compatibility View	586
SYSCOMMENTS ASE Compatibility View	587

SYSCONSTRAINT system view	587
SYSDBFILe system view	588
SYSDbSPACE system view	589
SYSDbSPACEPERM System View	589
SYSDEPENDENCY system view	590
SYSDOMAIN system view	590
SYSEVENT system view	591
SYSEVENTTYPE System View	592
SYSEXTERNENV system view	592
SYSEXTERNENVOBJECT system view	593
SYSEXTERNLOGIN system view	594
SYSFILE compatibility view (deprecated)	595
SYSFKCOL compatibility view (deprecated)	595
SYSFKEY system view	596
SYSFOREIGNKEY compatibility view (deprecated)	597
SYSFOREIGNKEYS consolidated view	597
SYSGROUP system view	598
SYSGROUPS consolidated view	598
SYSHISTORY system view	600
SYSIDX system view	603
SYSIDXCOL system view	605
SYSINDEX compatibility view (deprecated)	606
SYSINDEXES consolidated view	606
SYSINDEXES ASE Compatibility View	607
SYSINFO compatibility view (deprecated)	608
SYSIQBACKUPHISTORY System View	608
SYSIQBACKUPHISTORYDETAIL System View	609
SYSIQCOLUMN System View (Deprecated) ...	610
SYSIQDBFILE System View	611
SYSIQDBSPACE System View	612
SYSIQFILE System View (Deprecated)	612
SYSIQIDX System View	613
SYSIQINFO System View	614

SYSIQJOINIDX System View	615
SYSIQJOININDEX System View (Deprecated)	616
SYSIQJOINIXCOLUMN System View	616
SYSIQJOINIXTABLE System View	617
SYSIQLOGICALSERVER System View	618
SYSIQLOGINPOLICYLSINFO System View	618
SYSIQLSLOGINPOLICYOPTION System View	618
SYSIQLSMEMBER System View	618
SYSIQLSMEMBERS Consolidated View	619
SYSIQLSLOGINPOLICIES Consolidated View	619
SYSIQLSPOLICY System View	619
SYSIQLSPOLICYOPTION System View	619
SYSIQMPXSERVER System View	619
SYSIQOBJECTS ASE Compatibility View	619
SYSIQPARTITIONCOLUMN System View	620
SYSIQTAB System View	620
SYSIQTABCOL System View	621
SYSIQTABLE System View (Deprecated)	622
SYSIQVINDEX ASE Compatibility View	623
SYSIXCOL compatibility view (deprecated)	623
SYSJAR system view	623
SYSJARCOMPONENT system view	624
SYSJAVACLASS system view	625
SYSLOGINMAP system view	625
SYSLOGINPOLICY system view	626
SYSLOGINPOLICYOPTION system view	627
SYSLOGINS ASE Compatibility View	627
SYSMVOPTION system view	627
SYSMVOPTIONNAME system view	628
SYSOBJECT system view	628
SYSOBJECTS ASE Compatibility View	629
SYSOPTION system view	630

SYSOPTIONS consolidated view630
 SYSOPTSTAT system view630
 SYSPARTITION System View631
 SYSPARTITIONKEY System View632
 SYSPARTITIONScheme System View632
 SYSPHYSIDX system view633
 SYSPROCAUTH consolidated view633
 SYSPROCEDURE system view634
 SYSPROCPARM system view635
 SYSPROCPARMS consolidated view637
 SYSPROCPERM system view638
 SYSPROCS consolidated view638
 SYSPROXYTAB system view638
 SYSPUBLICATION system view639
 SYSPUBLICATIONS consolidated view640
 SYSREMARK system view641
 SYSREMOTEOPTION System View641
 SYSREMOTEOPTION2 consolidated view641
 SYSREMOTEOPTIONS consolidated view642
 SYSREMOTEOPTIONTYPE system view642
 SYSREMOTETYPE system view643
 SYSREMOTETYPES consolidated view643
 SYSREMOTEUSER system view644
 SYSREMOTEUSERS consolidated view645
 SYSSCHEDULE System View646
 SYSSERVER system view646
 SYSSOURCE system view647
 SYSSQLSERVERTYPE system view647
 SYSSUBPARTITIONKEY System View648
 SYSSUBSCRIPTION system view648
 SYSSUBSCRIPTIONS consolidated view648
 SYSSYNC system view649
 SYSSYNC2 consolidated view650
 SYSSYNCPUBLICATIONDEFAULTS
 consolidated view651

SYSSYNCS consolidated view	651
SYSSYNCS SCRIPT system view	652
SYSSYNCS SCRIPTS consolidated view	652
SYSSYNCS SUBSCRIPTIONS consolidated view	653
SYSSYNCS USERS consolidated view	653
SYSTAB system view	654
SYSTAB AUTH consolidated view	657
SYSTAB COL system view	658
SYSTABLE compatibility view (deprecated)	660
SYSTABLE PERM system view	661
SYSTEXT CONFIG system view	664
SYSTEXT IDX system view	665
SYSTEXT IDXTAB system view	666
SYSTRIGGER system view	667
SYSTRIGGERS consolidated view	669
SYSTYPE MAP system view	670
SYSTYPES ASE Compatibility View	670
SYSUSER system view	671
SYSUSER AUTH compatibility view (deprecated)	672
SYSUSER AUTHORITY system view	672
SYSUSER LIST compatibility view (deprecated)	672
SYSUSER MESSAGE system view	673
SYSUSER OPTIONS consolidated view	673
SYSUSER PERM compatibility view (deprecated)	674
SYSUSER PERMS compatibility view (deprecated)	674
SYSUSER TYPE system view	675
SYSUSERS ASE Compatibility View	676
SYSVIEW System View	676
SYSVIEWS consolidated view	677
SYSWEBSERVICE system view	677

Transact-SQL Compatibility Views	678
Compatibility with Other Sybase Databases	683
About SQL Anywhere	683
An Overview of Transact-SQL Support	683
Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ Architectures	684
Servers and Databases	684
Space Allocation and Device Management	685
System Tables, Catalog Store, and IQ Store	685
Administrative Roles	686
Data Types	687
Bit Data Type	687
Character Data Types	688
Binary Data Types	689
Date, Time, Datetime, and Timestamp Data Types	689
Numeric Data Types	691
Text Data Type	691
Image Data Type	692
Java Data Types	692
Data Definition Language	692
Creating a Transact-SQL Compatible Database from Sybase Central	692
Creating a Transact-SQL Compatible Database Using the CREATE DATABASE statement	692
Case-sensitivity	692
Ensuring Compatible Object Names	693
Considerations When Using the CREATE TABLE Statement	694
Considerations When Using the CREATE DEFAULT, CREATE RULE, and CREATE DOMAIN Statements	697
Considerations When Using the CREATE TRIGGER Statement	697

Considerations When Using the CREATE	
INDEX Statement	697
Users, Groups, and Permissions	698
Load Formats	699
Options for Transact-SQL Compatibility	700
Data Manipulation Language	700
General Guidelines for Writing Portable SQL	
.....	700
Criteria for Writing Compatible Queries	700
Subquery Support	701
GROUP BY Clause Support	702
COMPUTE Clause Support	702
WHERE Clause Support	702
Transact-SQL Outer Joins Support	702
ANSI Joins Support	703
Null Comparisons Support	703
Zero-length Strings Support	703
HOLDLOCK, SHARED, and FOR BROWSE	
Support	704
SQL Function Support	704
OLAP Function Support	705
System Function Support	706
User Defined Function Support	706
Differences Interpreting Arithmetic Expressions	
on Dates	706
SELECT INTO Statement Support	707
Updatable Views Support	707
Support for FROM Clause in UPDATE and	
DELETE	707
Transact-SQL Procedure Language Overview	707
Transact-SQL Stored Procedure Overview	708
Transact-SQL Batch Overview	708
SQL Statements in Procedures and Batches	
.....	708
Automatic Translation of Stored Procedures	710

Result Sets from Transact-SQL Procedures	710
Variables in Transact-SQL Procedures	711
Error Handling in Transact-SQL Procedures	712
Using the RAISERROR Statement in Procedures	713
Transact-SQL-like Error Handling in the Watcom-SQL Dialect	713
SQL Anywhere and Sybase IQ Differences and Shared Functionality	714
SQL Anywhere Server and Database Startup and Administration	714
SQL Anywhere Database Options	715
SQL Anywhere Data Definition Language (DDL) Differences	715
SQL Anywhere Data Manipulation Language (DML) Differences	715
Adaptive Server Enterprise and Sybase IQ Differences and Shared Functionality	716
Adaptive Server Enterprise Stored Procedures	716
Adaptive Server Enterprise System Views	717
Index	719

Contents

Audience

This book is intended for Sybase® IQ users who require reference material for SQL statements, language elements, data types, functions, system procedures, and system tables.

Other books provide more context on how to perform particular tasks. Use this book to get information about SQL syntax, parameters, and options. For command line utility start-up parameters, see the *Utility Guide*.

Audience

File Locations and Installation Settings

These topics describe the installation and operating system settings used by Sybase IQ.

Depending on the operating system, these settings may be stored as environment variables, initialization file entries, or registry entries.

Installation Directory Structure

When you install Sybase IQ, several directories may be created. This section describes the directory structure.

The directories created depend on which options you choose during installation and which directories already exist in your Sybase directory (the directory defined by *\$SYBASE* on UNIX or *%SYBASE%* on Windows).

By default, Sybase IQ software is installed in a unique subdirectory under the Sybase directory. This subdirectory is called the installation directory. Other tools provided with Sybase IQ have unique subdirectories under the Sybase directory. This section describes only the subdirectory structure for Sybase IQ.

By default, the Sybase IQ directory is *IQ-15_4*. The location of *IQ-15_4* varies, depending on where you install Sybase IQ. The *IQ-15_4* directory is also referred to by the environment variable *\$IQDIR15* on UNIX or *%IQDIR15%* on Windows.

The Sybase IQ directory holds a number of directories and files:

- **Demo directory** (*%ALLUSERSPROFILE%/SybaseIQ/demo*) – Holds the tools required to build the *iqdemo* database. The *iqdemo* database files are *iqdemo.db*, *iqdemo.iq*, *iqdemo.iqmsg*, and *iqdemo.iqtmp*. The demo database itself is not shipped with Sybase IQ.
- **The subdirectory /demo/adata** – Holds 15.x data to allow the creation of the 15.x *iqdemo* database. The subdirectory */demo/demodata* holds Sybase IQ 12.7 data to allow the creation of an *iqdemo* database that has the same schema layout and data as the IQ 12.7 *asiqdemo* database. Use */demo/mkqiqdemo.bat* on Windows and *demo/mkqiqdemo.sh* on UNIX to create the 15.x *iqdemo* database. The *iqdemo* database can be used to demonstrate problems to Technical Support.
- **Scripts directory** (*IQ-15_4/scripts*) – Holds some scripts used in examples and when creating catalog objects like stored procedures. Do not edit these scripts. If you edit, delete, or move these scripts, the server will not operate properly.
- **Samples directories** – The *samples* directory contains SQL samples and user-defined function (UDF) samples. *%ALLUSERSPROFILE%/SybaseIQ/samples/*

`sqlanywhere` contains directories of SQL samples. The `sqlanywhere/c` directory holds C++ examples that illustrate using ESQL (embedded SQL) and C with SQL Anywhere®. Because Sybase IQ and SQL Anywhere share common code, you can modify these examples for use with Sybase IQ. The `%ALLUSERSPROFILE%/SybaseIQ/samples/udf` directory holds sample C++ scalar and aggregate UDFs.

- **Executable directories** – Hold executables, libraries, help files, and so on. On UNIX, executable subdirectories include IQ-15_4 subdirectories `/bin64`, `/lib64`, `/logfiles`, `/res`, and `/tix`. On Windows, these include IQ-15_4 subdirectories `\h`, `\install`, `\java`, and `\bin32`.

How Sybase IQ Locates Files

To start and run, Sybase IQ must find and access several types of files. Several directories or files with identical names may reside on a system.

Understanding how Sybase IQ finds these files is important to ensure that you use the correct files. The types of files include but are not limited to:

- **Libraries** – might include product libraries or system libraries. File name extensions include `.so.nnn` or `.so` on UNIX, or `.dll` or `.lib` on Windows. These files are required to run Sybase IQ. If an incorrect DLL is found, a version mismatch error may occur. For example, library files might be found in `$IQDIR15/lib64` or `$SYBASE/$SYBASE_OCS/lib64` on UNIX, or `%IQDIR15%\bin32` or `%SYBASE%\%SYBASE_OCS\dll` on Windows. An empty directory, `$IQDIR15/usrlib`, lets you supersede default libraries with custom libraries and patches, because `start_iq` includes `usrlib` before regular library directories. Sybase IQ uses both Adaptive Server Enterprise and SQL Anywhere libraries. If either of these products have already been installed on your system, note the directory where they are installed to avoid confusion.
- **Interface files** – required to run Sybase IQ. For example, `.odbc.ini` and `utility_db.ini` on UNIX, and `util_db.ini` on Windows. For more information about these files, see the *System Administration Guide: Volume 1* and the *Installation and Configuration Guide*.
- **Configuration files** – used to specify connection parameters. Examples include `default.cfg` on Windows or `iqdemo.cfg`.
- **Database files** – store the data and metadata. For example: `iqdemo.db`, `iqdemo.iq`, `iqdemo.iqmsg`, `iqdemo.iqtmp`.
- **Log files** – store information about the current session on the server and connected database. For example, a server log might be named `%ALLUSERSPROFILE%/SybaseIQ/IQ15_4/logfiles/yourservername.0006.srvlog`. The database log (for example, `%ALLUSERSPROFILE%/SybaseIQ/IQ-15_4/demo/iqdemo.log`) is created when you connect to the database. For more information about these files, see the *Installation and Configuration Guide*.

- Product scripts – are sample files that show how to create, populate, and upgrade databases.
- User files – include flat files used with the LOAD command and SQL scripts used with tools such as Interactive SQL.
- Temporary files – created by Sybase IQ to store temporary information for operations like performing sorts for queries.

Some file names are specified in SQL statements and must be located at runtime. Examples of SQL statements that use file names include the following:

- **INSTALL** statement – the name of the file that holds Java classes.
- **LOAD TABLE** statement – the name of the file from which data should be loaded.
- **CREATE DATABASE** statement – A file name is needed for this statement and similar statements that can create files.

In some cases, Sybase IQ uses a simple algorithm to locate files. In other cases, a more extensive search is carried out.

Simple File Searching

In many SQL statements such as **LOAD TABLE** or **CREATE DATABASE**, the file name is interpreted as relative to the current working directory of the database server; that is, where the server was started.

Also, when a database server is started and a database file name (DBF parameter) is supplied, the path is interpreted as relative to the directory in which the server was started.

Extensive File Searching

Sybase IQ programs, including the database server and administration utilities, carry out extensive searches for required files, such as DLLs or shared libraries. In these cases, Sybase IQ programs look for files in the following order:

1. *The executable directory* – the directory in which the program executable is held. Also, directories with the following paths relative to the program executable directory:
 - Parent of the executable directory.
 - A child of the parent directory named `scripts`.
2. *Current working directory* – when a program is started, it has a current working directory (the directory from which it is started). This directory is searched for required files.
3. *Location registry entry* – on a Windows installation, Sybase IQ adds a LOCATION registry entry. The indicated directory is searched, followed by the following:
 - A child named `scripts`
 - A child with the operating system name (`bin32`, `bin`, and so on)
4. *System-specific directories* – this includes directories where common operating system files are held, such as the Windows directory and the Windows\system directory on Windows.

File Locations and Installation Settings

5. *CLASSPATH* directories – for Java files, directories listed in the CLASSPATH environment variable are searched to locate files.
6. *PATH* directories – directories in the system path and the user's path are searched to locate files.
7. *LIBRARY PATH* directories – directories listed in the LIBPATH environment variable are searched for shared libraries.

Environment Variables

Sybase IQ uses environment variables to store various types of information; not all variables need to be set in all circumstances.

Setting Environment Variables on Windows

On Windows platforms, the installation program automatically sets all environmental variables, so no changes are necessary. However, if you must set optional variables or change defaults, use this procedure.

1. On your desktop, right-click My Computer and select Properties from the submenu.
2. Click the Advanced tab.
3. Click the Environment Variables button.

The Environment Variables dialog opens.

- a) If the environment variable does not already exist, click New and type the variable name and its value in the spaces provided; then click OK.
- b) If the variable does exist, select it from the list of System Variables or User Variables, click Edit, and make any modifications in the Variable Value field. Then click OK to capture the setting.

Note: See the Microsoft Windows documentation for an explanation of user variables and system variables.

Running UNIX Environment Source Files

Environment source files set the required environment variables on UNIX.

Issue the following command to set all required environment variables.

1. For the Bourne/Korn shell:

```
. $SYBASE/IQ-15_4/IQ-15_4.sh
```

2. For the C shell:

```
source $SYBASE/IQ-15_4/IQ-15_4.csh;  
rehash
```


Setting Environment Variables on UNIX

On UNIX platforms, running environment source files sets the required environment variables. However, if you must set optional variables or change defaults, use this procedure.

1. To check the setting for an environment variable, use:

```
echo $variable-name
```

For example, to see the setting for the \$SYBASE variable:

```
% echo $SYBASE
```

```
/server1/users/test/sybase
```

2. In one of your startup files (.cshrc, .shrc, .login), add a line that sets the variable.

In some shells (such as sh, bash, ksh) the line is:

```
VARIABLE=value;export VARIABLE
```

In other shells (such as csh, tsch) the line is:

```
setenv VARIABLE "value"
```

IQCHARSET Environment Variable

IQCHARSET sets the default character set.

Charset is a character set name. For example, setting **IQCHARSET=cp1252** sets the default character set to cp1252.

The first of the following values set determines the default character set.

- IQCHARSET environment variable
- Query the operating system

If no character set information is specified, use `iso_1` for UNIX, or `cp850` otherwise.

Setting

```
IQCHARSET=charset
```

IQDIR15 Environment Variable

IQDIR15 identifies the location of the Sybase IQ directory and is the location for other directories and files under that directory.

- `IQDIR15/bin[64]/util_db.ini` holds the login ID and password for the utility database, `utility_db`. The installation program lets you change these from their default values, login ID “DBA” and password “sql.”
- `IQDIR15/logfiles` is the default location for the server log and backup/restore log (the backup history file). You can override this default by setting the `IQLOGDIR15` environment variable.

File Locations and Installation Settings

- `IQDIR15/demo` is the location for the `iqdemo` database files.

Setting

```
IQDIR15 = ${SYBASE}/IQ-15_4
```

Operating System

(Required) Set by the environment source file or the installation program. This default setting can be changed on Windows.

IQPORT Environment Variable

IQPORT overrides the default value for the Sybase IQ Agent port number, which is used for communications between the Sybase IQ plug-in and Agent.

Note: Once the agent starts, you cannot change the port value.

1099 is the plug-in default value when searching for an agent process on any given port. If the plug-in finds no agent on this port, it displays a prompt so that you can specify the correct port value.

Setting

```
IQPORT = 5556
```

Operating System

Optional. If the user did not specify IQPORT in the environment source file, the port number defaults to 1099. You can change this default value, provided you do so before the plug-in starts. You can set this variable by supplying the **-DIQPORT** argument to **scjview** when starting Sybase Central. For example:

```
scjview -DIQPORT=3345
```

IQLANG Environment Variable

IQLANG sets the default language.

Language_code is the two-letter combination representing a language. For example, setting **IQLANG=DE** sets the default language to German.

The first of the following values set determines the default language.

- IQLANG environment variable
- Registry (Windows only) as set by the installer
- Query the operating system

If no language information is set, English is the default.

Setting

```
IQLANG=language_code
```

Operating System

Optional but recommended in non-English environments.

List of Language Label Values

Valid language label values and their equivalent ISO 639 language codes. Set the two-letter ISO_639 language code in the IQ_LANG environment variable.

Language	ISO_639 Language Code	Language Label	Alternative Label
Arabic	AR	arabic	N/A
Czech	CS	czech	N/A
Danish	DA	danish	N/A
Dutch	NL	dutch	N/A
English	EN	us_english	english
Finnish	FI	finnish	N/A
French	FR	french	N/A
German	DE	german	N/A
Greek	EL	greek	N/A
Hebrew	HE	hebrew	N/A
Hungarian	HU	hungarian	N/A
Italian	IT	italian	N/A
Japanese	JA	japanese	N/A
Korean	KO	korean	N/A
Lithuanian	LT	lithuanian	N/A
Norwegian	NO	norwegian	norweg
Polish	PL	polish	N/A
Portuguese	PT	portuguese	portugue
Russian	RU	russian	N/A
Simplified Chinese	ZH	chinese	simpchin
Spanish	ES	spanish	N/A

Language	ISO_639 Language Code	Language Label	Alternative Label
Swedish	SV	swedish	N/A
Thai	TH	thai	N/A
Traditional Chinese	TW	tchinese	tradchin
Turkish	TR	turkish	N/A
Ukrainian	UK	ukrainian	N/A

IQLOGDIR15 Environment Variable

The IQLOGDIR15 environment variable defines the location of various log files. IQLOGDIR15 is not set by the installation program.

- The server log is in the file `servername.nnnn.srvlog` (where `nnnn` is the number of times the server has been started) in the directory specified by `$IQLOGDIR15`.

If IQLOGDIR15 is not set to a valid, write-enabled directory, then most utilities, including `start_iq`, use the default location `$IQDIR15/logfiles` for all server logs.

Setting

```
IQLOGDIR15 = path
```

Operating System

Optional.

IQTIMEOUT Environment Variable

IQTIMEOUT overrides the default wait time of five minutes.

The argument *nnn* is the number of minutes for the Sybase IQ Agent to wait. For example:

- To wait 45 minutes (Korn or Bourne shell):

```
IQTIMEOUT=45
export IQTIMEOUT
```

- To wait an hour (C shell):

```
setenv IQTIMEOUT 60
```

The Sybase IQ Agent waits indefinitely for a process to complete. Setting a wait time is recommended when creating or synchronizing query servers for a multiplex with a very large catalog store. Large catalog stores extend the time needed for the `dbbackup` part of synchronization, and increasing the wait time accommodates a larger synchronize.

Note: Set IQTIMEOUT before you start the Sybase IQ Agent.

Setting

```
IQTIMEOUT = nnn
```

Operating System

Optional but recommended in multiplex environments. See *Using Sybase IQ Multiplex*.

IQTMP15 Environment Variable

The IQTMP15 environment variable is not set by the installation program. IQTMP15 is used by Sybase IQ to indicate a directory where temporary files are kept.

The IQTMP15 environment variable should point to a local directory for those using NFS (network file system), which permits the IQTMP15 directory to purge directories and files that are no longer needed as client connections are closed. Each client connection creates several directories and files in a temporary directory. These are needed only for the duration of the connection. The directory must have write permissions for all users who connect to the server.

Note: The temporary files whose location is defined by IQTMP15 are files used by the client and server. This variable does not control the default location of your IQ temporary store. The **CREATE DATABASE** statement controls the default location of your IQ temporary store. See *Reference: Statements and Options*.

Warning! Do not set IQTMP15 to \$SYBASE or \$IQDIR15.

If you do not explicitly set IQTMP15, then the Sybase IQ Agent sets IQTMP15 to a subdirectory in the UNIX directory /tmp.

If more than one database server is running on a machine, each server and associated local client needs a separate temporary directory to avoid conflicts. When you do not specify a port or engine number for connection, Sybase IQ uses shared memory connectivity instead of network connectivity.

To avoid conflicts when using shared memory:

- Create a temporary directory dedicated to each server. Make sure that each local client uses the same temporary directory as its server by setting the IQTMP15 environment variable explicitly in both environments.
- Create a data source name in the .odbc.ini file (on UNIX) for each server and provide detailed connection information. See the *Installation and Configuration Guide*.
- Use connection strings that specify explicit parameters instead of relying on defaults.
- Confirm connections by issuing:

```
SELECT "database name is" = db_name(), "servername_is" =  
@@servername
```

Setting

```
IQTMP15 = temp_directory
```

Operating System

Optional on UNIX. Not used on Windows platforms.

JAVA_HOME Environment Variable

Defines the JRE home which points to directory containing bin/java.

Used if the location of the Java VM is not set in the \$SYBASE_JRE6_32, \$SYBASE_JRE6_64, or \$SYBASE_JRE5_64 environment variables.

JAVA_HOME is commonly created when installing a VM.

On UNIX, run the SYBASE.csh (C shell) or SYBASE.sh (Bourne or Korn shell) environment source file to find and start the correct JRE for the IQ engine. The Java VM location specified in JAVA_HOME takes precedence over the location returned by SYBASE.csh or SYBASE.sh. If neither JAVA_HOME nor the SYBASE.csh or SYBASE.sh scripts locate the Java VM, IQ will not load a Java VM.

Settings

```
JAVA_HOME = Sybase/shared/JRE<version>
```

Operating System

Required.

LIBPATH or LD_LIBRARY_PATH Environment Variable

LIBPATH or LD_LIBRARY_PATH specifies the directories where Sybase IQ shared libraries are located.

On UNIX, set the library path variable by running the environment source file.

Settings

For AIX:

```
LIBPATH = installation_path/lib
```

For all other UNIX/LINUX platforms:

```
LD_LIBRARY_PATH = installation_path/lib
```

Operating System

Required. Variable name is platform dependent. UNIX only.

PATH Environment Variable

PATH is an operating system required variable that includes the directories where Sybase IQ executables are located.

On Windows, the installation program modifies PATH. On UNIX, run the environment source file to include the necessary directories.

On Windows, PATH takes the place of the LIBRARY_PATH variable, so executables and DLLs are located using the PATH variable.

Setting

```
PATH = installation_path
```

Operating System

Required.

SQLCONNECT Environment Variable

SQLCONNECT specifies connection parameters that are used by several of the database administration utilities, such as Interactive SQL, **dbinfo**, and **dbstop**, when connecting to a database server.

The SQLCONNECT environment variable is optional, and is not set by the installation program.

This string is a list of parameter settings, of the form **parameter=value**, delimited by semicolons.

The number sign “#” is an alternative to the equals sign; use it if you are setting the connection parameters string in the SQLCONNECT environment variable. Using “=” inside an environment variable setting is a syntax error. The = sign is allowed only in Windows.

Note: Specifying connection parameters in SQLCONNECT rather than on the command line offers greater security on UNIX systems. It prevents users from being able to display your password with **ps -ef**. This is especially useful if you run Interactive SQL or other utilities in quiet mode. Note that specifying connection parameters in SQLCONNECT rather than on the command line is more secure, but is not entirely secure. Because the password is in plain text, it possible for hostile action to extract it from the environment context. For more information, see *System Administration Guide: Volume 1 > Connection and Communication Parameters > Connection Parameters*.

Settings

```
SQLCONNECT = parameter#value ; ...
```

Operating System

Optional.

SYBASE Environment Variable

SYBASE identifies the location of Sybase applications, such as Open Client and Open Server.

You must set the SYBASE variable before you can install Sybase IQ on UNIX systems. This variable is required for using Sybase Central on UNIX systems.

Setting

```
SYBASE = path
```

Operating System

Required.

\$\$SYBASE_JRE6_32, \$\$SYBASE_JRE6_64, \$\$SYBASE_JRE5_64 **Environment Variables**

This variable specifies the location of the Java Runtime Environment used by the Sybase Central plug-in for Sybase IQ.

For Windows and UNIX, the environment variable is \$\$SYBASE_JRE6_32 or \$\$SYBASE_JRE6_64. For AIX/Linux/IBM the variable is \$\$SYBASE_JRE5_64.

On UNIX, run the SYBASE.csh (C shell) or SYBASE.sh (Bourne or Korn shell) environment source file and locate the correct JRE. JAVA_HOME takes precedence. On Windows, the installation program sets the variable when it installs Open Client Software Developer's Kit.

Setting

```
SYBASE_JRE= "${SYBASE}/shared/jre-6_0"
```

SYBASE_OCS Environment Variable

SYBASE_OCS specifies the home directory for the Open Client product.

This variable is only used on Windows. On Windows, the installation program sets SYBASE_OCS when it installs Open Client/Server Software Developers Kit.

Setting

```
SYBASE_OCS = "OCS-15_4"
```

Operating System

Required.

Registry Entries

On Windows operating systems, Sybase IQ uses several Registry settings.

These settings are made for you by the software, and in general operation, you should not need to access the registry. The settings are provided here if you modify your operating environment.

Warning! Sybase recommends *not* modifying the Registry, as incorrect changes might damage your system.

Current User and Local Machine Settings

Some operating systems, such as Windows, hold two levels of system settings: user settings and local machine settings.

Current user settings are specific to an individual user and are used only when that user is logged on. Local machine settings are global to the machine and are available no matter which user is logged on;. You must have administrator permissions on your machine to make local machine settings.

Sybase IQ permits the use of both current user and local machine settings. For Windows, these settings are held in the HKEY_CURRENT_USER registry and HKEY_LOCAL_MACHINE registry, respectively.

The Sybase IQ installation lets you choose whether the settings it makes are for the current user only or at the local machine level.

If you make settings in both current user and local machine registries, the current user setting takes precedence over the local machine setting.

If you are running a Sybase IQ program as a service on Windows, ensure that the settings are made at the *local machine* level.

Services can continue to run under a special account when you log off a machine, as long as you do not shut the machine down entirely. Services can be made independent of individual accounts and need access to local machine settings.

In general, Sybase recommends using local machine settings.

Registry Structure

On Windows, you can access the Registry directly using the Registry Editor.

Note: Read Only Mode protects your Registry data from accidental changes. To use it, open the Registry Editor, select Edit | Permissions, and then check Read permission.

The Sybase IQ registry entry is held in the HKEY_LOCAL_MACHINE key, in the following location:

```
SOFTWARE
  Sybase
    IQ 15.4
```

Starting the Registry Editor

Start the registry editor to access the Windows registry.

1. Select **Start > Run**.
2. In the Open box, type:

```
regedt32
```

Registry Settings on Installation

The installation program automatically makes these registry settings in the Sybase registry.

- Location – In the Sybase IQ registry, this entry holds the installation directory location. For example:

```
Location:REG_SZ:C:\Program Files\Sybase  
\IQ-15_4
```

The Sybase IQ registry includes other entries for installed applications. The Sybase Central registry holds information about the Sybase Central version and installed plug-ins.

SQL Language Elements

These topics provide detailed descriptions of the language elements and conventions of Sybase IQ SQL.

Keywords

Each SQL statement contains one or more keywords.

SQL is not case-sensitive to keywords, but throughout the Sybase IQ documentation, keywords are indicated in uppercase. For example, in this statement, `SELECT` and `FROM` are keywords:

```
SELECT *
FROM Employees
```

The following statements are equivalent to the one above:

```
Select *
From Employees
select * from Employees
sELECT * fRoM Employees
```

Reserved Words

Some keywords in SQL are also reserved words.

To use a reserved word in a SQL statement as an identifier, you must enclose the word in double quotes. Many, but not all, of the keywords that appear in SQL statements are reserved words. For example, you must use the following syntax to retrieve the contents of a table named `SELECT`.

```
SELECT *
FROM "SELECT"
```

If you are using Embedded SQL, you can use the database library function `sql_needs_quotes` to determine whether a string requires quotation marks. A string requires quotes if it is a reserved word or if it contains a character not ordinarily allowed in an identifier.

This table lists the SQL reserved words in Sybase IQ. Because SQL is not case-sensitive with respect to keywords, each of the words in this table may appear in uppercase, lowercase, or any combination of the two. All strings that differ only in capitalization from these words are reserved words.

Table 1. SQL reserved words

add	all	alter	and
-----	-----	-------	-----

SQL Language Elements

any	as	asc	attach
backup	begin	between	bigint
binary	bit	bottom	break
by	call	capability	cascade
case	cast	char	char_convert
character	check	checkpoint	close
comment	commit	compressed	conflict
connect	constraint	contains	continue
convert	create	cross	cube
current	current_timestamp	current_user	cursor
date	datetimeoffset	dbspace	deallocate
dec	decimal	declare	default
delete	deleting	desc	detach
distinct	do	double	drop
dynamic	else	elseif	encrypted
end	endif	escape	except
exception	exec	execute	existing
exists	externlogin	fetch	first
float	for	force	foreign
forward	from	full	goto
grant	group	having	holdlock
identified	if	in	index
inner	inout	insensitive	insert
inserting	install	instead	int
integer	integrated	intersect	into
is	isolation	join	kerberos
key	lateral	left	like
limit	lock	login	long
match	membership	merge	message
mode	modify	natural	nchar

new	no	noholdlock	not
notify	null	numeric	nvarchar
of	off	on	open
openstring	openxml	option	options
or	order	others	out
outer	over	passthrough	precision
prepare	primary	print	privileges
proc	procedure	publication	raiserror
readtext	real	reference	references
refresh	release	remote	remove
rename	reorganize	resource	restore
restrict	return	revoke	right
rollback	rollup	save	savepoint
scroll	select	sensitive	session
set	setuser	share	smallint
some	spatial	sqlcode	sqlstate
start	stop	subtrans	subtransaction
synchronize	table	temporary	then
time	timestamp	tinyint	to
top	tran	treat	trigger
truncate	tsequal	unbounded	union
unique	uniqueidentifier	unknown	unsigned
update	updating	user	using
validate	values	varbinary	varbit
varchar	variable	varying	view
wait	waitfor	when	where
while	window	with	within
work	writetext	xml	

See also

- *Identifiers* on page 20

- *The quoted_identifier Option* on page 31
- *Subqueries in Search Conditions* on page 35
- *Column Names in Expressions* on page 24

Identifiers

Identifiers are names of objects in the database, such as user IDs, tables, and columns.

Identifiers have a maximum length of 128 bytes. They must be enclosed in double quotes or square brackets if any of these conditions are true:

- The identifier contains spaces.
- The first character of the identifier is not an alphabetic character (as defined below).
- The identifier contains a reserved word.
- The identifier contains characters other than alphabetic characters and digits.

Alphabetic characters include the alphabet, as well as the underscore character (`_`), at sign (`@`), number sign (`#`), and dollar sign (`$`). The database collation sequence dictates which characters are considered alphabetic or digit characters.

You can represent an apostrophe (single quote) inside an identifier by following it with another apostrophe.

If the `QUOTED_IDENTIFIER` database option is set to `OFF`, double quotes are used to delimit SQL strings and cannot be used for identifiers. However, you can always use square brackets to delimit identifiers, regardless of the setting of `QUOTED_IDENTIFIER`.

The default setting for the `QUOTED_IDENTIFIER` option is `OFF` for Open Client and jConnect connections; otherwise the default is `ON`.

Limitations

Identifiers have the following limitations:

- Table names cannot contain double quotes.
- User names cannot contain double quote or semi-colon characters (single quote characters are allowed).
- Database names cannot contain double quote, single quote, and semi-colon characters.
- User names and database names cannot start or end with a space.
- DbSPACE names are always case-insensitive, regardless of the **CREATE DATABASE...CASE IGNORE** or **CASE RESPECT** specification.

For additional restrictions on server and database names, see the `-n start_iq` server option in the *Utility Guide*.

Examples

The following are all valid identifiers.

```
Surname
"Surname"
[Surname]
SomeBigName
"Client Number"
```

See also

- *Reserved Words* on page 17
- *The quoted_identifier Option* on page 31
- *Subqueries in Search Conditions* on page 35
- *Column Names in Expressions* on page 24

Strings

Strings are either literal strings, or expressions with CHAR or VARCHAR data types.

A literal string is any sequence of characters enclosed in apostrophes ('single quotes'). A SQL variable of character data type can hold a string. This is a simple example of a literal string:

An expression with a CHAR data type might be a built-in or user-defined function, or one of the many other kinds of expressions available.

```
'This is a string.'
```

Special Characters in Strings

Represent special characters in strings by escape sequences, as follows:

- To represent an apostrophe inside a string, use two apostrophes in a row. For example:

```
'John' 's database'
```
- To represent a newline character, use a backslash followed by n (\n). For example:

```
'First line:\nSecond line:'
```
- To represent a backslash character, use two backslashes in a row (\\). For example:

```
'c:\\temp'
```
- Hexadecimal escape sequences can be used for any character, printable or not. A hexadecimal escape sequence is a backslash followed by an x followed by two hexadecimal digits (for example, \x6d represents the letter m). For example:

```
'\x00\x01\x02\x03'
```

Compatibility

For compatibility with Adaptive Server® Enterprise, you can set the QUOTED_IDENTIFIER database option to OFF. With this setting, you can also use double quotes to mark the beginning and end of strings. The option is ON by default.

See also

- *Comparison Conditions* on page 33

- *Expressions* on page 22
- *NULL Value* on page 66
- *Search Conditions* on page 32
- *Three-Valued Logic* on page 46
- *SQL Operators* on page 24
- *Subqueries in Search Conditions* on page 35

Expressions

Expressions are formed from several different kinds of elements, such as constants, column names, SQL operators, and subqueries.

Syntax

```
expression:  
case-expression  
| constant  
| [ correlation-name. ] column-name [ java-ref ]  
| - expression  
| expression operator expression  
| ( expression )  
| function-name ( expression, ... )  
| if-expression  
| [ java-package-name. ] java-class-name java-ref  
| ( subquery )  
| variable-name [ java-ref ]
```

Parameters

```
case-expression:  
{ CASE search-condition  
... WHEN expression THEN expression [ , ... ]  
... [ ELSE expression ]  
END  
| CASE  
... WHEN search-condition THEN expression [ , ... ]  
... [ ELSE expression ]  
END }
```

```
constant:  
{ integer | number | 'string' | special-constant | host-variable }
```

```
special-constant:  
{ CURRENT { DATE | TIME | TIMESTAMP | USER }  
| LAST USER  
| NULL  
| SQLCODE  
| SQLSTATE }
```

```
if-expression:  
IF condition  
... THEN expression
```



```
... [ ELSE expression ]
ENDIF
```

```
java-ref:
{ .field-name [ java-ref ]
| >> field-name [ java-ref ]
| .method-name ( [ expression ] [ , ... ] ) [ java-ref ]
| >> method-name ( [ expression ] [ , ... ] ) [ java-ref ] }
```

```
operator:
{ + | - | * | / | || | % }
```

Usage

Anywhere.

Authorization

Must be connected to the database.

Side effects

None.

Compatibility

- The IF condition is not supported in Adaptive Server Enterprise.
- Java expressions are not currently supported in Adaptive Server Enterprise.
- For other differences, see the separate descriptions of each class of expression, in the following sections.

See also

- *Comparison Conditions* on page 33
- *NULL Value* on page 66
- *Search Conditions* on page 32
- *Strings* on page 21
- *Three-Valued Logic* on page 46
- *SQL Operators* on page 24
- *Subqueries in Search Conditions* on page 35
- *Special Values* on page 54
- *CASE Statement Support* on page 709

Constants in Expressions

Constants are numbers or strings.

String constants are enclosed in apostrophes. An apostrophe is represented inside the string by two apostrophes in a row.

Column Names in Expressions

A column name is an identifier preceded by an optional correlation name. A correlation name is usually a table name.

If a column name has characters other than letters, digits, and underscores, the name must be surrounded by quotation marks (“”). For example, the following are valid column names:

```
Employees.Surname
City
"StartDate"
```

See also

- *Subqueries in Search Conditions* on page 35
- *Reserved Words* on page 17
- *Identifiers* on page 20

Subqueries in Expressions

A subquery is a **SELECT** statement enclosed in parentheses. The **SELECT** statement can contain one and only one select list item. When used as an expression, a scalar subquery is allowed to return only zero or one value;

Within the **SELECT** list of the top level **SELECT**, or in the **SET** clause of an **UPDATE** statement, you can use a scalar subquery anywhere that you can use a column name. However, the subquery cannot appear inside a conditional expression (**CASE**, **IF**, **NULLIF**, **ARGN**).

For example, the following statement returns the number of employees in each department, grouped by department name:

```
SELECT DepartmentName, COUNT(*), 'out of',
(SELECT COUNT(*) FROM Employees)
FROM Departments AS D, Employees AS E
WHERE D.DepartmentID = E.DepartmentID
GROUP BY DepartmentName;
```

SQL Operators

These topics describe the arithmetic, string, and bitwise operators available in Sybase IQ.

The normal precedence of operations applies. Expressions in parentheses are evaluated first; then multiplication and division before addition and subtraction. String concatenation occurs after addition and subtraction.

See also

- *Comparison Conditions* on page 33
- *Expressions* on page 22
- *NULL Value* on page 66
- *Search Conditions* on page 32

- *Strings* on page 21
- *Three-Valued Logic* on page 46

Arithmetic Operators

These arithmetic operators are available in Sybase IQ.

Table 2. Arithmetic operators

Operator	Description
expression + expression	Addition. If either expression is the NULL value, the result is the NULL value.
expression - expression	Subtraction. If either expression is the NULL value, the result is the NULL value.
- expression	Negation. If the expression is the NULL value, the result is the NULL value.
expression * expression	Multiplication. If either expression is the NULL value, the result is the NULL value.
expression / expression	Division. If either expression is the NULL value or if the second expression is 0, the result is the NULL value.
expression % expression	Modulo finds the integer remainder after a division involving two whole numbers. For example, 21 % 11 = 10 because 21 divided by 11 equals 1 with a remainder of 10.

String Operators

These string operators are available in Sybase IQ.

Table 3. String operators

Operator	Description
expression expression	String concatenation (two vertical bars). If either string is the NULL value, the string is treated as the empty string for concatenation.
expression + expression	Alternative string concatenation. When using the + concatenation operator, you must ensure the operands are explicitly set to character data types rather than relying on implicit data conversion.

The result data type of a string concatenation operator is a LONG VARCHAR. If you use string concatenation operators in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **LEFT** to the correct data type and size.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. The || operator is the ISO/ANSI SQL string concatenation operator.
- Sybase—The + operator is supported by Adaptive Server Enterprise.

See also

- *REVERSE Function [String]* on page 283

Bitwise Operators

You can use these bitwise operators on all unscaled integer data types, in both Sybase IQ and Adaptive Server Enterprise.

Operator	Description
&	AND
	OR
^	EXCLUSIVE OR
~	NOT

The AND Operator (&)

The AND operator compares 2 bits. If they are both 1, the result is 1.

Bit 1	Bit 2	Bit 1 & Bit 2
0	0	0
0	1	0
1	0	0
1	1	1

Bitwise OR (|)

The OR operator compares 2 bits. If one or the other bit is 1, the result is 1.

Bit 1	Bit 2	Bit 1 Bit 2
0	0	0
0	1	1
1	0	1
1	1	1

EXCLUSIVE OR (^)

The EXCLUSIVE OR operator results in a 1 when either, but not both, of its two operands is 1.

Bit 1	Bit 2	Bit 1 ^Bit 2
0	0	0
0	1	1
1	0	1
1	1	0

NOT (~)

The NOT operator is a unary operator that returns the inverse of its operand.

Bit	~ Bit
1	0
0	1

Join Operators

The Transact-SQL™ outer join operators *= and =* are supported in Sybase IQ, in addition to the ISO/ANSI SQL join syntax using a table expression in the **FROM** clause.

Compatibility

- Modulo—The default value is OFF for new databases.
- String concatenation—When you are using the + concatenation operator in Sybase IQ, ensure the operands are explicitly set to strings rather than relying on implicit data conversion. For example, the following query returns the integer value 579:

```
SELECT 123 + 456
```

whereas the following query returns the character string 123456:

```
SELECT '123' + '456'
```

You can use the **CAST** or **CONVERT** function to explicitly convert data types.

Note: When used with BINARY or VARBINARY data types, the + operator is concatenation, not addition.

The || concatenation operator is not supported by Adaptive Server Enterprise.

Operator Precedence

Follow this recommendation to make the order of operation explicit.

When you are using more than one operator in an expression, Sybase recommends that you use parentheses to make the order of operation explicit, rather than relying on an identical operator precedence between Adaptive Server Enterprise and Sybase IQ.

IF Expressions

The IF expression provides IF-THEN-ELSE SQL expressions.

The syntax of the IF expression is as follows:

```
IF condition
THEN expression1
[ ELSE expression2 ]
ENDIF
```

This expression returns:

- If *condition* is TRUE, the IF expression returns *expression1*.
- If *condition* is FALSE, the IF expression returns *expression2*.
- If *condition* is FALSE, and there is no *expression2*, the IF expression returns NULL.
- If *condition* is NULL, the IF expression returns NULL.

Note: IF statement is different from IF expression.

Do not confuse the syntax of the IF expression with that of the IF statement.

CASE Expressions

The **CASE** expression provides conditional SQL expressions.

You can use case expressions anywhere you can use an expression. The syntax of the **CASE** expression is as follows:

```
CASE expression
WHEN expression THEN expression [, ...]
[ ELSE expression ] END
```

You cannot use a subquery as a value expression in a **CASE** statement.

If the expression following the **CASE** statement is equal to the expression following the **WHEN** statement, then the expression following the **THEN** statement is returned. Otherwise, the expression following the **ELSE** statement is returned, if it exists.

For example, the following code uses a case expression as the second clause in a **SELECT** statement.

```
SELECT ID,
       (CASE name
        WHEN 'Tee Shirt' THEN 'Shirt'
        WHEN 'Sweatshirt' THEN 'Shirt'
```

```

WHEN 'Baseball Cap' THEN 'Hat'
ELSE 'Unknown'
END) as Type
FROM "GROUPO".Products

```

An alternative syntax is:

```

CASE
WHEN search-condition THEN expression [, ...]
[ ELSE expression ] END

```

If the search condition following the **WHEN** statement is satisfied, the expression following the **THEN** statement is returned. Otherwise the expression following the **ELSE** statement is returned, if it exists.

For example, the following statement uses a case expression as the third clause of a **SELECT** statement to associate a string with a search condition.

```

SELECT ID, name,
  (CASE
   WHEN name='Tee Shirt' THEN 'Sale'
   WHEN quantity >= 50 THEN 'Big Sale'
   ELSE 'Regular price'
  END) as Type
FROM "GROUPO".Products

```

See also

- *NULLIF Function [Miscellaneous]* on page 246
- *NULLIF Function for Abbreviated CASE Expressions* on page 29

NULLIF Function for Abbreviated CASE Expressions

The **NULLIF** function provides a way to write some **CASE** statements in short form.

The syntax for **NULLIF** is as follows:

```

NULLIF ( expression-1, expression-2 )

```

NULLIF compares the values of the two expressions. If the first expression equals the second expression, **NULLIF** returns NULL. If the first expression does not equal the second expression, **NULLIF** returns the first expression.

See also

- *CASE Expressions* on page 28
- *NULLIF Function [Miscellaneous]* on page 246

Compatibility of Expressions and Constants

These topics describe the compatibility of expressions and constants between Adaptive Server Enterprise and Sybase IQ.

Compatibility of Expressions

This table describes the compatibility of expressions between Adaptive Server Enterprise and Sybase IQ.

This table is a guide only, and a marking of **Both** may not mean that the expression performs in an identical manner for all purposes under all circumstances. For detailed descriptions, see the Adaptive Server Enterprise documentation and the Sybase IQ documentation on the individual expression.

In this table, **expr** represents an expression, and **op** represents an operator.

Table 4. Compatibility of expressions between Adaptive Server Enterprise and Sybase IQ

Expression	Supported by
constant	Both
column name	Both
variable name	Both
function (expr)	Both
- expr	Both
expr op expr	Both
(expr)	Both
(subquery)	Both
if-expression	Sybase IQ only

Compatibility of Constants

This table describes the compatibility of constants between Adaptive Server Enterprise and Sybase IQ.

This table is a guide only, and a marking of **Both** may not mean that the expression performs in an identical manner for all purposes under all circumstances. For detailed descriptions, see the Adaptive Server Enterprise documentation and the Sybase IQ documentation on the individual expression.

Table 5. Compatibility of constants between Adaptive Server Enterprise and Sybase IQ

Constant	Supported by
integer	Both
number	Both

Constant	Supported by
'string'	Both
special-constant	Both
host-variable	Sybase IQ

Default Interpretation of Delimited Strings

By default, Adaptive Server Enterprise and Sybase IQ give different meanings to delimited strings: that is, strings enclosed in apostrophes (single quotes) and in quotation marks (double quotes).

Sybase IQ employs the SQL92 convention, that strings enclosed in apostrophes are constant expressions, and strings enclosed in quotation marks (double quotes) are delimited identifiers (names for database objects). Adaptive Server Enterprise employs the convention that strings enclosed in quotation marks are constants, whereas delimited identifiers are not allowed by default and are treated as strings.

The quoted_identifier Option

Both Adaptive Server Enterprise and Sybase IQ provide a **quoted_identifier** option that allows the interpretation of delimited strings to be changed. By default, the **quoted_identifier** option is set to OFF in Adaptive Server Enterprise, and to ON in Sybase IQ.

You cannot use SQL reserved words as identifiers if the **quoted_identifier** option is off.

Although the Transact-SQL **SET** statement is not supported for most Adaptive Server Enterprise connection options, **SET** is supported for the **quoted_identifier** option.

The following statement in either Sybase IQ or Adaptive Server Enterprise changes the setting of the **quoted_identifier** option to ON:

```
SET quoted_identifier ON
```

With the **quoted_identifier** option set to ON, Adaptive Server Enterprise allows table, view, and column names to be delimited by quotes. Other object names cannot be delimited in Adaptive Server Enterprise.

The following statement in Sybase IQ or Adaptive Server Enterprise changes the setting of the **quoted_identifier** option to OFF:

```
SET quoted_identifier OFF
```

You can choose to use either the SQL92 or the default Transact-SQL convention in both Adaptive Server Enterprise and Sybase IQ as long as the **quoted_identifier** option is set to the same value in each DBMS.

Examples

If you operate with the **quoted_identifier** option ON (the default Sybase IQ setting), the following statements involving the SQL keyword **user** are valid for both types of DBMS.

```
CREATE TABLE "user" (  
    coll char(5)  
) ;  
INSERT "user" ( coll )  
VALUES ( 'abcde' ) ;
```

If you operate with the **quoted_identifier** option OFF (the default Adaptive Server Enterprise setting), the following statements are valid for both types of DBMS.

```
SELECT *  
FROM Employees  
WHERE Surname = "Chin"
```

See also

- *Reserved Words* on page 17
- *Identifiers* on page 20

Search Conditions

Conditions are used to choose a subset of the rows from a table, or in a control statement such as an **IF** statement to determine control of flow.

SQL conditions do not follow Boolean logic, where conditions are either true or false. In SQL, every condition evaluates as one of TRUE, FALSE, or UNKNOWN. This is called three-valued logic. The result of a comparison is UNKNOWN if either value being compared is the NULL value.

Rows satisfy a search condition if and only if the result of the condition is TRUE. Rows for which the condition is UNKNOWN do not satisfy the search condition.

Subqueries form an important class of expression that is used in many search conditions.

The different types of search conditions are discussed in the following sections.

You specify a search condition for a **WHERE** clause, a **HAVING** clause, a **CHECK** clause, a **JOIN** clause, or an **IF** expression.

Syntax

```
{ expression compare expression  
| expression compare { ANY | SOME | ALL } ( subquery )  
| expression IS [ NOT ] DISTINCT FROM  
| expression IS [ NOT ] NULL expression  
| expression [ NOT ] BETWEEN expression AND expression  
| expression [ NOT ] LIKE expression [ ESCAPE expression ]  
| expression [ NOT ] IN ( { expression | subquery |  
... value-expr1 , value-expr2 [ , value-expr3 ] ... } )  
| column-name [ NOT ] CONTAINS ( ... word1 [ , word2, ] [ , word3 ] ... )  
| CONTAINS ( column-name [ ,... ], contains-query string )  
| EXISTS ( subquery )  
| NOT condition  
| condition AND condition
```

```

| condition OR condition
| ( condition )
| ( condition , estimate )
| condition IS [ NOT ] { TRUE | FALSE | UNKNOWN }

```

Parameters

```

compare:
{ = | > | < | >= | <= | <> | != | !< | !> }

```

Usage

Anywhere

Authorization

Must be connected to the database

Example

For example, the following query retrieves the names and birth years of the oldest employees:

```

SELECT Surname, BirthDate
FROM Employees
WHERE BirthDate <= ALL (SELECT BirthDate FROM Employees);

```

The subqueries that provide comparison values for quantified comparison predicates might retrieve multiple rows but can have only one column.

Side Effects

None

See also

- *Comparison Conditions* on page 33
- *Expressions* on page 22
- *NULL Value* on page 66
- *Strings* on page 21
- *Three-Valued Logic* on page 46
- *SQL Operators* on page 24
- *Subqueries in Search Conditions* on page 35

Comparison Conditions

Comparison conditions in search conditions use a comparison operator.

The syntax for comparison conditions is as follows:

```

expression compare expression

```

where *compare* is a comparison operator. This table lists the comparison operators available in Sybase IQ.

Operator	Description
=	Equal to
>	Greater than
<	Less than
>=	Greater than or equal to
<=	Less than or equal to
!=	Not equal to
<>	Not equal to
!>	Not greater than
!<	Not less than

Example

For example, the following query retrieves the names and birth years of the oldest employees:

```
SELECT Surname, BirthDate
FROM Employees
WHERE Surname <= ALL (SELECT MIN(BirthDate) FROM Employees);
```

The subqueries that provide comparison values for quantified comparison predicates, as in the preceding example, might retrieve multiple rows but can only have one column.

Note: All string comparisons are:

- *Case-sensitive* if the database was created as case respect (the default)
 - *Case-insensitive* if the database was created as case ignore
-

Compatibility

- **Trailing blanks**—Any trailing blanks in character data are ignored for comparison purposes by Adaptive Server Enterprise. The behavior of Sybase IQ when comparing strings is controlled by the Ignore Trailing Blanks in String Comparisons database creation option.
- **Case sensitivity**—By default, Sybase IQ databases, like Adaptive Server Enterprise databases, are created as case-sensitive. Comparisons are carried out with the same attention to case as the database they are operating on. You can control the case sensitivity of Sybase IQ databases when creating the database.

See also

- *Expressions* on page 22
- *NULL Value* on page 66
- *Search Conditions* on page 32

- *Strings* on page 21
- *Three-Valued Logic* on page 46
- *SQL Operators* on page 24
- *Subqueries in Search Conditions* on page 35

Subqueries in Search Conditions

A subquery is a **SELECT** statement enclosed in parentheses. Such a **SELECT** statement must contain one and only one select list item.

A column can be compared to a subquery in a comparison condition (for example, >, <, or !=) as long as the subquery returns no more than one row. If the subquery (which must have one column) returns one row, the value of that row is compared to the expression. If a subquery returns no rows, its value is NULL.

Subqueries that return exactly one column and any number of rows can be used in **IN** conditions, **ANY** conditions, **ALL** conditions, or **EXISTS** conditions. These conditions are discussed in the following sections.

Sybase IQ supports **UNION** only in uncorrelated subquery predicates, not in scalar value subqueries or correlated subquery predicates.

Subqueries cannot be used inside a **CONTAINS** or **LIKE** predicate.

Sybase IQ does not support multiple subqueries in a single **OR** clause. For example, the following query has two subqueries joined by an **OR**:

```
CREATE VARIABLE @ln int;SELECT @ln = 1;select count(*) FROM
lineitemWHERE l_shipdate IN (select l_shipdate FROM lineitem WHERE
l_orderkey IN (2,4,6))OR l_shipdate IN (select l_shipdate FROM
lineitem WHERE l_orderkey IN (1,3,5))OR l_linenummer = @ln;
```

Similar subqueries joined by **AND** and **BETWEEN** are allowed.

See also

- *Column Names in Expressions* on page 24
- *Reserved Words* on page 17
- *Identifiers* on page 20
- *Comparison Conditions* on page 33
- *Expressions* on page 22
- *NULL Value* on page 66
- *Search Conditions* on page 32
- *Strings* on page 21
- *Three-Valued Logic* on page 46

Disjunction of Subquery Predicates

The SQL89 standard allows for several forms of subquery predicates.

Each subquery can appear within the **WHERE** or **HAVING** clause with other predicates, and can be combined using the **AND** or **OR** operators. Sybase IQ supports these subqueries, which can be correlated (contain references to a table that appears in the outer query and cannot be evaluated independently) and uncorrelated (do not contain references to remote tables).

The forms of subquery predicates include:

- Unquantified comparison predicates:

```
<scalar-expression> <comparison-operator> <subquery>
```

The comparison operator is: =, <, >, >=, <, or <=

Unquantified comparison subqueries return exactly one value. If the subquery returns more than one value, an error message appears. This type of query is also called a scalar subquery predicate.

- **IN** predicates:

```
<scalar-expression> [NOT] IN <subquery>
```

The **IN** subquery predicate returns a list of values or a single value. This type is also called a quantified subquery predicate.

- Existence predicates:

```
[NOT] EXISTS <subquery>
```

The **EXISTS** predicate represents the existence of the subquery. The expression **EXISTS** <subquery> evaluates to true only if the subquery result is not empty. The **EXISTS** predicate does not compare results with any column or expressions in the outer query block, and is typically used with correlated subqueries.

- Quantified comparison predicates:

```
<scalar-expression> <comparison-operator> [ANY | ALL] <subquery>
```

A quantified comparison predicate compares one or a collection of values returned from a subquery.

The types of queries you can run include:

- Disjunction of uncorrelated scalar subqueries or **IN** subqueries that cannot be executed vertically within the **WHERE** or **HAVING** clause
- Disjunction of correlated/uncorrelated **EXISTS** subqueries within the **WHERE** or **HAVING** clause
- Disjunction of arbitrary correlated/uncorrelated scalar subqueries, **IN** or **EXISTS** subqueries, or quantified comparison subqueries within the **WHERE** or **HAVING** clause
- Arbitrary uncorrelated/correlated subquery predicates combined with **AND/OR** (conjunct/disjunct) and simple predicates or subquery predicates
- Conjunction/disjunction of subquery predicates on top of a view/derived table

- Disjunction of subquery predicates in **UPDATE**, **DELETE**, and **SELECT INTO** statements

The **SUBQUERY_CACHING_PREFERENCE** option lets experienced DBAs choose which subquery caching method to use. See *Reference: Statements and Options*.

Examples

Disjunction of uncorrelated **EXISTS** and **IN** subqueries:

```
SELECT COUNT(*)
FROM supplier
WHERE s_suppkey IN (SELECT MAX(l_suppkey)
                   FROM lineitem
                   GROUP BY l_linenumbr)
OR EXISTS (SELECT p_brand
          FROM part
          WHERE p_brand = 'Brand#43');
```

Disjunction of uncorrelated **EXISTS** subqueries:

```
SELECT COUNT(*)
FROM supplier
WHERE EXISTS (SELECT l_suppkey
             FROM lineitem
             WHERE l_suppkey = 12345)
OR EXISTS (SELECT p_brand
          FROM part
          WHERE p_brand = 'Brand#43');
```

Disjunction of uncorrelated scalar or **IN** subquery predicates:

```
SELECT COUNT(*)
FROM supplier
WHERE s_acctbal*10 > (SELECT MAX(o_totalprice)
                   FROM orders
                   WHERE o_custkey = 12345)
OR substring(s_name, 1, 6) IN (SELECT c_name
                             FROM Customers
                             WHERE c_nationkey = 10);
```

Disjunction of correlated/uncorrelated quantified comparison subqueries:

```
SELECT COUNT(*)
FROM lineitem
WHERE l_suppkey > ANY (SELECT MAX(s_suppkey)
                    FROM supplier
                    WHERE s_acctbal >100
                    GROUP BY s_nationkey)
OR l_partkey >= ANY (SELECT MAX(p_partkey)
                   FROM part
                   GROUP BY p_mfgr);
```

Disjunction of any correlated subquery predicates:

```
SELECT COUNT(*)
FROM supplier S
WHERE EXISTS (SELECT l_suppkey
            FROM lineitem
```

```

        WHERE l_suppkey = S.s_suppkey)
OR EXISTS (SELECT p_brand FROM part
           WHERE p_brand = 'Brand#43'
           AND p_partkey > S.s_suppkey);

```

Before support for disjunction of subqueries, users were required to write queries in two parts, and then use **UNION** to merge the final results.

The following query illustrates a merged query that gets the same results as the example for disjunction of any correlated subquery predicates. Performance of the merged query is suboptimal because it scans the supplier table twice and then merges the results from each **UNION** to return the final result.

```

SELECT COUNT(*)
FROM (SELECT s_suppkey FROM supplier S
      WHERE EXISTS (SELECT l_suppkey
                    FROM lineitem
                    WHERE l_suppkey = S.s_suppkey)
UNION
SELECT s_suppkey
FROM supplier S
WHERE EXISTS (SELECT p_brand
              FROM part
              WHERE p_brand = 'Brand#43'
              AND p_partkey > S.s_suppkey)) as UD;

```

ALL or ANY Conditions

Use ALL or ANY conditions in subqueries in search conditions.

The syntax for **ALL** conditions is:

```
expression compare ALL ( subquery )
```

where *compare* is a comparison operator.

The syntax for **ANY** conditions is:

```
expression compare ANY ( subquery )
```

where *compare* is a comparison operator.

For example, an **ANY** condition with an equality operator is TRUE if *expression* is equal to any of the values in the result of the subquery, and FALSE if the expression is not NULL and does not equal any of the columns of the subquery:

```
expression = ANY ( subquery )
```

The **ANY** condition is UNKNOWN if *expression* is the NULL value, unless the result of the subquery has no rows, in which case the condition is always FALSE.

You can use the keyword **SOME** instead of **ANY**.

Restrictions

If there is more than one expression on either side of a quantified comparison predicate, an error message is returned. For example:

```
Subquery allowed only one select list item
```

Queries of this type can always be expressed in terms of **IN** subqueries or scalar subqueries using **MIN** and **MAX** set functions.

Compatibility

ANY and **ALL** subqueries are compatible between Adaptive Server Enterprise and Sybase IQ. Only Sybase IQ supports **SOME** as a synonym for **ANY**.

BETWEEN Conditions

Use **BETWEEN** conditions in subqueries to retrieve values within a range.

The syntax for **BETWEEN** conditions is as follows:

```
expr [ NOT ] BETWEEN start-expr AND end-expr
```

The **BETWEEN** condition can evaluate as TRUE, FALSE, or UNKNOWN. Without the **NOT** keyword, the condition evaluates as TRUE if *expr* is between *start-expr* and *end-expr*. The **NOT** keyword reverses the meaning of the condition but leaves UNKNOWN unchanged.

The **BETWEEN** condition is equivalent to a combination of two inequalities:

```
expr >= start-expr AND expr <= end-expr
```

A **BETWEEN** predicate is of the form “A between B and C.” Either “B” or “C” or both “B” and “C” can be subqueries. “A” must be a value expression or column.

Compatibility

The **BETWEEN** condition is compatible between Sybase IQ and Adaptive Server Enterprise.

LIKE Conditions

Use **LIKE** conditions in subqueries to use wildcards in the WHERE clause to perform pattern matching.

The syntax for **LIKE** conditions is:

```
expression [ NOT ] LIKE pattern [ ESCAPE escape-expr ]
```

The **LIKE** condition can evaluate as TRUE, FALSE, or UNKNOWN. You can use **LIKE** only on string data.

You cannot use subqueries inside a **LIKE** predicate.

LIKE predicates that start with characters other than wildcard characters may execute faster if an **HG** or **LF** index is available.

Certain **LIKE** predicates execute faster, if a WD index is available.

Without the **NOT** keyword, the condition evaluates as **TRUE** if *expression* matches the *pattern*. If either *expression* or *pattern* is the **NULL** value, this condition is **UNKNOWN**. The **NOT** keyword reverses the meaning of the condition but leaves **UNKNOWN** unchanged.

The pattern might contain any number of wildcard characters. The wildcard characters are:

Wildcard	Matches
_ (underscore)	Any one character
% (percent)	Any string of zero or more characters
[]	Any single character in the specified range or set
[^]	Any single character not in the specified range or set

All other characters must match exactly.

For example, the search condition:

```
name LIKE 'a%b_'
```

is **TRUE** for any row where name starts with the letter a and has the letter b as its second-to-last character.

If you specify an *escape-expr*, it must evaluate to a single character. The character can precede a percent, an underscore, a left square bracket, or another escape character in the *pattern* to prevent the special character from having its special meaning. When escaped in this manner, a percent matches a percent, and an underscore matches an underscore.

Patterns Supported

All patterns of 126 characters or less are supported.

Some patterns between 127 and 254 characters are supported, but only under certain circumstances. See the following subsections for examples.

All patterns 255 characters or greater are not supported.

Patterns Between 127 and 254 Characters: Example 1

Under specific circumstances where adjacent constant characters exist in your pattern, patterns of length between 127 and 254 characters are supported. Each constant character in the string pattern requires two bytes, even if the character is a single-byte character. The string pattern in the **LIKE** predicate must be less than 256 bytes (or 255/2 characters) or else the following error appears:

```
There was an error reading the results of the SQL statement.
The displayed results may be incorrect or incomplete.
Cannot compile Like pattern: either bad pattern or pattern too long.
```

Sybase IQ collapses adjacent constant characters into a single character. For example, consider the following **LIKE** predicate with a string length of 130 characters:

Searching for One of a Set of Characters

You can specify a set of characters to look for by listing the characters inside square brackets. For example, the following condition finds the strings *smith* and *smyth*:

```
LIKE 'sm[iy]th'
```

Searching for One of a Range of Characters

Specify a range of characters to look for by listing the ends of the range inside square brackets, separated by a hyphen. For example, the following condition finds the strings *bough* and *rough*, but not *tough*:

```
LIKE '[a-r]ough'
```

The range of characters [a-z] is interpreted as “greater than or equal to a, and less than or equal to z,” where the greater than and less than operations are carried out within the collation of the database. For information on ordering of characters within a collation, see *System Administration Guide: Volume 1* > International Languages and Character Sets.

The lower end of the range must precede the higher end of the range. For example, a **LIKE** condition containing the expression [z-a] returns no rows, because no character matches the [z-a] range.

Unless the database is created as case-sensitive, the range of characters is case-insensitive. For example, the following condition finds the strings *Bough*, *rough*, and *TOUGH*:

```
LIKE '[a-z]ough'
```

If the database is created as a case-sensitive database, the search condition is case-sensitive also.

Combining Searches for Ranges and Sets

You can combine ranges and sets within square brackets. For example, the following condition finds the strings *bough*, *rough*, and *tough*:

```
LIKE '[a-rt]ough'
```

The bracket [*a-mpqs-z*] is interpreted as “exactly one character that is either in the range *a* to *m* inclusive, or is *p*, or is *q*, or is in the range *s* to *z* inclusive.”

Searching for One Character Not in a Range

Use the caret character (^) to specify a range of characters that is excluded from a search. For example, the following condition finds the string *tough*, but not the strings *rough*, or *bough*:

```
LIKE '[^a-r]ough'
```

The caret negates the entire contents of the brackets. For example, the bracket [*^a-mpqs-z*] is interpreted as “exactly one character that is not in the range *a* to *m* inclusive, is not *p*, is not *q*, and is not in the range *s* to *z* inclusive.”

Special Cases of Ranges and Sets

Any single character in square brackets indicates that character. For example, `[a]` matches just the character `a`. `[^]` matches just the caret character, `[%]` matches only the percent character (the percent character does not act as a wildcard character in this context), and `[_]` matches just the underscore character. Also, `[/]` matches only the character `/`.

Other special cases are:

- The expression `[a-]` matches either of the characters `a` or `-`.
- The expression `[]` is never matched and always returns no rows.
- The expressions `[or [abp-q` are ill-formed expressions, and give syntax errors.
- You cannot use wildcard characters inside square brackets. The expression `[a%b]` finds one of `a`, `%`, or `b`.
- You cannot use the caret character to negate ranges except as the first character in the bracket. The expression `[a^b]` finds one of `a`, `^`, or `b`.

Compatibility

The **ESCAPE** clause is supported by Sybase IQ only.

Note: For information on support of the **LIKE** predicate with large object data and variables, see *Unstructured Data Analytics in Sybase IQ > Unstructured Data Queries*.

Users must be specifically licensed to use the large object data types `LONG BINARY` and `LONG VARCHAR`. For details on the Unstructured Data Analytics Option, see *Unstructured Data Analytics in Sybase IQ*.

See also

- *PATINDEX Function [String]* on page 250
- *LOCATE Function [String]* on page 225

IN conditions

Use **IN** conditions in subqueries to reduce the need to use multiple **OR** conditions:

The syntax for **IN** conditions is:

```
{ expression [ NOT ] IN ( subquery )
  | expression [ NOT ] IN ( expression )
  | expression [ NOT ] IN ( value-expr1 , value-expr2
  | , value-expr3 ] ... ) }
```

Without the **NOT** keyword, the **IN** condition is **TRUE** if *expression* equals any of the listed values, **UNKNOWN** if *expression* is the **NULL** value, and **FALSE** otherwise. The **NOT** keyword reverses the meaning of the condition but leaves **UNKNOWN** unchanged.

The maximum number of values allowed in an **IN** condition list is 250,000.

Compatibility

IN conditions are compatible between Adaptive Server Enterprise and Sybase IQ.

CONTAINS Conditions

Use **CONTAINS** conditions in subqueries to define text-matching.

The syntax for **CONTAINS** conditions for a column with a **WD** index is as follows:

```
{ column-name [ NOT ] CONTAINS ( ( word1 [ , word2 ] [ , word3 ] ... )
```

The *column-name* must be a CHAR, VARCHAR, or LONG VARCHAR (CLOB) column in a base table, and must have a **WD** index. The *word1*, *word2* and *word3* expressions must be string constants no longer than 255 bytes, each containing exactly one word. The length of that word cannot exceed the maximum permitted word length of the word index of the column.

Without the **NOT** keyword, the **CONTAINS** condition is TRUE if *column-name* contains each of the words, UNKNOWN if *column-name* is the NULL value, and FALSE otherwise. The **NOT** keyword reverses these values but leaves UNKNOWN unchanged.

For example, this search condition:

```
varchar_col CONTAINS ('cat', 'mat')
```

is TRUE if the value of *varchar_col* is The cat is on the mat. If the value of *varchar_col* is The cat chased the mouse, this condition is FALSE.

When Sybase IQ executes a statement containing both **LIKE** and **CONTAINS**, the **CONTAINS** condition takes precedence.

Avoid using the **CONTAINS** predicate in a view that has a user-defined function, because the **CONTAINS** criteria are ignored. Use the **LIKE** predicate with wildcards instead, or issue the query outside of a view.

For information on using **CONTAINS** conditions with **TEXT** indexes, see *Unstructured Data Analytics in Sybase IQ*.

EXISTS Conditions

An **EXISTS** condition is met if the subquery returns at least one row.

The syntax for **EXISTS** conditions is as follows:

```
EXISTS( subquery )
```

The **EXISTS** condition is TRUE if the subquery result contains at least one row, and FALSE if the subquery result does not contain any rows. The **EXISTS** condition cannot be UNKNOWN.

Compatibility

The **EXISTS** condition is compatible between Adaptive Server Enterprise and Sybase IQ.

IS DISTINCT FROM and IS NOT DISTINCT FROM Search Conditions

The **IS DISTINCT FROM** and **IS NOT DISTINCT FROM** search conditions are sargable and evaluate to TRUE or FALSE.

The syntax for **IS [NOT] DISTINCT FROM** conditions is:

```
expression1 IS [ NOT ] DISTINCT FROM expression2
```

Remarks

The **IS NOT DISTINCT FROM** search condition evaluates to TRUE if *expression1* is equal to *expression2*, or if both expressions are NULL. This is equivalent to a combination of two search conditions, as follows:

```
expression1 = expression2 OR ( expression1 IS NULL AND expression2 IS NULL )
```

The **IS DISTINCT FROM** syntax reverses the meaning. That is, **IS DISTINCT FROM** evaluates to TRUE if *expression1* is not equal to *expression2*, and at least one of the expressions is not NULL. This is equivalent to the following:

```
NOT( expression1 = expression2 OR ( expression1 IS NULL AND expression2 IS NULL ) )
```

Standards and Compatibility

- **SQL/2008** – The **IS [NOT] DISTINCT FROM** predicate is defined in SQL/2008 standard. The **IS DISTINCT FROM** predicate is *Feature T151, "DISTINCT predicate"*, of the SQL/2008 standard. The **IS NOT DISTINCT FROM** predicate is *Feature T152, "DISTINCT predicate with negation"*, of the SQL/2008 standard.

IS NULL Conditions

Use **IS NULL** conditions in subqueries to NULL values represent missing unknown data.

The syntax for **IS NULL** conditions is:

```
expression IS [ NOT ] NULL
```

Without the **NOT** keyword, the **IS NULL** condition is TRUE if the expression is the NULL value, and FALSE otherwise. The **NOT** keyword reverses the meaning of the condition.

Compatibility

The **IS NULL** condition is compatible between Adaptive Server Enterprise and Sybase IQ.

Conditions with Logical Operators

Combine search conditions in subqueries using **AND**, **OR**, and **NOT**.

Conditions are combined using **AND** as follows:

```
condition1 AND condition2
```

If both conditions are TRUE, the combined condition is TRUE. If either condition is FALSE, the combined condition is FALSE. If otherwise, the combined condition is UNKNOWN.

Conditions are combined using **OR** as follows:

```
condition1 OR condition2
```

If both conditions are TRUE, the combined condition is TRUE. If either condition is FALSE, the combined condition is FALSE. If otherwise, the combined condition is UNKNOWN.

There is no guaranteed order as to which condition, *condition1* or *condition2*, is evaluated first.

Compatibility

The **AND** and **OR** operators are compatible between Sybase IQ and Adaptive Server Enterprise.

NOT Conditions

The NOT condition can be either TRUE, FALSE, or UNKNOWN.

The syntax for **NOT** conditions is:

```
NOT condition1
```

The **NOT** condition is TRUE if *condition1* is FALSE, FALSE if *condition1* is TRUE, and UNKNOWN if *condition1* is UNKNOWN.

Truth Value Conditions

The truth value of a condition is either TRUE or FALSE.

The syntax for truth value conditions is:

```
IS [ NOT ] truth-value
```

Without the **NOT** keyword, the condition is TRUE if the *condition* evaluates to the supplied *truth-value*, which must be one of TRUE, FALSE, or UNKNOWN. Otherwise, the value is FALSE. The **NOT** keyword reverses the meaning of the condition but leaves UNKNOWN unchanged.

Compatibility

Truth-valued conditions are supported by Sybase IQ only.

Three-Valued Logic

The AND, OR, NOT, and IS logical operators of SQL work in three-valued logic.

These tables show the three-valued logic.

AND Operator

AND	TRUE	FALSE	UNKNOWN
TRUE	TRUE	FALSE	UNKNOWN
FALSE	FALSE	FALSE	FALSE
UNKNOWN	UNKNOWN	FALSE	UNKNOWN

OR Operator

OR	TRUE	FALSE	UNKNOWN
TRUE	TRUE	TRUE	TRUE
FALSE	TRUE	FALSE	UNKNOWN
UNKNOWN	TRUE	UNKNOWN	UNKNOWN

NOT Operator

TRUE	FALSE	UNKNOWN
FALSE	TRUE	UNKNOWN

IS Operator

IS	TRUE	FALSE	UNKNOWN
TRUE	TRUE	FALSE	FALSE
FALSE	FALSE	TRUE	FALSE
UNKNOWN	FALSE	FALSE	TRUE

See also

- *Comparison Conditions* on page 33
- *Expressions* on page 22
- *NULL Value* on page 66
- *Search Conditions* on page 32
- *Strings* on page 21
- *SQL Operators* on page 24
- *Subqueries in Search Conditions* on page 35

User-Supplied Condition Hints

The selectivity of a condition is the fraction of the table's rows that satisfy that condition.

The Sybase IQ query optimizer uses information from available indexes to select an appropriate strategy for executing a query. For each condition in the query, the optimizer decides whether the condition can be executed using indexes, and if so, the optimizer chooses which index and in what order with respect to the other conditions on that table. The most important factor in these decisions is the selectivity of the condition; that is, the fraction of the table's rows that satisfy that condition.

The optimizer normally decides without user intervention, and it generally makes optimal decisions. In some situations, however, the optimizer might not be able to accurately determine the selectivity of a condition before it has been executed. These situations normally occur only where either the condition is on a column with no appropriate index available, or where the condition involves some arithmetic or function expression and is, therefore, too complex for the optimizer to accurately estimate.

If you have a query that is run frequently, then you may want to experiment to see whether you can improve the performance of that query by supplying the optimizer with additional information to aid it in selecting the optimal execution strategy.

User-Supplied Condition Selectivity

The simplest form of condition hint is to supply a selectivity value that will be used instead of the value the optimizer would have computed.

Selectivity hints are supplied within the text of the query by wrapping the condition within parentheses. Then within the parentheses, after the condition, you add a comma and a numeric value to be used as the selectivity.

This selectivity value is expressed as a percentage of the table's rows, which satisfy the condition. Possible numeric values for selectivity thus range from 100.0 to 0.0.

Note: In query plans, selectivity is expressed as a fraction instead of as a percentage; so a user-supplied selectivity of 35.5 appears in that query's plan as a selectivity of 0.355000.

Examples

- The following query provides an estimate that one and one half percent of the `ship_date` values are earlier than 1994/06/30:

```
SELECT ShipDate
FROM SalesOrderItems
WHERE ( ShipDate < '2001/06/30', 1.5 )
ORDER BY ShipDate DESC
```

- The following query estimates that half a percent of the rows satisfy the condition:

```
SELECT *
FROM Customers c, SalesOrders o
```

```
WHERE (o.SalesRepresentative > 1000.0, 0.5)
      AND c.ID = o.customerID
```

Fractional percentages enable more precise user estimates to be specified and can be particularly important for large tables.

Compatibility

SQL Anywhere Studio® supports user-supplied selectivity estimates.

Adaptive Server Enterprise does not support user-supplied selectivity estimates.

See also

- *Selectivity Hints* on page 49

User-Supplied Condition Hint Strings

You can supply additional hint information to the optimizer through a condition hint string.

These per-condition hint strings let users specify additional execution preferences for a condition, which the optimizer follows, if possible. These preferences include which index to use for the condition, the selectivity of the condition, the phase of execution when the condition is executed, and the usefulness of the condition, which affects its ordering among the set of conditions executed within one phase of execution.

Condition hint strings, like the user-supplied selectivity estimates, are supplied within the text of the query by wrapping the condition within parentheses. Then within the parentheses and after the condition, you add a comma and a supply a quoted string containing the desired hints. Within that quoted string each hint appears as a hint type identifier, followed by a colon and the value for that hint type. Multiple hints within the same hint string are separated from each other by a comma, and multiple hints can appear in any order. White space is allowed between any of two elements within a hint string.

Selectivity Hints

The first hint type that can appear within a hint string is a selectivity hint. A selectivity hint is identified by a hint type identifier of either “S” or “s”.

Like user-supplied selectivity estimates, the selectivity value is always expressed as a percentage of the table’s rows, which satisfy the condition.

Example

The following example is exactly equivalent to the second user-supplied condition selectivity example.

```
SELECT *
FROM Customers c, SalesOrders o
WHERE (o.SalesRepresentative > 1000.0, 's: 0.5)
      AND c.ID = o.CustomerID
```

See also

- *User-Supplied Condition Selectivity* on page 48

Index Preference Hints

The second supported hint type is an index preference hint, which is identified by a hint type identifier of either “I” or “i”.

The value for an index preference hint can be any integer between -10 and 10. The meaning of each positive integer value is to prefer a specific index type, while negative values indicate that the specific index type is to be avoided.

The effect of an index preference hint is the same as that of the **INDEX_PREFERENCE** option, except that the preference applies only to the condition it is associated with rather than all conditions within the query. An index preference can only affect the execution of a condition if the specified index type exists on that column and that index type is valid for use when evaluating the associated condition; not all index types are valid for use with all conditions. See *Reference: Statements and Options*.

Example

The following example specifies a 3 percent selectivity and indicates that, if possible, the condition should be evaluated using an HG index:

```
SELECT *
FROM Customers c, SalesOrders o
WHERE (o.SalesRepresentative > 1000.0, 'S:3.00, I:+2')
      AND c.ID = o.CustomerID
```

The next example specifies a 37.5 percent selectivity and indicates that if possible the condition should not be evaluated using an HG index:

```
SELECT *
FROM Customers c, SalesOrders o
WHERE (o.SalesRepresentative > 1000.0, 'i:-2, s:37.500')
      AND c.ID = o.CustomerID
```

INDEX_PREFERENCE Option

Controls the choice of indexes to use for queries.

Allowed Values

-10 to 10

Default

0

Scope

DBA permissions are not required to set this option. Can be set temporary for an individual connection or for the PUBLIC group. Takes effect immediately.

Description

The Sybase IQ optimizer normally chooses the best index available to process local **WHERE** clause predicates and other operations that can be done within an IQ index.

INDEX_PREFERENCE is used to override the optimizer choice for testing purposes; under most circumstances, it should not be changed.

Table 6. INDEX_PREFERENCE Valid Values

Value	Action
0	Let the optimizer choose
1	Prefer LF indexes
2	Prefer HG indexes
3	Prefer HNG indexes
4	Prefer CMP indexes
5	Prefer the default index
6	Prefer WD indexes
8	Prefer DATE indexes
9	Prefer TIME indexes
10	Prefer DTTM indexes
-1	Avoid LF indexes
-2	Avoid HG indexes
-3	Avoid HNG indexes
-4	Avoid CMP indexes
-5	Avoid the default index
-6	Avoid WD indexes
-8	Avoid DATE indexes
-9	Avoid TIME indexes
-10	Avoid DTTM indexes

Execution Phase Hints

The third supported hint type is the execution phase hint, which is identified with a hint type identifier of either “E” or “e”.

Within the Sybase IQ query engine, there are distinct phases of execution where conditions can be evaluated: invariant, delayed, bound, and horizontal.

By default, the optimizer chooses to evaluate each condition within the earliest phase of execution where all the information needed to evaluate that condition is available. Every condition, therefore, has a default execution phase where it is evaluated.

Because no condition can be evaluated before the information it needs is available, the execution phase hint can only be used to delay the execution of a condition to a phase after its default phase. It cannot be used to force a condition to be evaluated within any phase earlier than its default phase.

The four phases of condition execution from earliest to latest are as follows:

Invariant— A condition that refers to only one column (or two columns from the same table) and that can be evaluated using an index is generally referred to as a simple invariant condition. Simple invariant conditions are normally evaluated early within the optimization process. This means that the number of rows satisfying all of those invariant conditions is available to guide the optimizer’s decisions on the best join order and join algorithms to use. Because this is the earliest phase of execution, a user can never force a condition into this phase, but conditions can be forced out of this phase into later phases.

Delayed—Some conditions cannot be evaluated until some other part of a query has been executed. These delayed conditions are evaluated once when the query node to which they are attached is first fetched. These conditions fall into two categories, uncorrelated subquery conditions and IN or PROBABLY_IN pushdown join conditions created by the optimizer.

Bound—Some conditions must be evaluated multiple times. These conditions generally fall into two categories: conditions containing outer references within a correlated subquery, and pushdown equality join conditions created by the optimizer. The outer reference conditions, for example, are reevaluated each time the outer reference value changes during the query's execution.

Horizontal—Some conditions, such as those which contain more than two columns from a table, must be evaluated one row at a time, rather than by using an index.

An execution phase hint accepts a value that identifies in which execution phase the user wants the condition to be evaluated. Each value is a case-insensitive single character:

- D – Delayed
- B – Bound
- H – Horizontal

Example

The following example shows a condition hint string which indicates that the condition should be moved into the “Delayed” phase of execution, and it indicates that if possible the condition should be evaluated using an LF index.:

```
SELECT *
FROM Customers c, SalesOrders o
WHERE (o.SalesRepresentative > 10000.0, 'E:D, I:1')
      AND c.id = o.CustomerID
```

Usefulness Hints

The final supported hint type is the usefulness hint, which is identified by a hint type identifier of either “U” or “u”.

The value for a usefulness hint can be any numeric value between 0.0 and 10.0. Within the optimizer a usefulness value is computed for every condition, and the usefulness value is then used to determine the order of evaluation among the set of conditions to be evaluated within the same phase of execution. The higher the usefulness value, the earlier it appears in the order of evaluation. Supplying a usefulness hint lets users place a condition at a particular point within the order of evaluation, but it cannot change the execution phase within which the condition is evaluated.

Example

The following example shows a condition hint string which indicates that the condition should be moved into the “Delayed” phase of execution, and that its usefulness should be set to 3.25 within that “Delayed” phase.

```
SELECT *
FROM Customers c, SalesOrders o
WHERE (co.SalesRepresentative > 10000.0, 'U: 3.25, E: D')
AND c.id = o.CustomerID
```

Compatibility

SQL Anywhere Studio does not support user-supplied condition hint strings.

Adaptive Server Enterprise does not support user-supplied condition hint strings.

User-Supplied Hints on Join Equality Conditions

Users can specify a join algorithm preference that does not affect every join in the query.

Simple equality join predicates can be tagged with a predicate hint that allows a join preference to be specified for just that one join. If the same join has more than one join condition with a local join preference, and if those hints are not the same value, then all local preferences are ignored for that join. Local join preferences do not affect the join order chosen by the optimizer.

The following example requests a hash join:

```
AND (T.X = 10 * R.x, 'J:4')
```

Guidelines for Usage of User-Supplied Condition Hints

Condition hints are generally appropriate only within frequently run queries.

Only advanced users should experiment with condition hints. The optimizer generally makes optimal decisions, except where it cannot infer accurate information about a condition from the available indexes.

The optimizer often rewrites or simplifies the original conditions, and it also infers new conditions from the original conditions. Condition hints are not carried through new to conditions inferred by the optimizer, nor are they carried through to simplified conditions.

Special Values

Special values can be used in expressions, and as column defaults when creating tables.

See also

- *Expressions* on page 22

CURRENT DATABASE Special Value

CURRENT DATABASE returns the name of the current database.

Data Type
STRING

CURRENT DATE Special Value

CURRENT DATE returns the current year, month and day.

Data Type
DATE

See also

- *TIMESTAMP Special Value* on page 57
- *CURRENT TIMESTAMP Special Value* on page 55
- *CURRENT TIME Special Value* on page 55
- *Date and Time Data Types* on page 84
- *Retrieve Dates and Times* on page 87

CURRENT PUBLISHER Special Value

CURRENT PUBLISHER returns a string that contains the publisher user ID of the database for SQL Remote replications.

Data Type

STRING

CURRENT PUBLISHER can be used as a default value in columns with character data types.

CURRENT TIME Special Value

CURRENT TIME returns the current hour, minute, second, and fraction of a second.

Data Type

TIME

Description

The fraction of a second is stored to 6 decimal places, but the accuracy of the current time is limited by the accuracy of the system clock.

See also

- *TIMESTAMP Special Value* on page 57
- *CURRENT TIMESTAMP Special Value* on page 55
- *CURRENT DATE Special Value* on page 54
- *Date and Time Data Types* on page 84
- *Retrieve Dates and Times* on page 87

CURRENT TIMESTAMP Special Value

Combines **CURRENT DATE** and **CURRENT TIME** to form a **TIMESTAMP** value containing the year, month, day, hour, minute, second and fraction of a second.

As with **CURRENT TIME**, the accuracy of the fraction of a second is limited by the system clock.

CURRENT TIMESTAMP defaults to 3 digits.

Data type

TIMESTAMP

See also

- *TIMESTAMP Special Value* on page 57
- *CURRENT TIME Special Value* on page 55
- *CURRENT DATE Special Value* on page 54

- *Date and Time Data Types* on page 84
- *Retrieve Dates and Times* on page 87
- *CURRENT USER Special Value* on page 56
- *LAST USER Special Value* on page 56
- *USER Special Value* on page 58

CURRENT USER Special Value

CURRENT USER returns a string that contains the user ID of the current connection.

On **UPDATE**, columns with a default value of **CURRENT USER** are not changed.

Data Type

STRING

CURRENT USER can be used as a default value in columns with character data types.

See also

- *CURRENT TIMESTAMP Special Value* on page 55
- *LAST USER Special Value* on page 56
- *USER Special Value* on page 58

LAST USER Special Value

LAST USER returns the name of the user who last modified the row.

On **INSERT** and **LOAD**, this constant has the same effect as **CURRENT USER**. On **UPDATE**, if a column with a default value of **LAST USER** is not explicitly modified, it is changed to the name of the current user.

When combined with the **DEFAULT TIMESTAMP**, a default value of **LAST USER** can be used to record (in separate columns) both the user and the date and time a row was last changed.

Data Type

STRING

LAST USER can be used as a default value in columns with character data types.

See also

- *CURRENT USER Special Value* on page 56
- *CURRENT TIMESTAMP Special Value* on page 55
- *USER Special Value* on page 58

SQLCODE Special Value

SQLCODE returns the current **SQLCODE** value.

The **SQLCODE** value is set after each statement. You can check the **SQLCODE** to see whether or not the statement succeeded.

DATA TYPE
STRING

SQLSTATE Special Value

SQLSTATE returns the current **SQLSTATE** value.

The **SQLSTATE** value is set after each statement. You can check the **SQLSTATE** to see whether or not the statement succeeded.

Data Type
STRING

TIMESTAMP Special Value

TIMESTAMP indicates when each row in the table was last modified.

When a column is declared with **DEFAULT TIMESTAMP**, a default value is provided for insert and load operations. The value is updated with the current date and time whenever the row is updated.

On **INSERT** and **LOAD**, **DEFAULT TIMESTAMP** has the same effect as **CURRENT TIMESTAMP**. On **UPDATE**, if a column with a default value of **TIMESTAMP** is not explicitly modified, the value of the column is changed to the current date and time.

Note: Sybase IQ does not support **DEFAULT** values of **UTC TIMESTAMP** or **CURRENT UTC TIMESTAMP**, nor does IQ support the database option **DEFAULT_TIMESTAMP_INCREMENT**. Sybase IQ generates an error every time an attempt is made to insert or update the **DEFAULT** value of a column of type **UTC TIMESTAMP** or **CURRENT UTC TIMESTAMP**.

Data Type
TIMESTAMP

See also

- *CURRENT TIMESTAMP Special Value* on page 55
- *CURRENT TIME Special Value* on page 55
- *CURRENT DATE Special Value* on page 54
- *Date and Time Data Types* on page 84
- *Retrieve Dates and Times* on page 87

USER Special Value

USER returns a string that contains the user ID of the current connection.

On **UPDATE**, columns with a default value of **USER** are not changed.

Data Type

STRING

USER can be used as a default value in columns with character data types.

See also

- *CURRENT USER Special Value* on page 56
- *CURRENT TIMESTAMP Special Value* on page 55
- *LAST USER Special Value* on page 56

Variables

Sybase IQ supports local variables, connection-level variables, and global variables.

All global variables have names beginning with two @ signs. For example, the global variable *@@version* has a value that is the current version number of the database server. Users cannot define global variables.

Local Variables

Local variables are declared by the user, and can be used in procedures or in batches of SQL statements to hold information.

Local variables are declared using the **DECLARE** statement, which can be used only within a compound statement (that is, bracketed by the **BEGIN** and **END** keywords). The variable is initially set as NULL. You can set the value of the variable using the **SET** statement, or you can assign the value using a **SELECT** statement with an **INTO** clause.

The syntax of the **DECLARE** statement is as follows:

```
DECLARE variable-name data-type
```

You can pass local variables as arguments to procedures, as long as the procedure is called from within the compound statement.

Examples

- The following batch illustrates the use of local variables:

```
BEGIN
  DECLARE local_var INT ;
  SET local_var = 10 ;
```

```
MESSAGE 'local_var = ', local_var ;
END
```

Running this batch from ISQL displays this message on the server window:

```
local_var = 10
```

- The variable `local_var` does not exist outside the compound statement in which it is declared. The following batch is invalid, and displays a column not found error:

```
-- This batch is invalid.
BEGIN
    DECLARE local_var INT ;
    SET local_var = 10 ;
    MESSAGE 'local_var = ', local_var ;
END;
MESSAGE 'local_var = ', local_var ;
```

- The following example illustrates the use of **SELECT** with an **INTO** clause to set the value of a local variable:

```
BEGIN
    DECLARE local_var INT ;
    SELECT 10 INTO local_var ;
    MESSAGE 'local_var = ', local_var ;
END
```

Running this batch from ISQL displays this message on the server window:

```
local_var = 10
```

Compatibility

Names—Adaptive Server Enterprise and Sybase IQ both support local variables. In Adaptive Server Enterprise, all variables must be prefixed with an @ sign. In Sybase IQ, the @ prefix is optional. To write compatible SQL, ensure all your variables have the @ prefix.

Scope—The scope of local variables differs between Sybase IQ and Adaptive Server Enterprise. Sybase IQ supports the use of the **DECLARE** statement to declare local variables within a batch. However, if the **DECLARE** is executed within a compound statement, the scope is limited to the compound statement.

Declaration—Only one variable can be declared for each **DECLARE** statement in Sybase IQ. In Adaptive Server Enterprise, more than one variable can be declared in a single statement.

Connection-Level Variables

Connection-level variables are declared by the user, and can be used in procedures or in batches of SQL statements to hold information.

Connection-level variables are declared with the **CREATE VARIABLE** statement. The **CREATE VARIABLE** statement can be used anywhere except inside compound statements. Connection-level variables can be passed as parameters to procedures.

The syntax for **CREATE VARIABLE** is:

```
CREATE VARIABLE variable-name data-type
```

SQL Language Elements

When a variable is created, it is initially set to NULL. You can set the value of connection-level variables in the same way as local variables, using the **SET** statement or using a **SELECT** statement with an **INTO** clause.

Connection-level variables exist until the connection is terminated, or until you explicitly drop the variable using the **DROP VARIABLE** statement. The following statement drops the variable *con_var*:

```
DROP VARIABLE con_var
```

Example

- The following batch of SQL statements illustrates the use of connection-level variables.

```
CREATE VARIABLE con_var INT;  
SET con_var = 10;  
MESSAGE 'con_var = ', con_var;
```

Running this batch from ISQL displays this message on the server window:

```
con_var = 10
```

Compatibility

Adaptive Server Enterprise does not support connection-level variables.

Global Variables

Global variables are system-supplied variables that provide system-supplied values.

Sybase IQ sets the values of global variables. For example, the global variable *@@version* has a value that is the current version number of the database server.

Global variables are distinguished from local and connection-level variables by two @ signs preceding their names. For example, *@@error* is a global variable. Users cannot create global variables, and cannot update the value of global variables directly.

Some global variables, such as *@@spid*, hold connection-specific information and therefore have connection-specific values. Other variables, such as *@@connections*, have values that are common to all connections.

Global Variable and Special Constants

The special constants such as **CURRENT DATE**, **CURRENT TIME**, **USER**, **SQLSTATE**, and so on are similar to global variables.

The following statement retrieves the value of the version global variable:

```
SELECT @@version
```

In procedures, global variables can be selected into a variable list. The following procedure returns the server version number in the *ver* parameter.

```
CREATE PROCEDURE VersionProc ( OUT ver  
                               VARCHAR ( 100) )  
BEGIN
```

```

SELECT @@version
INTO ver;
END

```

In Embedded SQL, global variables can be selected into a host variable list.

List of Global Variables

This table lists the global variables available in Sybase IQ.

Table 7. Sybase IQ global variables

Variable name	Meaning
<i>@@error</i>	Commonly used to check the error status (succeeded or failed) of the most recently executed statement. Contains 0 if the previous transaction succeeded; otherwise, contains the last error number generated by the system. A statement such as if <i>@@error</i> != 0 return causes an exit if an error occurs. Every SQL statement resets <i>@@error</i> , so the status check must immediately follow the statement whose success is in question.
<i>@@fetch_status</i>	Contains status information resulting from the last fetch statement. <i>@@fetch_status</i> may contain the following values <ul style="list-style-type: none"> • 0 The fetch statement completed successfully. • -1 The fetch statement resulted in an error. • -2 There is no more data in the result set. <p>This feature is the same as <i>@@sqlstatus</i>, except that it returns different values. It is for Microsoft SQL Server compatibility.</p>
<i>@@identity</i>	The last value inserted into an Identity/Autoincrement column by an insert, load or update statement. <i>@@identity</i> is reset each time a row is inserted into a table. If a statement inserts multiple rows, <i>@@identity</i> reflects the Identity/Autoincrement value for the last row inserted. If the affected table does not contain an Identity/Autoincrement column, <i>@@identity</i> is set to 0. The value of <i>@@identity</i> is not affected by the failure of an insert, load, or update statement, or the rollback of the transaction that contained the failed statement. <i>@@identity</i> retains the last value inserted into an Identity/Autoincrement column, even if the statement that inserted that value fails to commit.
<i>@@isolation</i>	Current isolation level. <i>@@isolation</i> takes the value of the active level.
<i>@@procid</i>	Stored procedure ID of the currently executing procedure.

Variable name	Meaning
<i>@@rowcount</i>	Number of rows affected by the last statement. The value of <i>@@rowcount</i> should be checked immediately after the statement. Inserts, updates, and deletes set <i>@@rowcount</i> to the number of rows affected. With cursors, <i>@@rowcount</i> represents the cumulative number of rows returned from the cursor result set to the client, up to the last fetch request. The <i>@@rowcount</i> is not reset to zero by any statement which does not affect rows, such as an IF statement.
<i>@@servername</i>	Name of the current database server.
<i>@@sqlstatus</i>	Contains status information resulting from the last FETCH statement.
<i>@@version</i>	Version number of the current version of Sybase IQ.

Adaptive Server Enterprise Global Variables Supported in Sybase IQ

This table includes all Adaptive Server Enterprise global variables that are supported in Sybase IQ. Adaptive Server Enterprise global variables that are not supported by Sybase IQ are not included in the list.

This list includes all global variables that return a value, including those for which the value is fixed at NULL, 1, -1, or 0, and might not be meaningful.

Table 8. Adaptive Server Enterprise global variables supported in Sybase IQ

Global variable	Returns
<i>@@char_convert</i>	Returns 0.
<i>@@client_csname</i>	In Adaptive Server Enterprise, the client’s character set name. Set to NULL if client character set has never been initialized; otherwise, contains the name of the most recently used character set. Returns NULL in Sybase IQ.
<i>@@client_csid</i>	In Adaptive Server Enterprise, the client’s character set ID. Set to -1 if client character set has never been initialized; otherwise, contains the most recently used client character set ID from syscharsets. Returns -1 in Sybase IQ.
<i>@@connections</i>	The number of logins since the server was last started.
<i>@@cpu_busy</i>	In Adaptive Server Enterprise, the amount of time, in ticks, that the CPU has spent performing Adaptive Server Enterprise work since the last time Adaptive Server Enterprise was started. In Sybase IQ, returns 0.

Global variable	Returns
@@error	<p>Commonly used to check the error status (succeeded or failed) of the most recently executed statement. Contains 0 if the previous transaction succeeded; otherwise, contains the last error number generated by the system. A statement such as:</p> <pre>if @@error != 0 return</pre> <p>causes an exit if an error occurs. Every statement resets <i>@@error</i>, including PRINT statements or IF tests, so the status check must immediately follow the statement whose success is in question.</p>
@@identity	<p>In Adaptive Server Enterprise, the last value inserted into an IDENTITY column by an INSERT, LOAD, or SELECT INTO statement. <i>@@identity</i> is reset each time a row is inserted into a table. If a statement inserts multiple rows, <i>@@identity</i> reflects the IDENTITY value for the last row inserted. If the affected table does not contain an IDENTITY column, <i>@@identity</i> is set to 0. The value of <i>@@identity</i> is not affected by the failure of an INSERT or SELECT INTO statement, or the rollback of the transaction that contained the failed statement. <i>@@identity</i> retains the last value inserted into an IDENTITY column, even if the statement that inserted that value fails to commit.</p>
@@idle	<p>In Adaptive Server Enterprise, the amount of time, in ticks, that Adaptive Server Enterprise has been idle since the server was last started. In Sybase IQ, returns 0.</p>
@@io_busy	<p>In Adaptive Server Enterprise, the amount of time, in ticks, that Adaptive Server Enterprise has spent performing input and output operations since the server was last started. In Sybase IQ, returns 0.</p>
@@isolation	<p>Current isolation level of the connection. In Adaptive Server Enterprise, <i>@@isolation</i> takes the value of the active level.</p>
@@langid	<p>In Adaptive Server Enterprise, defines the local language ID of the language currently in use. In Sybase IQ, returns 0.</p>
@@language	<p>In Adaptive Server Enterprise, defines the name of the language currently in use. In Sybase IQ, returns "English".</p>
@@maxcharlen	<p>In Adaptive Server Enterprise, maximum length, in bytes, of a character in the Adaptive Server Enterprise default character set. In Sybase IQ, returns 1.</p>
@@max_connections	<p>For the network server, the maximum number of active clients (not database connections, as each client can support multiple connections). For Adaptive Server Enterprise, the maximum number of connections to the server.</p>
@@ncharsize	<p>In Adaptive Server Enterprise, average length, in bytes, of a national character. In Sybase IQ, returns 1.</p>

Global variable	Returns
@@nestlevel	In Adaptive Server Enterprise, nesting level of current execution (initially 0). Each time a stored procedure or trigger calls another stored procedure or trigger, the nesting level is incremented. In Sybase IQ, returns -1.
@@pack_received	In Adaptive Server Enterprise, number of input packets read by Adaptive Server Enterprise since the server was last started. In Sybase IQ, returns 0.
@@pack_sent	In Adaptive Server Enterprise, number of output packets written by Adaptive Server Enterprise since the server was last started. In Sybase IQ, returns 0.
@@packet_errors	In Adaptive Server Enterprise, number of errors that have occurred while Adaptive Server Enterprise was sending and receiving packets. In Sybase IQ, returns 0.
@@procid	Stored procedure ID of the currently executing procedure.
@@servername	Name of the local Adaptive Server Enterprise or Sybase IQ server.
@@spid	In Adaptive Server Enterprise, server process ID number of the current process. In Sybase IQ, the connection handle for the current connection. This is the same value as that displayed by the sa_conn_info procedure.
@@sqlstatus	Contains status information resulting from the last FETCH statement. <i>@@sqlstatus</i> may contain the following values: <ul style="list-style-type: none"> • 0 – the FETCH statement completed successfully. • 1 – the FETCH statement resulted in an error. • 2 – there is no more data in the result set.
@@thresh_hysteresis	In Adaptive Server Enterprise, change in free space required to activate a threshold. In Sybase IQ, returns 0.
@@timeticks	In Adaptive Server Enterprise, number of microseconds per tick. The amount of time per tick is machine-dependent. In Sybase IQ, returns 0.
@@total_errors	In Adaptive Server Enterprise, number of errors that have occurred while Adaptive Server Enterprise was reading or writing. In Sybase IQ, returns 0.
@@total_read	In Adaptive Server Enterprise, number of disk reads by Adaptive Server Enterprise since the server was last started. In Sybase IQ, returns 0.
@@total_write	In Adaptive Server Enterprise, number of disk writes by Adaptive Server Enterprise since the server was last started. In Sybase IQ, returns 0.
@@tranchained	Current transaction mode of the Transact-SQL program. <i>@@tranchained</i> returns 0 for unchained or 1 for chained.

Global variable	Returns
@@trancount	Nesting level of transactions. Each BEGIN TRANSACTION in a batch increments the transaction count.
@@transtate	In Adaptive Server Enterprise, current state of a transaction after a statement executes. In Sybase IQ, returns -1.
@@version	Information about the current version of Adaptive Server Enterprise or Sybase IQ.

Comments

Use comments to attach explanatory text to SQL statements or statement blocks. The database server does not execute comments.

These comment indicators are available in Sybase IQ:

Comment indicator	Description
-- (Double hyphen)	The database server ignores any remaining characters on the line. This is the SQL92 comment indicator.
// (Double slash)	The double slash has the same meaning as the double hyphen.
/* ... */ (Slash-asterisk)	Any characters between the two comment markers are ignored. The two comment markers might be on the same or different lines. Comments indicated in this style can be nested. This style of commenting is also called C-style comments.
% (Percent sign)	The percent sign has the same meaning as the double hyphen. Sybase recommends that you do not use % as a comment indicator.

Note: The double-hyphen and the slash-asterisk comment styles are compatible with Adaptive Server Enterprise.

Examples

This example illustrates the use of double-dash comments:

```
CREATE FUNCTION fullname (firstname CHAR(30),
                          lastname CHAR(30))
RETURNS CHAR(61)
-- fullname concatenates the firstname and lastname
-- arguments with a single space between.
BEGIN
    DECLARE name CHAR(61);
    SET name = firstname || ' ' || lastname;
```

```
RETURN ( name );
END
```

This example illustrates the use of C-style comments:

```
/*
   Lists the names and employee IDs of employees
   who work in the sales department.
*/
CREATE VIEW SalesEmployee AS
SELECT emp_id, emp_lname, emp_fname
FROM "GROUP0".Employees
WHERE DepartmentID = 200
```

NULL Value

Use NULL to specify a value that is unknown, missing, or not applicable.

The NULL value is a special value that is different from any valid value for any data type. However, the NULL value is a legal value in any data type. These are two separate and distinct cases where NULL is used:

Situation	Description
missing	The field does have a value, but that value is unknown.
inapplicable	The field does not apply for this particular row.

SQL allows columns to be created with the NOT NULL restriction. This means that those particular columns cannot contain the NULL value.

The NULL value introduces the concept of three valued logic to SQL. The NULL value compared using any comparison operator with any value including the NULL value is UNKNOWN. The only search condition that returns TRUE is the IS NULL predicate. In SQL, rows are selected only if the search condition in the **WHERE** clause evaluates to TRUE; rows that evaluate to UNKNOWN or FALSE are not selected.

You can also use the **IS [NOT] truth-value** clause, where *truth-value* is one of TRUE, FALSE or UNKNOWN, to select rows where the NULL value is involved.

In the following examples, the column `Salary` contains the NULL value.

Condition	Truth value	Selected?
<code>Salary = NULL</code>	UNKNOWN	NO
<code>Salary <> NULL</code>	UNKNOWN	NO
<code>NOT (Salary = NULL)</code>	UNKNOWN	NO
<code>NOT (Salary <> NULL)</code>	UNKNOWN	NO

Condition	Truth value	Selected?
Salary = 1000	UNKNOWN	NO
Salary IS NULL	TRUE	YES
Salary IS NOT NULL	FALSE	NO
Salary = 1000 IS UNKNOWN	TRUE	YES

The same rules apply when comparing columns from two different tables. Therefore, joining two tables together does not select rows where any of the columns compared contain the NULL value.

The NULL value also has an interesting property when used in numeric expressions. The result of *any* numeric expression involving the NULL value is the NULL value. This means that if the NULL value is added to a number, the result is the NULL value—not a number. If you want the NULL value to be treated as 0, you must use the **ISNULL(expression, 0)** function.

Many common errors in formulating SQL queries are caused by the behavior of NULL. Be careful to avoid these problem areas. Note the effect of three-valued logic when combining search conditions.

Syntax

NULL

Usage

Anywhere

Permissions

Must be connected to the database

Side Effects

None

Example

The following **INSERT** statement inserts a NULL into the `date_returned` column of the `Borrowed_book` table.

```
INSERT
INTO Borrowed_book
( date_borrowed, date_returned, book )
VALUES ( CURRENT DATE, NULL, '1234' )
```

See also

- *Comparison Conditions* on page 33
- *Expressions* on page 22

SQL Language Elements

- *Search Conditions* on page 32
- *Strings* on page 21
- *Three-Valued Logic* on page 46
- *SQL Operators* on page 24
- *Subqueries in Search Conditions* on page 35

SQL Data Types

SQL data types define the type of data to be stored, such as character strings, numbers, and dates.

Character Data Types

Use character data types for storing strings of letters, numbers and symbols.

Syntax

Use this syntax for character data types.

```
CHAR [ ( max-length ) ]
```

```
CHARACTER [ ( max-length ) ]
```

```
CHARACTER VARYING [ ( max-length ) ]
```

```
VARCHAR [ ( max-length ) ]
```

```
UNIQUEIDENTIFIERSTR
```

Usage

This table describes the character data types.

Table 9. Character Data Type

Character Data Type	Description
CHAR	<p>Character data of maximum length <i>max-length</i> bytes. If <i>max-length</i> is omitted, the default is 1. The maximum size allowed is 32KB – 1. See Notes for restrictions on CHAR data greater than 255 bytes.</p> <p>See the notes below on character data representation in the database, and on storage of long strings.</p> <p>All CHAR values are blank padded up to <i>max-length</i>, regardless of whether the BLANK PADDING option is specified. When multibyte character strings are held as a CHAR type, the maximum length is still in bytes, not characters.</p>
CHARACTER	Same as CHAR.

Character Data Type	Description
CHARACTER VARYING	Same as VARCHAR.
LONG VARCHAR	Arbitrary length character data. The maximum size is limited by the maximum size of the database file (currently 2 gigabytes).
TEXT	This is a user-defined data type. It is implemented as a LONG VARCHAR allowing NULL.
VARCHAR	Same as CHAR, except that no blank padding is added to the storage of these strings, and VARCHAR strings can have a maximum length of (32KB – 1). See Notes for restrictions on VARCHAR data greater than 255 bytes.
UNIQUEIDENTIFIERSTR	Domain implemented as CHAR(36). This data type is used for remote data access, when mapping Microsoft SQL Server uniqueidentifier columns.

Note: As a separately licensed option, Sybase IQ supports character large object (CLOB) data with a length ranging from zero (0) to 512TB (terabytes) for an IQ page size of 128KB or 2PB (petabytes) for an IQ page size of 512KB. The maximum length is equal to 4GB multiplied by the database page size. See *Unstructured Data Analytics in Sybase IQ*.

See also

- *Binary Data Types* on page 78
- *NEWID Function [Miscellaneous]* on page 239
- *STRTOUUID Function [String]* on page 311
- *UUIDTOSTR Function [String]* on page 334
- *Binary Data Types* on page 689
- *Character Data Types* on page 688

Storage Sizes

The storage size of character data, given column definition size and input data size.

Table 10. Storage Size of Character Data

Data type	Column definition	Input data	Storage
CHARACTER, CHAR	width of (32K – 1) bytes	(32K – 1) bytes	(32K – 1) bytes
VARCHAR, CHARACTER VARYING	width of (32K – 1) bytes	(32K – 1) bytes	(32K – 1) bytes

Character Sets and Code Pages

Character data is placed in the database using the exact binary representation that is passed from the application.

This usually means that character data is stored in the database with the binary representation of the character set used by your system. You can find documentation about character sets in the documentation for your operating system.

On Windows, code pages are the same for the first 128 characters. If you use special characters from the top half of the code page (accented international language characters), you must be careful with your databases. In particular, if you copy the database to a different machine using a different code page, those special characters are retrieved from the database using the original code page representation. With the new code page, they appear on the window to be the wrong characters.

This problem also appears if you have two clients using the same multiuser server, but running with different code pages. Data inserted or updated by one client might appear incorrect to another.

This problem also shows up if a database is used across platforms. PowerBuilder and many other Windows applications insert data into the database in the standard ANSI character set. If non-Windows applications attempt to use this data, they do not properly display or update the extended characters.

This problem is quite complex. If any of your applications use the extended characters in the upper half of the code page, make sure that all clients and all machines using the database use the same or a compatible code page.

Indexes

All index types, except `DATE`, `TIME`, and `DTTM`, are supported for `CHAR` data and `VARCHAR` data less than or equal to 255 bytes in length.

VARCHAR Data and Trailing Blanks

For a column of data type `VARCHAR`, trailing blanks within the data being inserted are handled differently depending on whether or not the data is enclosed in quotes.

Data inserted using **INSERT**, **UPDATE**, or **LOAD TABLE** can be:

- Enclosed in quotes
- Not enclosed in quotes
- Binary

For a column of data type `VARCHAR`, trailing blanks within the data being inserted are handled as follows:

- For data enclosed in quotes, trailing blanks are never trimmed.

SQL Data Types

- For data not enclosed in quotes:
 - Trailing blanks always trimmed on insert and update.
 - For a **LOAD** statement, you can use the **STRIP RTRIM/OFF LOAD** option to specify whether to have the trailing blanks trimmed. The **STRIP RTRIM/OFF** option applies only to variable-length non-binary data. For example, assume the following schema:

```
CREATE TABLE t( c1 VARCHAR(3) );
LOAD TABLE t( c1 ',' ) ..... STRIP RTRIM // trailing
blanks trimmed

LOAD TABLE t( c1 ',' ) ..... STRIP OFF // trailing blanks
not trimmed

LOAD TABLE t( c1 ASCII(3) ) ... STRIP RTRIM // trailing
blanks not trimmed
LOAD TABLE t( c1 ASCII(3) ) ... STRIP OFF // trailing blanks
trimmed

LOAD TABLE t( c1 BINARY ) ..... STRIP RTRIM // trailing
blanks trimmed
LOAD TABLE t( c1 BINARY ) ..... STRIP OFF // trailing blanks
trimmed
```

- For binary data, trailing blanks are always trimmed.

When you write your applications, do not depend on the existence of trailing blanks in VARCHAR columns. If an application relies on trailing blanks, use a CHAR column instead of a VARCHAR column.

Restriction on CHAR and VARCHAR Data Over 255 Bytes

Only the default index, **WD**, **TEXT**, and **CMP** index types are supported for CHAR and VARCHAR columns over 255 bytes.

You cannot create an **LF**, **HG**, **HNG**, **DATE**, **TIME**, or **DTTM** index for these columns.

Character Data Compatibility

Character data compatibility differences exist between Sybase IQ and Adaptive Server Enterprise and SQL Anywhere.

- The CHARACTER (*n*) alternative for CHAR is not supported in Adaptive Server Enterprise.
- Sybase IQ does not support the NCHAR, NVARCHAR, UNICHAR, and UNIVARCHAR data types provided by Adaptive Server Enterprise. Sybase IQ supports Unicode in the CHAR and VARCHAR data types.
- Sybase IQ supports a longer LONG VARCHAR data type than SQL Anywhere. See *Unstructured Data Analytics in Sybase IQ*.
- For compatibility between Sybase IQ and Adaptive Server Enterprise, always specify a length for character data types.

Long strings

Values up to 254 characters are stored as short strings, with a preceding length byte. Any values that are longer than 255 bytes are considered long strings. Characters after the 255th are stored separate from the row containing the long string value.

SQL Anywhere treats CHAR, VARCHAR, and LONG VARCHAR columns all as the same type.

There are several functions (see SQL Functions) that will ignore the part of any string past the 255th character. They are **soundex**, **similar**, and all of the date functions. Also, any arithmetic involving the conversion of a long string to a number will work on only the first 255 characters. It would be extremely unusual to run in to one of these limitations.

All other functions and all other operators work with the full length of long strings.

Numeric Data Types

Use numeric data types for storing numerical data.

Syntax

Use this syntax for numeric data types.

```
[ UNSIGNED ] BIGINT
```

```
[ UNSIGNED ] { INT | INTEGER }
```

```
SMALLINT
```

```
TINYINT
```

```
DECIMAL [ ( precision [ , scale ] ) ]
```

```
NUMERIC [ ( precision [ , scale ] ) ]
```

```
DOUBLE
```

```
FLOAT [ ( precision ) ]
```

```
REAL
```

Usage for Numeric Data Types

Be aware of these points when using numeric data types.

- The INTEGER, NUMERIC, and DECIMAL data types are sometimes called exact numeric data types, in contrast to the approximate numeric data types FLOAT, DOUBLE, and REAL. Only exact numeric data is guaranteed to be accurate to the least significant digit specified after arithmetic operations.

SQL Data Types

- Do not fetch `TINYINT` columns into Embedded SQL variables defined as `CHAR` or `UNSIGNED CHAR`, since the result is an attempt to convert the value of the column to a string and then assign the first byte to the variable in the program.
- A period is the only decimal separator (decimal point); comma is not supported as a decimal separator.

Table 11. Numeric Data Type

Numeric Data Type	Description
BIGINT	<p>A signed 64-bit integer, requiring 8 bytes of storage.</p> <p>You can specify integers as <code>UNSIGNED</code>. By default the data type is signed. Its range is between -9223372036854775808 and 9223372036854775807 (signed) or from 0 to 18446744073709551615 (unsigned).</p>
INT or INTEGER	<p>A signed 32-bit integer with a range of values between -2147483648 and 2147483647 requiring 4 bytes of storage.</p> <p>The <code>INTEGER</code> data type is an exact numeric data type; its accuracy is preserved after arithmetic operations.</p> <p>You can specify integers as <code>UNSIGNED</code>; by default the data type is signed. The range of values for an unsigned integer is between 0 and 4294967295.</p>
SMALLINT	<p>A signed 16-bit integer with a range between -32768 and 32767, requiring 2 bytes of storage.</p> <p>The <code>SMALLINT</code> data type is an exact numeric data type; its accuracy is preserved after arithmetic operations.</p>
TINYINT	<p>An unsigned 8-bit integer with a range between 0 and 255, requiring 1 byte of storage.</p> <p>The <code>TINYINT</code> data type is an exact numeric data type; its accuracy is preserved after arithmetic operations.</p>

Numeric Data Type	Description
DECIMAL	A signed decimal number with <i>precision</i> total digits and with <i>scale</i> of the digits after the decimal point. The precision can equal 1 to 126, and the scale can equal 0 up to precision value. The defaults are scale = 38 and precision = 126. Results are calculated based on the actual data type of the column to ensure accuracy, but you can set the maximum scale of the result returned to the application using the MAX_CLIENT_NUMERIC_SCALE option.
NUMERIC	Same as DECIMAL.
DOUBLE	<p>A signed double-precision floating-point number stored in 8 bytes. The range of absolute, nonzero values is between 2.2250738585072014e-308 and 1.797693134862315708e+308. Values held as DOUBLE are accurate to 15 significant digits, but might be subject to rounding errors beyond the fifteenth digit.</p> <p>The DOUBLE data type is an approximate numeric data type; it is subject to rounding errors after arithmetic operations.</p>
FLOAT	<p>If <i>precision</i> is not supplied, the FLOAT data type is the same as the REAL data type. If <i>precision</i> supplied, then the FLOAT data type is the same as the REAL or DOUBLE data type, depending on the value of the precision. The cutoff between REAL and DOUBLE is platform-dependent, and it is the number of bits used in the mantissa of single-precision floating point number on the platform.</p> <p>When a column is created using the FLOAT data type, columns on all platforms are guaranteed to hold the values to at least the specified minimum precision. In contrast, REAL and DOUBLE do not guarantee a platform-independent minimum precision.</p> <p>The FLOAT data type is an approximate numeric data type; it is subject to rounding errors after arithmetic operations.</p>

Numeric Data Type	Description
REAL	<p>A signed single-precision floating-point number stored in 4 bytes. The range of absolute, nonzero values is 1.175494351e-38 to 3.402823466e+38. Values held as REAL are accurate to 6 significant digits, but might be subject to rounding errors beyond the sixth digit.</p> <p>The REAL data type is an approximate numeric data type; it is subject to rounding errors after arithmetic operations.</p>

This table lists the storage required for a decimal number.

Table 12. Storage size for a decimal number

Precision	Storage
1 to 4	2 bytes
5 to 9	4 bytes
10 to 18	8 bytes
19 to 126	See below

The storage requirement in bytes for a decimal value with a precision greater than 18 can be calculated using the following formula:

$$4 + 2 * (\text{int}(((\text{prec} - \text{scale}) + 3) / 4) + \text{int}((\text{scale} + 3) / 4) + 1)$$

where *int* takes the integer portion of its argument. The storage used by a column is based upon the precision and scale of the column. Each cell in the column has enough space to hold the largest value of that precision and scale. For example:

NUMERIC(18,4) takes 8 bytes per cell
 NUMERIC(19,4) takes 16 bytes per cell

The DECIMAL data type is an exact numeric data type; its accuracy is preserved to the least significant digit after arithmetic operations. Its maximum absolute value is the number of nines defined by [*precision* - *scale*], followed by the decimal point, and then followed by the number of nines defined by *scale*. The minimum absolute nonzero value is the decimal point, followed by the number of zeros defined by [*scale* - 1], then followed by a single one. For example:

NUMERIC (3,2) Max positive = 9.99 Min non-zero = 0.01 Max negative = -9.99

If neither precision nor scale is specified for the explicit conversion of NULL to NUMERIC, the default is NUMERIC(1,0). For example,

```
SELECT CAST( NULL AS NUMERIC ) A,
       CAST( NULL AS NUMERIC(15,2) ) B
```

is described as:

```
A NUMERIC(1,0)
B NUMERIC(15,2)
```

Note: The maximum value supported in SQL Anywhere for the numeric function is 255. If the precision of the numeric function exceeds the maximum value supported in SQL Anywhere, the following error occurs: "The result datatype for function '_funcname' exceeds the maximum supported numeric precision of 255. Please set the proper value for precision in numeric function, 'location'"

Numeric Data Compatibility

Numeric data compatibility differences exist between Sybase IQ and Adaptive Server Enterprise and SQL Anywhere.

- In embedded SQL, fetch TINYINT columns into 2-byte or 4-byte integer columns. Also, to send a TINYINT value to a database, the C variable should be an integer.
- Adaptive Server Enterprise 12.5.x versions do not support unsigned integers. You can map Sybase IQ unsigned integers to Adaptive Server Enterprise signed integers or numeric data, and the data are converted implicitly.
 - Map IQ UNSIGNED SMALLINT data to ASE INT
 - If you have negative values, map IQ UNSIGNED BIGINT to ASE NUMERIC (*precision, scale*)

To avoid performance issues for cross-database joins on UNSIGNED BIGINT columns, the best approach is to cast to a (signed) BIGINT on the Sybase IQ side.
- You should avoid default precision and scale settings for NUMERIC and DECIMAL data types, as these differ by product:

Database	Default precision	Default scale
Sybase IQ	126	38
Adaptive Server Enterprise	18	0
SQL Anywhere	30	6

- The FLOAT (*p*) data type is a synonym for REAL or DOUBLE, depending on the value of *p*. For Adaptive Server Enterprise, REAL is used for *p* less than or equal to 15, and DOUBLE for *p* greater than 15. For Sybase IQ, the cutoff is platform-dependent, but on all platforms, the cutoff value is greater than 22.
- Sybase IQ includes two user-defined data types, MONEY and SMALLMONEY, which are implemented as NUMERIC(19,4) and NUMERIC(10,4) respectively. They are provided primarily for compatibility with Adaptive Server Enterprise.

Indexes

This section describes the relationship between index types and numeric data types.

- The **CMP** and **HNG** index types do not support the **FLOAT**, **DOUBLE**, and **REAL** data types, and the **HG** index type is not recommended.
- The **WD**, **DATE**, **TIME**, and **DTTM** index types do not support the numeric data types.

Binary Data Types

Use binary data types for storing raw binary data, such as pictures, in a hexadecimal-like notation, up to a length of (32K – 1) bytes.

Syntax

```
BINARY [ ( length ) ]
```

```
VARBINARY [ ( max-length ) ]
```

```
UNIQUEIDENTIFIER
```

See also

- *NEWID Function [Miscellaneous]* on page 239
- *STRTOUUID Function [String]* on page 311
- *UUIDTOSTR Function [String]* on page 334
- *Character Data Types* on page 69
- *Binary Data Types* on page 689

Usage for Binary Data Types

Binary data begins with the characters “0x” or “0X” and can include any combination of digits and the uppercase and lowercase letters A through F.

You can specify the column length in bytes, or use the default length of 1 byte. Each byte stores 2 hexadecimal digits. Even though the default length is 1 byte, Sybase recommends that you always specify an even number of characters for **BINARY** and **VARBINARY** column length. If you enter a value longer than the specified column length, Sybase IQ truncates the entry to the specified length without warning or error.

Table 13. Binary Data Type

Binary Data Type	Description
BINARY	Binary data of length <i>length</i> bytes. If <i>length</i> is omitted, the default is 1 byte. The maximum size allowed is 32767 bytes. Use the fixed-length binary type BINARY for data in which all entries are expected to be approximately equal in length. Because entries in BINARY columns are zero-padded to the column length <i>length</i> , they might require more storage space than entries in VAR-BINARY columns.
VARBINARY	Binary data up to a length of <i>max-length</i> bytes. If <i>max-length</i> is omitted, the default is 1 byte. The maximum size allowed is (32K – 1) bytes. Use the variable-length binary type VARBINARY for data that is expected to vary greatly in length.
UNIQUEIDENTIFIER	The UNIQUEIDENTIFIER data type is used for storage of UUID (also known as GUID) values.

Treatment of Trailing Zeros

All BINARY columns are padded with zeros to the full width of the column. Trailing zeros are truncated in all VARBINARY columns.

The following example creates a table with all four variations of BINARY and VARBINARY data types defined with NULL and NOT NULL. The same data is inserted in all four columns and is padded or truncated according to the data type of the column.

```
CREATE TABLE zeros (bnot BINARY(5) NOT NULL,
    bnull BINARY(5) NULL,
    vbnnot VARBINARY(5) NOT NULL,
    vbnull VARBINARY(5) NULL);
INSERT zeros VALUES (0x12345000, 0x12345000,
    0x12345000, 0x12345000);
INSERT zeros VALUES (0x123, 0x123, 0x123, 0x123);
INSERT zeros VALUES (0x0, 0x0, 0x0, 0x0);
INSERT zeros VALUES ('0027100000000ae1b',
    '0027100000000ae1b', '0027100000000ae1b',
    '0027100000000ae1b');
SELECT * FROM zeros;
```

bnot	bnull	vbnnot	vbnull
0x1234500000	0x1234500000	0x12345000	0x12345000
0x0123000000	0x0123000000	0x0123	0x0123

bnot	bnull	vbnot	vbnull
0x0000000000	0x0000000000	0x00	0x00
0x3030323731	0x3030323731	0x3030323731	0x3030323731

Because each byte of storage holds 2 hexadecimal digits, Sybase IQ expects binary entries to consist of the characters “0x” followed by an even number of digits. When the “0x” is followed by an odd number of digits, Sybase IQ assumes that you omitted the leading 0 and adds it for you.

Input values “0x00” and “0x0” are stored as “0x00” in variable-length binary columns (VARBINARY). In fixed-length binary columns (BINARY), the value is padded with zeros to the full length of the field:

```
INSERT zeros VALUES (0x0, 0x0, 0x0, 0x0);
SELECT * FROM zeros
```

bnot	bnull	vbnot	vbnull
0x0000000000	0x0000000000	0x00	0x00

If the input value does not include “0x”, Sybase IQ assumes that the value is an ASCII value and converts it. For example:

```
CREATE TABLE sample (col_bin BINARY(8));
INSERT sample VALUES ('002710000000aeb');
SELECT * FROM sample;
```

col_bin
0x3030323731303030

Note: In the above example, ensure you set the `string_rtruncation` option to “off”.

When you select a BINARY value, you must specify the value with the padded zeros or use the **CAST** function. For example:

```
SELECT * FROM zeros WHERE bnot = 0x0123000000;
```

or :

```
SELECT * FROM zeros WHERE bnot = CAST(0x0123 as binary(5));
```

ASCII Data From a Flat File

Any ASCII data loaded from a flat file into a binary type column (BINARY or VARBINARY) is stored as nibbles.

For example, if 0x1234 or 1234 is read from a flat file into a binary column, Sybase IQ stores the value as hexadecimal 1234. Sybase IQ ignores the “0x” prefix. If the input data contains any characters out of the range 0 – 9, a – f, and A – F, the data is rejected.

Storage Size

Familiarize yourself with the storage size of binary data.

Table 14. Storage size of binary data

Data type	Column definition	Input data	Storage
VARBINARY	width of (32K – 1) bytes	(32K – 1) bytes binary	(32K – 1) bytes
VARBINARY	width of (32K – 1) bytes	(64K – 2) bytes ASCII	(32K – 1) bytes
BINARY	width of (32K – 1) bytes	(32K – 1) bytes binary	(32K – 1) bytes
BINARY	width of (32K – 1) bytes	(64K – 2) bytes ASCII	(32K – 1) bytes

The exact form in which you enter a particular value depends on the platform you are using. Therefore, calculations involving binary data might produce different results on different machines.

For platform-independent conversions between hexadecimal strings and integers, use the **INTTOHEX** and **HEXTOINT** functions rather than the platform-specific **CONVERT** function.

See also

- *Data Type Conversion Functions* on page 106
- *Data Type Conversions* on page 93

String Operators

The concatenation string operators `||` and `+` both support binary type data.

Explicit conversion of binary operands to character data types is not necessary with the `||` operator. Explicit and implicit data conversion produce different results, however.

Restrictions on BINARY and VARBINARY Data

Restrictions apply to columns containing **BINARY** and **VARBINARY** data.

- You cannot use the aggregate functions **SUM**, **AVG**, **STDDEV**, or **VARIANCE** with the binary data types. The aggregate functions **MIN**, **MAX**, and **COUNT** *do* support the binary data types **BINARY** and **VARBINARY**.
- **HNG**, **WD**, **DATE**, **TIME**, and **DTTM** indexes do not support **BINARY** or **VARBINARY** data.
- Only the default index, **CMP** index, and **TEXT** index types are supported for **BINARY** and **VARBINARY** data greater than 255 bytes in length.
- Bit operations are supported on **BINARY** and **VARBINARY** data that is 8 bytes or less in length.

Binary Data Compatibility

The treatment of trailing zeros in binary data differs between Sybase IQ, SQL Anywhere, and Adaptive Server Enterprise.

Table 15. Treatment of trailing zeros

Data type	Sybase IQ	SQL Anywhere	Adaptive Server Enterprise
BINARY NOT NULL	Padded	Not padded	Padded
BINARY NULL	Padded	Not padded	Not padded
VARBINARY NOT NULL	Truncated, not padded	Truncated, not padded	Truncated, not padded
VARBINARY NULL	Truncated, not padded	Truncated, not padded	Truncated, not padded

Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ all support the **STRING_RTRUNCATION** database option, which affects error message reporting when an **INSERT** or **UPDATE** string is truncated. For Transact-SQL compatible string comparisons, set the **STRING_RTRUNCATION** option to the same value in both databases.

You can also set the **STRING_RTRUNCATION** option ON when loading data into a table, to alert you that the data is too large to load into the field. The default value is ON.

Bit operations on binary type data are not supported by Adaptive Server Enterprise. SQL Anywhere only supports bit operations against the first four bytes of binary type data. Sybase IQ supports bit operations against the first eight bytes of binary type data.

UNIQUEIDENTIFIER

The **UNIQUEIDENTIFIER** data type is used for storage of UUID (also known as GUID) values.

The **UNIQUEIDENTIFIER** data type is often used for a primary key or other unique column to hold UUID (Universally Unique Identifier) values that can be used to uniquely identify rows. The **NEWID** function generates UUID values in such a way that a value produced on one computer does not match a UUID produced on another computer. **UNIQUEIDENTIFIER** values generated using **NEWID** can therefore be used as keys in a synchronization environment.

For example, the following statement updates the table `mytab` and sets the value of the column `uid_col` to a unique identifier generated by the **NEWID** function, if the current value of the column is `NULL`.

```
UPDATE mytab
  SET uid_col = NEWID()
  WHERE uid_col IS NULL
```

If you execute the following statement,

```
SELECT NEWID( )
```

the unique identifier is returned as a `BINARY (16)`. For example, the value might be `0xd3749fe09cf446e399913bc6434f1f08`. You can convert this string into a readable format using the `UUIDTOSTR()` function.

UUID values are also referred to as GUIDs (Globally Unique Identifier).

The `STRTOUUID` and `UUIDTOSTR` functions are used to convert values between `UNIQUEIDENTIFIER` and string representations.

`UNIQUEIDENTIFIER` values are stored and returned as `BINARY (16)`.

Because `UNIQUEIDENTIFIER` values are large, using `UNSIGNED BIGINT` or `UNSIGNED INT` identity columns instead of `UNIQUEIDENTIFIER` is more efficient, if you do not need cross database unique identifiers.

Standards and compatibility for UNIQUEIDENTIFIER

These standards and compatibilities apply to `UNIQUEIDENTIFIER` values.

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Supported by SQL Anywhere. Not supported by Adaptive Server Enterprise.
- Backwards compatibility—In databases created before Sybase IQ version 12.7, the `STRTOUUID`, `UUIDTOSTR`, and `NEWID` functions were supported through CIS functional compensation. In versions 15.3 and later, the `STRTOUUID`, `UUIDTOSTR`, and `NEWID` functions are native Sybase IQ functions.

Binary Large Object Data

As a separately licensed option, Sybase IQ supports binary large object (BLOB) data with a length ranging from zero (0) to 512TB (terabytes) for an IQ page size of 128KB or 2PB (petabytes) for an IQ page size of 512KB.

The maximum length is equal to 4GB multiplied by the database page size. See *Unstructured Data Analytics in Sybase IQ*

Bit Data Type

Use the `BIT` data type for storing Boolean values.

Data type	Values	Supported by
BIT	0 or 1	Sybase IQ and Adaptive Server Enterprise

Usage

`BIT` stores only the values 0 or 1.

SQL Data Types

Inserting any nonzero value into a BIT column stores a 1 in the column. Inserting any zero value into a BIT column stores a 0.

Only the default index type is supported for BIT data.

Bit Data Compatibility

Adaptive Server Enterprise BIT datatypes only allow 0 or 1 values.

Date and Time Data Types

Use date and time data types for storing dates and times.

Syntax

Use this syntax for date and time data types.

DATE

DATETIME

SMALLDATETIME

TIME

TIMESTAMP

See also

- *TIMESTAMP Special Value* on page 57
- *CURRENT TIMESTAMP Special Value* on page 55
- *CURRENT TIME Special Value* on page 55
- *CURRENT DATE Special Value* on page 54
- *Retrieve Dates and Times* on page 87
- *BIGTIME and BIGDATETIME Support* on page 690

Usage for Date and Time Data Types

Familiarize yourself with these usage considerations before using date and time data types.

Table 16.

Date and Time Data Type	Description
DATE	A calendar date, such as a year, month and day. The year can be from 0001 to 9999. The day must be a nonzero value, so that the minimum date is 0001-01-01. A DATE value requires 4 bytes of storage.

Date and Time Data Type	Description
DATETIME	A domain, implemented as <code>TIMESTAMP</code> . <code>DATETIME</code> is provided primarily for compatibility with Adaptive Server Enterprise.
SMALLDATETIME	A domain, implemented as <code>TIMESTAMP</code> . <code>SMALLDATETIME</code> is provided primarily for compatibility with Adaptive Server Enterprise.
TIME	Time of day, containing hour, minute, second, and fraction of a second. The fraction is stored to 6 decimal places. A <code>TIME</code> value requires 8 bytes of storage. (ODBC standards restrict <code>TIME</code> data type to an accuracy of seconds. For this reason, do not use <code>TIME</code> data types in WHERE clause comparisons that rely on a higher accuracy than seconds.)
TIMESTAMP	Point in time, containing year, month, day, hour, minute, second, and fraction of a second. The fraction is stored to 6 decimal places. The day must be a nonzero value. A <code>TIMESTAMP</code> value requires 8 bytes of storage.

The valid range of the `TIMESTAMP` data type is from 0001-01-01 00:00:00.000000 to 9999-12-31 23:59:59.999999. The display of `TIMESTAMP` data outside the range of 1600-02-28 23:59:59 to 7911-01-01 00:00:00 might be incomplete, but the complete datetime value is stored in the database; you can see the complete value by first converting it to a character string. You can use the **CAST()** function to do this, as in the following example, which first creates a table with `DATETIME` and `TIMESTAMP` columns, then inserts values where the date is greater 7911-01-01.

```
create table mydates (id int, descript char(20),
  datetime_null datetime, timestamp_null timestamp);

insert into mydates values (1, 'example', '7911-12-30
  23:59:59', '7911-12-30 06:03:44');
commit;
```

When you select without using **CAST**, hours and minutes are set to 00:00:

```
select * from mydates;

1, 'example', '7911-12-30 00:00:59.000', '7911-12-30
00:00:44.000'
```

When you select using cast, you see the complete timestamp:

```
select id, descript, cast(datetime_null as char(21)),
  cast(timestamp_null as char(21)) from mydates;
```

```
1, 'example', '7911-12-30 23:59:59.0', '7911-12-30 06:03:44.0'
```

See also

- *Compatibility of String to Datetime Conversions* on page 93

Index Types Supported

These index types are supported by date and time data.

- All date and time data types support the **CMP**, **HG**, **HNG**, and **LF** index types; the **WD** index type is not supported.
- DATE data supports the **DATE** index.
- TIME data supports the **TIME** index.
- DATETIME and TIMESTAMP data support the **DTTM** index.

Send Dates and Times

You send dates and times to the database in these ways.

- Using any interface, as a string
- Using ODBC, as a **TIMESTAMP** structure
- Using Embedded SQL, as a **SQLDATETIME** structure

When you send a time to the database as a string (for the **TIME** data type) or as part of a string (for **TIMESTAMP** or **DATE** data types), hours, minutes, and seconds must be separated by colons in the format *hh:mm:ss.sss*, but can appear anywhere in the string. As an option, a period can separate the seconds from fractions of a second, as in *hh:mm:ss.sss*. The following are valid and unambiguous strings for specifying times:

```
21:35 -- 24 hour clock if no am or pm specified
10:00pm -- pm specified, so interpreted as 12 hour clock
10:00 -- 10:00am in the absence of pm
10:23:32.234 -- seconds and fractions of a
                second included
```

When you send a date to the database as a string, conversion to a date is automatic. You can supply the string in one of two ways:

- As a string of format *yyyy/mm/dd* or *yyyy-mm-dd*, which is interpreted unambiguously by the database
- As a string interpreted according to the **DATE_ORDER** database option

Date format strings cannot contain any multibyte characters. Only single-byte characters are allowed in a date/time/datetime format string, even when the collation order of the database is a multibyte collation order like 932JPN.

Retrieve Dates and Times

You retrieve dates and times from the database in these ways.

- Using any interface, as a string
- Using ODBC, as a `TIMESTAMP` structure
- Using embedded SQL, as a `SQLDATETIME` structure

Usage

When a date or time is retrieved as a string, it is retrieved in the format specified by the database options `DATE_FORMAT`, `TIME_FORMAT` and `TIMESTAMP_FORMAT`.

The following operators are allowed on dates:

Table 17. Operators

Operator	Description
timestamp + integer	Add the specified number of days to a date or timestamp.
timestamp - integer	Subtract the specified number of days from a date or timestamp.
date - date	Compute the number of days between two dates or timestamps.
date + time	Create a timestamp combining the given date and time.

See also

- *TIMESTAMP Special Value* on page 57
- *CURRENT_TIMESTAMP Special Value* on page 55
- *CURRENT_TIME Special Value* on page 55
- *CURRENT_DATE Special Value* on page 54
- *Date and Time Data Types* on page 84

Compare Dates and Times

To compare a date to a string *as a string*, use the `DATEFORMAT` function or `CAST` function to convert the date to a string before comparing.

Usage

```
DATEFORMAT(invoice_date, 'yyyy/mm/dd') = '1992/05/23'
```

You can use any allowable date format for the `DATEFORMAT` string expression.

SQL Data Types

Date format strings must not contain any multibyte characters. Only single-byte characters are allowed in a date/time/datetime format string, even when the collation order of the database is a multibyte collation order like 932JPN.

If '?' represents a multibyte character, the following query fails:

```
SELECT DATEFORMAT ( StartDate, 'yy?') FROM Employees;
```

Instead, move the multibyte character outside of the date format string using the concatenation operator:

```
SELECT DATEFORMAT (StartDate, 'yy') + '?' FROM Employees;
```

Unambiguous Dates and Times

Using the unambiguous date format prevents misinterpretation of dates according to the user's DATE_ORDER setting.

Usage

Dates in the format *yyyy/mm/dd* or *yyyy-mm-dd* are always recognized as dates regardless of the DATE_ORDER setting. You can use other characters as separators; for example, a question mark, a space character, or a comma. You should use this format in any context where different users might be employing different DATE_ORDER settings. For example, in stored procedures, use of the unambiguous date format prevents misinterpretation of dates according to the user's DATE_ORDER setting.

A string of the form *hh.mm.ss.sss* is also interpreted unambiguously as a time.

For combinations of dates and times, any unambiguous date and any unambiguous time yield an unambiguous date-time value. Also, the following form is an unambiguous date-time value:

```
YYYY-MM-DD HH.MM.SS.SSSSSS
```

You can use periods in the time only in combination with a date.

In other contexts, you can use a more flexible date format. Sybase IQ can interpret a wide range of strings as formats. The interpretation depends on the setting of the DATE_ORDER database option. The DATE_ORDER database option can have the value 'MDY', 'YMD', or 'DMY'. For example, to set the DATE_ORDER option to 'DMY' enter:

```
SET OPTION DATE_ORDER = 'DMY' ;
```

The default DATE_ORDER setting is 'YMD'. The ODBC driver sets the DATE_ORDER option to 'YMD' whenever a connection is made. Use the **SET OPTION** statement to change the value.

The database option DATE_ORDER determines whether the string 10/11/12 is interpreted by the database as Oct 11 1912, Nov 12 1910, or Nov 10 1912. The year, month, and day of a date string should be separated by some character (for example "/", "-", or space) and appear in the order specified by the DATE_ORDER option.

You can supply the year as either 2 or 4 digits. The value of the `NEAREST_CENTURY` option [TSQL] affects the interpretation of 2-digit years: 2000 is added to values less than `NEAREST_CENTURY`, and 1900 is added to all other values. The default value of this option is 50. Thus, by default, 50 is interpreted as 1950, and 49 is interpreted as 2049.

The month can be the name or number of the month. The hours and minutes are separated by a colon, but can appear anywhere in the string.

Sybase recommends that you always specify the year using the 4-digit format.

With an appropriate setting of `DATE_ORDER`, the following strings are all valid dates:

```
99-05-23 21:35
99/5/23
1999/05/23
May 23 1999
23-May-1999
Tuesday May 23, 1999 10:00pm
```

If a string contains only a partial date specification, default values are used to fill out the date. The following defaults are used:

- year—1900
- month—No default
- day—1 (useful for month fields; for example, ‘May 1999’ is the date ‘1999-05-01 00:00’)
- hour, minute, second, fraction—0

Domains

Domains are aliases for built-in data types, including precision and scale values where applicable.

Domains, also called user-defined data types, allow columns throughout a database to be defined automatically on the same data type, with the same `NULL` or `NOT NULL` condition. This encourages consistency throughout the database. Domain names are case-insensitive. Sybase IQ returns an error if you attempt to create a domain with the same name as an existing domain except for case.

Simple Domains

You create domains using the `CREATE DOMAIN` statement.

The following statement creates a data type named `street_address`, which is a 35-character string:

```
CREATE DOMAIN street_address CHAR( 35 )
```

You can use `CREATE DATATYPE` as an alternative to `CREATE DOMAIN`, but this is not recommended, as `CREATE DOMAIN` is the syntax used in the ISO/ANSI SQL standard.

SQL Data Types

Resource authority is required to create data types. Once a data type is created, the user ID that executed the **CREATE DOMAIN** statement is the owner of that data type. Any user can use the data type, and unlike other database objects, the owner name is never used to prefix the data type name.

The `street_address` data type may be used in exactly the same way as any other data type when defining columns. For example, the following table with two columns has the second column as a `street_address` column:

```
CREATE TABLE twocol (id INT,  
street street_address)
```

Owners or DBAs can drop domains by issuing a **COMMIT** and then using the **DROP DOMAIN** statement:

```
DROP DOMAIN street_address
```

You can carry out this statement only if no tables in the database are using data type.

Constraints and Defaults with User-Defined Data Types

Many of the attributes associated with columns, such as allowing NULL values, having a DEFAULT value, and so on, can be built into a user-defined data type. Any column that is defined on the data type automatically inherits the NULL setting, CHECK condition, and DEFAULT values. This allows uniformity to be built into columns with a similar meaning throughout a database.

For example, many primary key columns in the demo database are integer columns holding ID numbers. The following statement creates a data type that may be useful for such columns:

```
CREATE DOMAIN id INT  
NOT NULL  
DEFAULT AUTOINCREMENT  
CHECK( @col > 0 )
```

Any column created using the data type `id` is not allowed to hold NULLs, defaults to an autoincremented value, and must hold a positive number. Any identifier could be used instead of `col` in the `@col` variable.

The attributes of the data type can be overridden if needed by explicitly providing attributes for the column. A column created on data type `id` with NULL values explicitly allowed does allow NULLs, regardless of the setting in the `id` data type.

CREATE DOMAIN Statement

Creates a user-defined data type in the database.

Syntax

```
CREATE { DOMAIN | DATATYPE } domain-name data-type  
... [ NOT ] NULL ]  
... [ DEFAULT default-value ]
```

Parameters

- **domain-name:** – identifier
- **data-type:** – built-in data type, with precision and scale
- **default-value:** – *special-value* | *string* | *global variable* | [-] *number* | (*constant-expression*) | *built-in-function*(*constant-expression*) | **AUTOINCREMENT** | **CURRENT DATABASE** | **CURRENT REMOTE USER** | **NULL** | **TIMESTAMP** | **LAST USER**
- **special-value:** – **CURRENT** { **DATE** | **TIME** | **TIMESTAMP** | **USER** | **PUBLISHER** } | **USER**

Examples

- **Example 1** – Create a data type named **address**, which holds a 35-character string, and which may be NULL:

```
CREATE DOMAIN address CHAR( 35 ) NULL
```

Usage

User-defined data types are aliases for built-in data types, including precision and scale values, where applicable. They improve convenience and encourage consistency in the database.

Sybase recommends that you use **CREATE DOMAIN**, rather than **CREATE DATATYPE**, as **CREATE DOMAIN** is the ANSI/ISO SQL3 term.

The user who creates a data type is automatically made the owner of that data type. No owner can be specified in the **CREATE DATATYPE** statement. The user-defined data type name must be unique, and all users can access the data type without using the owner as prefix.

User-defined data types are objects within the database. Their names must conform to the rules for identifiers. User-defined data type names are always case-insensitive, as are built-in data type names.

By default, user-defined data types allow NULLs unless the **allow_nulls_by_default** option is set to OFF. In this case, new user-defined data types by default do not allow NULLs. The nullability of a column created on a user-defined data type depends on the setting of the definition of the user-defined data type, not on the setting of the **allow_nulls_by_default** option when the column is referenced. Any explicit setting of NULL or NOT NULL in the column definition overrides the user-defined data type setting.

The **CREATE DOMAIN** statement allows you to specify DEFAULT values on user-defined data types. The DEFAULT value specification is inherited by any column defined on the data type. Any DEFAULT value explicitly specified on the column overrides that specified for the data type. For more information on the use of column DEFAULT values, see *System Administration Guide: Volume 1 > Data Integrity > Column Defaults Encourage Data Integrity*.

The **CREATE DOMAIN** statement lets you incorporate a rule, called a CHECK condition, into the definition of a user-defined data type.

SQL Data Types

Sybase IQ enforces CHECK constraints for base, global temporary, local temporary tables, and user-defined data types.

To drop the data type from the database, use the **DROP** statement. You must be either the owner of the data type or have DBA authority in order to drop a user-defined data type.

See also *Reference: Building Blocks, Tables, and Procedures > SQL Data Types*.

Side effects:

- Automatic commit

Standards

- SQL—ISO/ANSI SQL compliant.
- Sybase—Not supported by Adaptive Server Enterprise. Transact-SQL provides similar functionality using the **sp_addtype** system procedure and the **CREATE DEFAULT** and **CREATE RULE** statements.

Permissions

Must have RESOURCE authority.

Domain Compatibility

Domain compatibility differences exist between Sybase IQ and Adaptive Server Enterprise and SQL Anywhere.

- Named constraints and defaults—In Sybase IQ, user-defined data types are created with a base data type, and optionally, a NULL or NOT NULL condition. Named constraints and named defaults are not supported.
- Creating data types—In Sybase IQ, you can use either the **sp_addtype** system procedure or the **CREATE DOMAIN** statement to add a domain. In Adaptive Server Enterprise, you must use **sp_addtype**. The owner of the **sp_addtype** and other stored procedures inherited from SQL Anywhere is dbo. The creator for any object created using SQL Anywhere stored procedure is also dbo, and thus a user without DBA authority may not alter or drop domains created using **sp_addtype**. You must have DBA authority to alter or drop domains created with **CREATE DOMAIN**.

Data Type Conversions

Type conversions happen automatically, or you can explicitly request them using the **CAST** or **CONVERT** function.

Usage

If a string is used in a numeric expression or as an argument to a function expecting a numeric argument, the string is converted to a number before use.

If a number is used in a string expression or as a string function argument, then the number is converted to a string before use.

All date constants are specified as strings. The string is automatically converted to a date before use.

There are certain cases where the automatic data type conversions are not appropriate.

```
'12/31/90' + 5 -- Tries to convert the string to a number
'a' > 0       -- Tries to convert 'a' to a number
```

You can use the **CAST** or **CONVERT** function to force type conversions.

The following functions can also be used to force type conversions:

- **DATE(expression)** – converts the expression into a date, and removes any hours, minutes or seconds. Conversion errors might be reported.
- **DATETIME(expression)** – converts the expression into a timestamp. Conversion errors might be reported.
- **STRING(expression)** – similar to `CAST(value AS CHAR)`, except that `string(NULL)` is the empty string (`''`), whereas `CAST(NULL AS CHAR)` is the `NULL` value.

See also

- *Data Type Conversion Functions* on page 106
- *Storage Size* on page 81

Compatibility of String to Datetime Conversions

There are some differences in behavior between Sybase IQ and Adaptive Server Enterprise when converting strings to date and time data types.

If you convert a string containing only a time value (no date) to a date/time data type, Sybase IQ and Adaptive Server Enterprise both use a default date of January 1, 1900. SQL Anywhere uses the current date.

Usage

If the milliseconds portion of a time is less than 3 digits, Adaptive Server Enterprise interprets the value differently depending on whether it was preceded by a period or a colon. If preceded by a colon, the value means thousandths of a second. If preceded by a period, 1 digit means tenths, 2 digits mean hundredths, and 3 digits mean thousandths. Sybase IQ and SQL Anywhere interpret the value the same way, regardless of the separator.

- Adaptive Server Enterprise converts the values below as shown.

```
12:34:56.7 to 12:34:56.700
12.34.56.78 to 12:34:56.780
12:34:56.789 to 12:34:56.789
12:34:56:7 to 12:34:56.007
12.34.56:78 to 12:34:56.078
12:34:56:789 to 12:34:56.789
```

- Sybase IQ converts the milliseconds value in the manner that Adaptive Server Enterprise does for values preceded by a period, in both cases:

```
12:34:56.7 to 12:34:56.700
12.34.56.78 to 12:34:56.780
12:34:56.789 to 12:34:56.789
12:34:56:7 to 12:34:56.700
12.34.56:78 to 12:34:56.780
12:34:56:789 to 12:34:56.789
```

See also

- *Usage for Date and Time Data Types* on page 84

Compatibility of Exported Dates

For dates in the first 9 days of a month and hours less than 10, Adaptive Server Enterprise supports a blank for the first digit; Sybase IQ supports a zero or a blank.

For details on how to load such data from Adaptive Server Enterprise into Sybase IQ, see *System Administration Guide: Volume 1 > Data Import and Export*.

Conversion of BIT to BINARY Data Type

Sybase IQ supports BIT to BINARY and BIT to VARBINARY implicit and explicit conversion and is compatible with Adaptive Server Enterprise support of these conversions.

Sybase IQ implicitly converts BIT to BINARY and BIT to VARBINARY data types for comparison operators, arithmetic operations, and **INSERT** and **UPDATE** statements.

For BIT to BINARY conversion, bit value 'b' is copied to the first byte of the binary string and the rest of the bytes are filled with zeros. For example, BIT value 1 is converted to BINARY(n) string 0x0100...00 having 2ⁿ nibbles. BIT value 0 is converted to BINARY string 0x00...00.

For BIT to VARBINARY conversion, BIT value 'b' is copied to the first byte of the BINARY string and the remaining bytes are not used; that is, only one byte is used. For example, BIT value 1 is converted to VARBINARY(n) string 0x01 having 2 nibbles.

The result of both implicit and explicit conversions of BIT to BINARY and BIT to VARBINARY data types is the same. The following table contains examples of BIT to BINARY and VARBINARY conversions.

Conversion of BIT value '1' to	Result
BINARY(3)	0x010000
VARBINARY(3)	0x01
BINARY(8)	0x0100000000000000
VARBINARY(8)	0x01

These examples illustrate both implicit and explicit conversion of BIT to BINARY and BIT to VARBINARY data types.

Given the following tables and data:

```
CREATE TABLE tbin(c1 BINARY(9))
CREATE TABLE tvarbin(c2 VARBINARY(9))
CREATE TABLE tbar(c2 BIT)

INSERT tbar VALUES(1)
INSERT tbar VALUES(0)
```

Implicit conversion of BIT to BINARY:

```
INSERT tbin SELECT c2 FROM tbar

c1
---
0x010000000000000000    (18 nibbles)
0x000000000000000000    (18 nibbles)
```

Implicit conversion of BIT to VARBINARY:

```
INSERT tvarbin SELECT c2 FROM tbar

c2
---
0x01
0x00
```

Explicit conversion of BIT to BINARY:

```
INSERT tbin SELECT CONVERT (BINARY(9), c2) FROM tbar

c1
---
0x010000000000000000    (18 nibbles)
0x000000000000000000    (18 nibbles)
```

SQL Data Types

Explicit conversion of BIT to VARBINARY:

```
INSERT tvarbin SELECT CONVERT(VARBINARY(9), c2) FROM tbar

c2
---
0x01
0x00
```

Conversion Between BIT and CHAR/VARCHAR Data Types

Sybase IQ supports implicit conversion between BIT and CHAR, and BIT and VARCHAR data types for comparison operators, arithmetic operations, and **INSERT** and **UPDATE** statements

.These examples illustrate both implicit and explicit conversions between BIT and CHAR, and BIT and VARCHAR data types.

Given the following tables and data:

```
CREATE TABLE tchar(c1 CHAR(9))
CREATE TABLE tvarchar(c2 VARCHAR(9))
CREATE TABLE tbar(c2 BIT)
CREATE TABLE tbit(c2 BIT)

INSERT tbar VALUES(1)
INSERT tbar VALUES(0)
```

Implicit conversion of BIT to VARCHAR / VARCHAR to BIT and implicit conversion of BIT to VARCHAR:

```
INSERT tvarchar SELECT c2 FROM tbar
SELECT c2, char_length(c2) FROM tvarchar

c2,char_length(tvarchar.c2)
-----
'1',1
'0',1
```

Implicit conversion of VARCHAR to BIT:

```
INSERT tbit SELECT c2 FROM tvarchar
SELECT c2 FROM tbit

c2
--
0
1
```

Explicit conversion of BIT to CHAR / CHAR to BIT and explicit conversion of BIT to CHAR:

```
INSERT tchar SELECT CONVERT (CHAR(9), c2) FROM tbar
SELECT c1, char_length(c1) FROM tchar

c1,char_length(tchar.c1)
```

```
-----
\1',9
\0',9
```

Explicit conversion of CHAR to BIT:

```
INSERT tbit SELECT CONVERT (BIT, c1) FROM tchar
SELECT c2 FROM tbit

c2
--
0
1
```

Explicit conversion of BIT to VARCHAR / VARCHAR to BIT and explicit conversion of BIT to VARCHAR:

```
INSERT tvvarchar SELECT CONVERT(VARCHAR(9), c2)
FROM tbar
SELECT c2, char_length(c2) FROM tvvarchar

c2,char_length(tvvarchar.c2)
-----
\1',1
\0',1
```

Explicit conversion of VARCHAR to BIT:

```
INSERT tbit SELECT CONVERT (BIT, c2) FROM tvvarchar
SELECT c2 FROM tbit

c2
--
0
1
```


SQL Functions

Functions return information from the database and are allowed anywhere an expression is allowed.

When using functions with Sybase IQ, unless otherwise stated, any function that receives the NULL value as a parameter returns a NULL value.

If you omit the **FROM** clause, or if all tables in the query are in the SYSTEM dbspace, SQL Anywhere processes the query, instead of Sybase IQ, and might behave differently, especially with regard to syntactic and semantic restrictions and the effects of option settings. See the SQL Anywhere documentation for rules that might apply to processing.

If you have a query that does not require a **FROM** clause, you can force Sybase IQ to process the query by adding the clause "FROM iq_dummy," where `iq_dummy` is a one-row, one-column table that you create in your database.

See also

- *Additional Functions* on page 351

Aggregate Functions

Aggregate functions summarize data over a group of rows from the database. The groups are formed using the **GROUP BY** clause of the **SELECT** statement.

Usage

Simple aggregate functions, such as **SUM()**, **MIN()**, **MAX()**, **AVG()** and **COUNT()** are allowed only in the select list and in the **HAVING** and **ORDER BY** clauses of a **SELECT** statement. These functions summarize data over a group of rows from the database. Groups are formed using the **GROUP BY** clause of the **SELECT** statement.

A new class of aggregate functions, called *window functions*, provides moving averages and cumulative measures that compute answers to queries such as, "What is the quarterly moving average of the Dow Jones Industrial average," or "List all employees and their cumulative salaries for each department."

- Simple aggregate functions, such as **AVG()**, **COUNT()**, **MAX()**, **MIN()**, and **SUM()** summarize data over a group of rows from the database. The groups are formed using the **GROUP BY** clause of the **SELECT** statement.
- Newer statistical aggregate functions that take one argument include **STDDEV()**, **STDDEV_SAMP()**, **STDDEV_POP()**, **VARIANCE()**, **VAR_SAMP()**, and **VAR_POP()**.

Both the simple and newer categories of aggregates can be used as a windowing function that incorporates a **<window clause>** in a SQL query specification (a *window*) that conceptually creates a moving window over a result set as it is processed.

Another class of window aggregate functions supports analysis of time series data. Like the simple aggregate and statistical aggregate functions, you can use these window aggregates with a SQL query specification (or *window-spec*). The time series window aggregate functions calculate correlation, linear regression, ranking, and weighted average results:

- ISO/ANSI SQL:2008 OLAP functions for time series analysis include: **CORR()**, **COVAR_POP()**, **COVAR_SAMP()**, **CUME_DIST()**, **FIRST_VALUE()**, **LAST_VALUE()**, **REGR_AVGX()**, **REGR_AVGY()**, **REGR_COUNT()**, **REGR_INTERCEPT()**, **REGR_R2()**, **REGR_SLOPE()**, **REGR_SXX()**, **REGR_SXY()**, and **REGR_SYY()**.
- Non-ISO/ANSI SQL:2008 OLAP aggregate function extensions used in the database industry include **FIRST_VALUE()**, **MEDIAN()**, and **LAST_VALUE()**.
- Weighted OLAP aggregate functions that calculate weighted moving averages include **EXP_WEIGHTED_AVG()** and **WEIGHTED_AVG()**.

Time series functions designed exclusively for financial time series forecasting and analysis have names beginning with “**TS_**”.

Note: Time series functionality is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide*.

For more information about using OLAP, see *System Administration Guide: Volume 2*.

For information on aggregate function support of the LONG BINARY and LONG VARCHAR data types, see *Unstructured Data Analytics in Sybase IQ*.

Table 18. Aggregate functions

Aggregate function	Parameters
AVG	([DISTINCT] { <i>column-name</i> <i>numeric-expr</i> })
CORR	(dependent-expression, independent-expression)
COUNT	(*)
COUNT	([DISTINCT] { <i>column-name</i> <i>numeric-expr</i> })
COVAR_POP	(dependent-expression, independent-expression)
COVAR_SAMP	(dependent-expression, independent-expression)
CUME_DIST	()
EXP_WEIGHTED_AVG	(<i>expression</i> , <i>period-expression</i>)
FIRST_VALUE	(expression)

Aggregate function	Parameters
LAST_VALUE	(expression)
LIST	(([DISTINCT] <i>string-expression</i> [, <i>'delimiter-string'</i>] [ORDER BY <i>order-by-expression</i> [ASC DESC], ...)
MAX	(([DISTINCT] { <i>column-name</i> <i>numeric-expr</i> })
MEDIAN	(expression)
MIN	(([DISTINCT] { <i>column-name</i> <i>numeric-expr</i> })
REGR_AVGX	(dependent-expression, independent-expression)
REGR_AVGY	(dependent-expression, independent-expression)
REGR_COUNT	(dependent-expression, independent-expression)
REGR_INTERCEPT	(dependent-expression, independent-expression)
REGR_R2	(dependent-expression, independent-expression)
REGR_SLOPE	(dependent-expression, independent-expression)
REGR_SXX	(dependent-expression, independent-expression)
REGR_SXY	(dependent-expression, independent-expression)
REGR_SYY	(dependent-expression, independent-expression)
STDDEV	([ALL] <i>expression</i>)
SUM	(([DISTINCT] { <i>column-name</i> <i>numeric-expr</i> })
VARIANCE	([ALL] <i>expression</i>)
WEIGHTED_AVG	(<i>expression</i> , <i>period-expression</i>)

The aggregate functions **AVG**, **SUM**, **STDDEV**, and **VARIANCE** do not support the binary data types (BINARY and VARBINARY).

See also

- *Analytical Functions* on page 101

Analytical Functions

Analytical functions include simple aggregates, window functions, and numeric functions.

- Simple aggregates — **AVG**, **COUNT**, **MAX**, **MIN**, and **SUM**, **STDDEV**, and **VARIANCE**

Note: You can use all simple aggregates, except the Grouping() function, with an OLAP windowed function.

- Window functions:
 - Windowing aggregates — **AVG, COUNT, MAX, MIN, and SUM.**
 - Ranking functions — **RANK, DENSE_RANK, PERCENT_RANK, ROW_NUMBER, and NTILE.**
 - Statistical functions — **STDDEV, STDDEV_SAMP, STDDEV_POP, VARIANCE, VAR_SAMP, and VAR_POP.**
 - Distribution functions — **PERCENTILE_CONT and PERCENTILE_DISC.**
 - Inter-row functions — **LAG and LEAD.**
- Numeric functions — **WIDTH_BUCKET, CEIL, and LN, EXP, POWER, SQRT, and FLOOR.**

Note: The ranking and inverse distribution analytical functions are not supported by Adaptive Server Enterprise.

This table lists the analytical functions and their parameters. Unlike some aggregate functions, you cannot specify **DISTINCT** in window functions.

Table 19. Analytical functions

Function	Parameters
AVG	({ <i>column-name</i> <i>numeric-expr</i> })
COUNT	(*)
COUNT	({ <i>column-name</i> <i>expression</i> })
DENSE_RANK	()
GROUPING *	({ GROUPING <i>group-by-expression</i> })
MAX	({ <i>column-name</i> <i>expression</i> })
MIN	({ <i>column-name</i> <i>expression</i> })
NTILE	(<i>integer</i>)
PERCENT_RANK	()
PERCENTILE_CONT	(<i>numeric-expr</i>)
PERCENTILE_DISC	(<i>numeric-expr</i>)
RANK	()
ROW_NUMBER	()
STDDEV	([ALL] <i>expression</i>)
STDDEV_POP	([ALL] <i>expression</i>)

Function	Parameters
STDDEV_SAMP	([ALL] <i>expression</i>)
SUM	({ <i>column-name</i> <i>expression</i> })
VAR_POP	([ALL] <i>expression</i>)
VAR_SAMP	([ALL] <i>expression</i>)
VARIANCE	([ALL] <i>expression</i>)

* The OLAP SQL standard allows Grouping() in **GROUP BY CUBE**, or **GROUP BY ROLLUP** operations only.

See also

- *Aggregate Functions* on page 99

Windowing Aggregate Function Usage

A major feature of the ISO/ANSI SQL extensions for OLAP is a construct called a *window*. This windowing extension let users divide result sets of a query (or a logical partition of a query) into groups of rows called partitions and determine subsets of rows to aggregate with respect to the current row.

You can use three classes of window functions with a window: ranking functions, the row numbering function, and window aggregate functions.

Windowing extensions specify a window function type over a window name or specification and are applied to partitioned result sets within the scope of a single query expression. A window partition is a subset of rows returned by a query, as defined by one or more columns in a special **OVER** clause:

```
OVER (PARTITION BY col1, col2...)
```

Windowing operations let you establish information such as the ranking of each row within its partition, the distribution of values in rows within a partition, and similar operations.

Windowing also lets you compute moving averages and sums on your data, enhancing the ability to evaluate your data and its impact on your operations.

A window partition is a subset of rows returned by a query, as defined by one or more columns in a special **OVER()** clause:

```
OVER (PARTITION BY col1, col2...)
```

For information on analytical function support of the LONG BINARY and LONG VARCHAR data types, see Unstructured Data Analytics in Sybase IQ > Function Support.

See also

- *CORR Function [Aggregate]* on page 148

- *COUNT Function [Aggregate]* on page 153
- *EXP_WEIGHTED_AVG Function [Aggregate]* on page 188
- *FIRST_VALUE Function [Aggregate]* on page 190
- *GROUPING Function [Aggregate]* on page 196
- *LAST_VALUE Function [Aggregate]* on page 215
- *MAX Function [Aggregate]* on page 230
- *MEDIAN Function [Aggregate]* on page 231
- *MIN Function [Aggregate]* on page 232
- *REGR_AVGX Function [Aggregate]* on page 266
- *REGR_COUNT Function [Aggregate]* on page 269
- *REGR_INTERCEPT Function [Aggregate]* on page 270
- *REGR_R2 Function [Aggregate]* on page 271
- *REGR_SLOPE Function [Aggregate]* on page 272
- *REGR_SXX Function [Aggregate]* on page 274
- *REGR_SXY Function [Aggregate]* on page 275
- *REGR_SYY Function [Aggregate]* on page 276
- *STDDEV Function [Aggregate]* on page 303
- *SUM Function [Aggregate]* on page 315
- *STDDEV_POP Function [Aggregate]* on page 305
- *STDDEV_SAMP Function [Aggregate]* on page 306
- *VAR_POP Function [Aggregate]* on page 335
- *VAR_SAMP Function [Aggregate]* on page 336
- *VARIANCE Function [Aggregate]* on page 338
- *WEIGHTED_AVG Function [Aggregate]* on page 341

Ranking Functions Usage

The OLAP ranking functions let application developers compose single-statement SQL queries that answer questions such as “Name the top 10 products shipped this year by total sales,” or “Give the top 5% of salespeople who sold orders to at least 15 different companies.”

These functions include the ranking functions, **RANK()**, **DENSE_RANK()**, **PERCENT_RANK()**, **ROW_NUMBER()**, and **NTILE()**.

Rank analytical functions rank items in a group, compute distribution, and divide a result set into a number of groupings. The rank analytical functions, **RANK()**, **DENSE_RANK()**, **PERCENT_RANK()**, **ROW_NUMBER()**, and **NTILE()** all require an **OVER (ORDER BY)** clause. For example:

```
RANK() OVER ( [PARTITION BY] ORDER BY <expression>
[ ASC | DESC ] )
```

The **ORDER BY** clause specifies the parameter on which ranking is performed and the order in which the rows are sorted in each group. This **ORDER BY** clause is used only within the **OVER**

clause and is *not* an **ORDER BY** for **SELECT**. No aggregation functions in the rank query **ROW** are allowed to specify **DISTINCT**.

Note: The **OVER (ORDER BY)** clause of the **ROW_NUMBER()** function cannot contain a **ROWS** or **RANGE** clause.

The **OVER** clause indicates that the function operates on a query result set. The result set is the rows that are returned after the **FROM**, **WHERE**, **GROUP BY**, and **HAVING** clauses have all been evaluated. The **OVER** clause defines the data set of the rows to include in the computation of the rank analytical function.

The value *expression* is a sort specification that can be any valid expression involving a column reference, aggregates, or expressions invoking these items.

The **ASC** or **DESC** parameter specifies the ordering sequence as ascending or descending. Ascending order is the default.

Rank analytical functions are only allowed in the select list of a **SELECT** or **INSERT** statement or in the **ORDER BY** clause of the **SELECT** statement. Rank functions can be in a view or a union. You cannot use rank functions in a subquery, a **HAVING** clause, or in the select list of an **UPDATE** or **DELETE** statement. More than one rank analytical function is allowed per query in Sybase IQ 15.3.

Statistical Aggregate Analytic Function Usage

Statistical aggregate analytic functions summarize data over a group of rows from the database.

The groups are formed using the **GROUP BY** clause of the **SELECT** statement. Aggregate functions are allowed only in the select list and in the **HAVING** and **ORDER BY** clauses of a **SELECT** statement. These functions include **STDDEV**, **STDDEV_POP**, **STDDEV_SAMP**, **VARIANCE**, **VAR_POP**, and **VAR_SAMP**.

The OLAP functions can be used as a window function with an **OVER()** clause in a SQL query specification that conceptually creates a moving window over a result set as it is processed.

Distribution Functions Usage

The inverse distribution analytical functions **PERCENTILE_CONT** and **PERCENTILE_DISC** take a percentile value as the function argument and operate on a group of data specified in the **WITHIN GROUP** clause, or operate on the entire data set.

These functions return one value per group. For **PERCENTILE_DISC**, the data type of the results is the same as the data type of its **ORDER BY** item specified in the **WITHIN GROUP** clause. For **PERCENTILE_CONT**, the data type of the results is either numeric, if the **ORDER BY** item in the **WITHIN GROUP** clause is a numeric, or double, if the **ORDER BY** item is an integer or floating-point.

The inverse distribution analytical functions require a **WITHIN GROUP (ORDER BY)** clause. For example:

SQL Functions

```
PERCENTILE_CONT ( expression1 ) WITHIN GROUP ( ORDER BY expression2  
[ASC | DESC ] )
```

The value of *expression1* must be a constant of numeric data type and range from 0 to 1 (inclusive). If the argument is NULL, then a “wrong argument for percentile” error is returned. If the argument value is less than 0, or greater than 1, then a “data value out of range” error is returned.

The **ORDER BY** clause, which must be present, specifies the expression on which the percentile function is performed and the order in which the rows are sorted in each group. This **ORDER BY** clause is used only within the **WITHIN GROUP** clause and is *not* an **ORDER BY** for the **SELECT**.

The **WITHIN GROUP** clause distributes the query result into an ordered data set from which the function calculates a result.

The value *expression2* is a sort specification that must be a single expression involving a column reference. Multiple expressions are not allowed and no rank analytical functions, set functions, or subqueries are allowed in this sort expression.

The ASC or DESC parameter specifies the ordering sequence as ascending or descending. Ascending order is the default.

Inverse distribution analytical functions are allowed in a subquery, a **HAVING** clause, a view, or a union. The inverse distribution functions can be used anywhere the simple non analytical aggregate functions are used. The inverse distribution functions ignore the NULL value in the data set.

Inter-Row Functions Usage

The inter-row functions **LAG** and **LEAD** enable access to previous values or subsequent values in a data series.

These functions provide access to more than one row of a table or partition simultaneously without a self join. The **LAG** function provides access to a row at a given physical offset prior to the **CURRENT ROW** in the table or partition. The **LEAD** function provides access to a row at a given physical offset after the **CURRENT ROW** in the table or partition. Use the **LAG** and **LEAD** functions to create queries such as “What was the stock price two intervals before the current row,” and “What was the stock price one interval after the current row.” See *System Administration Guide: Volume 2 > Using OLAP > Analytical Functions > Windowing > Interrow Functions*.

Inter-row functions require an **OVER (ORDER_BY)** clause.

Data Type Conversion Functions

Data type conversion functions convert arguments from one data type to another.

This table lists the data type conversion functions and their parameters.

Table 20. Data type conversion functions

Data type conversion function	Parameters
BIGINTTOHEX	(<i>integer-expression</i>)
CAST	(<i>expression AS data type</i>)
CONVERT	(<i>data type, expression [, format-style]</i>)
HEXTOBIGINT	(<i>hexadecimal-string</i>)
HEXTOINT	(<i>hexadecimal-string</i>)
INTTOHEX	(<i>integer-expr</i>)
ISDATE	(<i>string</i>)
ISNUMERIC	(<i>string</i>)

Description

The database server carries out many type conversions automatically. For example, if a string is supplied where a numerical expression is required, the string is automatically converted to a number.

See also

- *Data Type Conversions* on page 93
- *Storage Size* on page 81

Date and Time Functions

Date and time functions perform conversion, extraction, or manipulation operations on date and time data types and can return date and time information.

These tables list the date and time functions and their parameters.

*Syntax 1***Table 21. Date and time functions**

Date and time functions	Parameters
DATE	(<i>expression</i>)
DATECEILING	(<i>date-part, datetime-expr, [multiple-expr]</i>)
DATEFLOOR	(<i>date-part, datetime-expr, [multiple-expr]</i>)
DATEFORMAT	(<i>datetime-expr, string-expr</i>)

Date and time functions	Parameters
DATENAME	(<i>date-part</i> , <i>date-expr</i>)
DATEROUND	(<i>date-part</i> , <i>datetime-expr</i> , [<i>multiple-expr</i>])
DATETIME	(<i>expression</i>)
DAY	(<i>date-expr</i>)
DAYNAME	(<i>date-expr</i>)
DAYS	(<i>date-expr</i>)
DAYS	(<i>date-expr</i> , <i>date-expr</i>)
DAYS	(<i>date-expr</i> , <i>integer-expr</i>)
DOW	(<i>date-expr</i>)
HOUR	(<i>datetime-expr</i>)
HOURS	(<i>datetime-expr</i>)
HOURS	(<i>datetime-expr</i> , <i>datetime-expr</i>)
HOURS	(<i>datetime-expr</i> , <i>integer-expr</i>)
ISDATE	(<i>string</i>)
MINUTE	(<i>datetime-expr</i>)
MINUTES	(<i>datetime-expr</i>)
MINUTES	(<i>datetime-expr</i> , <i>datetime-expr</i>)
MINUTES	(<i>datetime-expr</i> , <i>integer-expr</i>)
MONTH	(<i>date-expr</i>)
MONTHNAME	(<i>date-expr</i>)
MONTHS	(<i>date-expr</i>)
MONTHS	(<i>date-expr</i> , <i>date-expr</i>)
MONTHS	(<i>date-expr</i> , <i>integer-expr</i>)
NOW	(*)
QUARTER	(<i>date-expr</i>)
SECOND	(<i>datetime-expr</i>)
SECONDS	(<i>datetime-expr</i>)

Date and time functions	Parameters
SECONDS	(<i>datetime-expr</i> , <i>datetime-expr</i>)
SECONDS	(<i>datetime-expr</i> , <i>integer-expr</i>)
TODAY	(*)
WEEKS	(<i>date-expr</i>)
WEEKS	(<i>date-expr</i> , <i>date-expr</i>)
WEEKS	(<i>date-expr</i> , <i>integer-expr</i>)
YEAR	(<i>date-expr</i>)
YEARS	(<i>date-expr</i>)
YEARS	(<i>date-expr</i> , <i>date-expr</i>)
YEARS	(<i>date-expr</i> , <i>integer-expr</i>)
YMD	(<i>year-num</i> , <i>month-num</i> , <i>day-num</i>)

Syntax 2

Table 22. Transact-SQL-compatible date and time functions

Transact-SQL compatible date and time functions	Parameters
DATEADD	(<i>date-part</i> , <i>numeric-expression</i> , <i>date-expr</i>)
DATEDIFF	(<i>date-part</i> , <i>date-expr1</i> , <i>date-expr2</i>)
DATENAME	(<i>date-part</i> , <i>date-expr</i>)
DATEPART	(<i>date-part</i> , <i>date-expr</i>)
GETDATE	()

Description

Sybase IQ provides two classes of date and time functions that can be used interchangeably, but have different styles. One set is Transact-SQL-compatible.

The date and time functions listed for Syntax 1 allow manipulation of time units. Most time units (such as MONTH) have four functions for time manipulation, although only two names are used (such as MONTH and MONTHS).

The functions listed for Syntax 2 are the Transact-SQL date and time functions. They allow an alternative way of accessing and manipulating date and time functions.

SQL Functions

You should convert arguments to date functions to dates before used them. For example, this is incorrect:

```
days ( '1995-11-17', 2 )
```

This is correct:

```
days ( date( '1995-11-17' ), 2 )
```

Sybase IQ does not have the same constants or data type promotions as SQL Anywhere, with which it shares a common user interface. If you issue a **SELECT** statement without a **FROM** clause, the statement is passed to SQL Anywhere. The following statement is handled exclusively by SQL Anywhere:

```
SELECT WEEKS( '1998/11/01' );
```

The following statement, processed by Sybase IQ, uses a different starting point for the **WEEKS** function and returns a different result than the statement above:

```
SELECT WEEKS( '1998/11/01' ) FROM iq_dummy;
```

Consider another example. The **MONTHS** function returns the number of months since an “arbitrary starting date.” The “arbitrary starting date” of Sybase IQ, the imaginary date 0000-01-01, is chosen to produce the most efficient date calculations and is consistent across various data parts. SQL Anywhere does not have a single starting date. The following statements, the first processed by SQL Anywhere, the second by Sybase IQ, both return the answer 12:

```
SELECT MONTHS( '0001/01/01' );
```

```
SELECT MONTHS( '0001/01/01' ) FROM iq_dummy;
```

However, also consider these statements:

```
SELECT DAYS( '0001/01/01' );
```

```
SELECT DAYS( '0001/01/01' ) FROM iq_dummy;
```

The first, processed by SQL Anywhere, yields the value 307, but the second, processed by Sybase IQ, yields 166.

For the most consistent results, therefore, always include the table name in the **FROM** clause whether you need it or not.

Note: Create a dummy table with only one column and row. You can then reference this table in the **FROM** clause for any **SELECT** statement that uses date or time functions, thus ensuring processing by Sybase IQ, and consistent results.

Date Parts

Many of the date functions use dates built from date parts.

This table displays allowed values of *date-part*.

Table 23. Date part values

Date part	Abbreviation	Values
Year	yy	0001 – 9999
Quarter	qq	1 – 4
Month	mm	1 – 12
Week	wk	1 – 54
Day	dd	1 – 31
Dayofyear	dy	1 – 366
Weekday	dw	1 – 7 (Sun. – Sat.)
Hour	hh	0 – 23
Minute	mi	0 – 59
Second	ss	0 – 59
Millisecond	ms	0 – 999
Microsecond	mcs or us	0 – 999999
Calyearofweek	cyr	Integer. The year in which the week begins. The week containing the first few days of the year can be part of the last week of the previous year, depending upon which day it begins. If the new year starts on a Thursday through Saturday, its first week starts on the last Sunday of the previous year. If the new year starts on a Sunday through Wednesday, none of its days are part of the previous year.
Calweekofyear	cwk	An integer from 1 to 54 representing the week number within the year that contains the specified date.
Caldayofweek	cdw	The day number within the week (Sunday = 1, Saturday = 7).

Note: By default, Sunday is the first day of the week. To make Monday be the first day, use:

```
set option 'Date_First_Day_of_Week' = '1'
```

For more information on specifying which day is the first day of the week, see *Reference: Statements and Options > Database Options > Alphabetical List of Options > DATE_FIRST_DAY_OF_WEEK Option.*

Compatibility

For compatibility with Adaptive Server Enterprise, use the Transact-SQL date and time functions.

See also

- *DATEADD Function [Date and Time]* on page 157
- *DATECEILING Function [Date and Time]* on page 158
- *DATEDIFF Function [Date and Time]* on page 161
- *DATEFLOOR Function [Date and Time]* on page 164
- *DATEPART Function [Date and Time]* on page 169
- *DATENAME Function [Date and Time]* on page 168
- *DATEROUND Function [Date and Time]* on page 170

HTTP Functions

HTTP functions facilitate the handling of HTTP requests within Web services.

This table lists all HTTP functions and their parameters.

Table 24. HTTP functions

HTTP function	Parameters
HTML_DECODE	(<i>string</i>)
HTML_ENCODE	(<i>string</i>)
HTTP_DECODE	(<i>string</i>)
HTTP_ENCODE	(<i>string</i>)
HTTP_HEADER	(<i>header-field-name</i>)
HTTP_VARIABLE	(<i>var-name</i> [[, <i>instance</i>], <i>header-field</i>])
NEXT_HTTP_HEADER	(<i>header-name</i>)
NEXT_HTTP_VARIABLE	(<i>var-name</i>)

Numeric Functions

Numeric functions perform mathematical operations on numerical data types or return numeric information.

Function

Sybase IQ does not have the same constants or data type promotions as SQL Anywhere, with which it shares a common user interface. If you issue a **SELECT** statement without a **FROM** clause, the statement is passed through to SQL Anywhere. For the most consistent results, include the table name in the **FROM** clause whether you need it or not.

Note: Consider creating a dummy table to use in such cases.

This table lists numeric functions and their parameters.

Table 25. Numeric functions

Numeric function	Parameters
ABS	(<i>numeric-expr</i>)
ACOS	(<i>numeric-expr</i>)
ASIN	(<i>numeric-expr</i>)
ATAN	(<i>numeric-expr</i>)
ATAN2	(<i>numeric-expr1</i> , <i>numeric-expr2</i>)
CEIL	(<i>numeric-expr</i>)
CEILING	(<i>numeric-expr</i>)
COS	(<i>numeric-expr</i>)
COT	(<i>numeric-expr</i>)
DEGREES	(<i>numeric-expr</i>)
EXP	(<i>numeric-expr</i>)
FLOOR	(<i>numeric-expr</i>)
LN	(<i>numeric-expr</i>)
LOG	(<i>numeric-expr</i>)
LOG10	(<i>numeric-expr</i>)
MOD	(<i>dividend</i> , <i>divisor</i>)
PI	(*)
POWER	(<i>numeric-expr1</i> , <i>numeric-expr2</i>)
RADIANS	(<i>numeric-expr</i>)
RAND	([<i>integer-expr</i>])
REMAINDER	(<i>numeric-expr</i> , <i>numeric-expr</i>)
ROUND	(<i>numeric-expr</i> , <i>integer-expr</i>)
SIGN	(<i>numeric-expr</i>)
SIN	(<i>numeric-expr</i>)
SQRT	(<i>numeric-expr</i>)

Numeric function	Parameters
SQUARE	(<i>numeric-expr</i>)
TAN	(<i>numeric-expr</i>)
“TRUNCATE”	(<i>numeric-expr, integer-expr</i>)
TRUNCNUM	(<i>numeric-expression, integer-expression</i>)
WIDTH_BUCKET	(<i>expression, min_value, max_value, num_buckets</i>)

String Functions

String functions perform conversion, extraction, or manipulation operations on strings, or return information about strings.

When working in a multibyte character set, check carefully whether the function being used returns information concerning characters or bytes.

Most of the string functions accept binary data (hexadecimal strings) in the *string-expr* parameter, but some of the functions, such as **LCASE**, **UCASE**, **LOWER**, and **LTRIM**, expect the string expression to be a character string.

Unless you supply a constant **LENGTH** argument to a function that produces a LONG VARCHAR result (such as **SPACE** or **REPEAT**), the default length is the maximum allowed.

Sybase IQ queries containing one or more of these functions might return one of the following errors:

```
ASA Error -1009080: Key doesn't fit on a single database page:
65560(4, 1)
```

```
ASA Error -1009119: Record size too large for database page size
```

For example:

```
SELECT COUNT(*) FROM test1 a WHERE (a.col1 + SPACE(4-LENGTH(a.col1))
+ a.col2 + space(2- LENGTH(a.col2))) IN (SELECT (b.col3) FROM test1
b);
```

To avoid such errors, cast the function result with an appropriate maximum length; for example:

```
SELECT COUNT(*) FROM test1 a WHERE (a.col1 + CAST(SPACE(4-
LENGTH(a.col1)) AS VARCHAR(4)) + a.col2 + CAST(SPACE(2-LENGTH
(a.col2)) AS VARCHAR(4))) IN (SELECT (b.col3) FROM test1 b);
```

The errors are more likely with an IQ page size of 64K or a multibyte collation.

This table lists string functions and their parameters.

Table 26. String functions

String function	Parameters
ASCII	(<i>string-expr</i>)
BIT_LENGTH	(<i>column-name</i>)
BYTE_LENGTH	(<i>string-expr</i>)
CHAR	(<i>integer-expr</i>)
CHAR_LENGTH	(<i>string-expr</i>)
CHARINDEX	(<i>string-expr1</i> , <i>string-expr2</i>)
DIFFERENCE	(<i>string-expr1</i> , <i>string-expr2</i>)
GRAPHICAL_PLAN	(<i>string-expr</i>)
HTML_PLAN	(<i>string-expr</i>)
INSERTSTR	(<i>numeric-expr</i> , <i>string-expr1</i> , <i>string-expr2</i>)
LCASE	(<i>string-expr</i>)
LEFT	(<i>string-expr</i> , <i>numeric-expr</i>)
LEN	(<i>string-expr</i>)
LENGTH	(<i>string-expr</i>)
LOCATE	(<i>string-expr1</i> , <i>string-expr2</i> [, <i>numeric-expr</i>])
LOWER	(<i>string-expr</i>)
LTRIM	(<i>string-expr</i>)
OCTET_LENGTH	(<i>column-name</i>)
PATINDEX	('% <i>pattern</i> %' , <i>string_expr</i>)
REPEAT	(<i>string-expr</i> , <i>numeric-expr</i>)
REPLACE	(<i>original-string</i> , <i>search-string</i> , <i>replace-string</i>)
REVERSE	(<i>expression</i> <i>uchar_expr</i>)
REPLICATE	(<i>string-expr</i> , <i>integer-expr</i>)
RIGHT	(<i>string-expr</i> , <i>numeric-expr</i>)
RTRIM	(<i>string-expr</i>)
SIMILAR	(<i>string-expr1</i> , <i>string-expr2</i>)
SORTKEY	(<i>string-expression</i> [, { <i>collation-id</i> <i>collation-name</i> [(<i>collation-tailoring-string</i>)] }])

String function	Parameters
SOUNDEX	(<i>string-expr</i>)
SPACE	(<i>integer-expr</i>)
STR	(<i>numeric_expr</i> [, <i>length</i> [, <i>decimal</i>]])
STR_REPLACE	(<i>string_expr1</i> , <i>string_expr2</i> , <i>string_expr3</i>)
STRING	(<i>string1</i> [, <i>string2</i> , ..., <i>string99</i>])
STUFF	(<i>string-expr1</i> , <i>start</i> , <i>length</i> , <i>string-expr2</i>)
SUBSTRING	(<i>string-expr</i> , <i>integer-expr</i> [, <i>integer-expr</i>])
TRIM	(<i>string-expr</i>)
UCASE	(<i>string-expr</i>)
UPPER	(<i>string-expr</i>)

For information on string functions that support the LONG BINARY and LONG VARCHAR data types, see *Unstructured Data Analytics in Sybase IQ > Function Support*.

See also

- *Physical Limitations* on page 353

System Functions

System functions return system information.

This table lists the system functions and their parameters.

Table 27. System functions

System function	Parameters
COL_LENGTH	(<i>table-name</i> , <i>column-name</i>)
COL_NAME	(<i>table-id</i> , <i>column-id</i> [, <i>database-id</i>])
CONNECTION_PROPERTY	({ <i>property-id</i> <i>property-name</i> } [, <i>connection-id</i>])
DATALENGTH	(<i>expression</i>)
DB_ID	([<i>database-name</i>])
DB_NAME	([<i>database-id</i>])

System function	Parameters
DB_PROPERTY	(({ <i>property-id</i> <i>property-name</i> } [, { <i>database-id</i> <i>database-name</i> }]))
EVENT_CONDITION	(<i>condition-name</i>)
EVENT_CONDITION_NAME	(<i>integer</i>)
EVENT_PARAMETER	(<i>context-name</i>)
GROUP_MEMBER	(<i>group-name-string-expression</i> [, <i>user-name-string-expression</i>])
INDEX_COL	(<i>table-name</i> , <i>index-id</i> , <i>key_#</i> [, <i>user-id</i>])
NEXT_CONNECTION	([<i>connection-id</i>] [, <i>database-id</i>])
NEXT_DATABASE	(({ NULL <i>database-id</i> }))
OBJECT_ID	(<i>object-name</i>)
OBJECT_NAME	(<i>object-id</i> [, <i>database-id</i>])
PROPERTY	(({ <i>property-id</i> <i>property-name</i> }))
PROPERTY_DESCRIPTION	(<i>property-id</i> <i>property-name</i>)
PROPERTY_NAME	(<i>property-id</i>)
PROPERTY_NUMBER	(<i>property-name</i>)
SUSER_ID	([<i>user-name</i>])
SUSER_NAME	([<i>user-id</i>])
USER_ID	([<i>user-name</i>])
USER_NAME	([<i>user-id</i>])

Description

Databases currently running on a server are identified by a database name and a database ID number. The `db_id` and `db_name` functions provide information on these values.

A set of system functions provides information about properties of a currently running database, or of a connection, on the database server. These system functions take the database name or ID, or the connection name, as an optional argument to identify the database or connection for which the property is requested.

Performance

System functions are processed differently than other Sybase IQ functions. When queries to Sybase IQ tables include system functions, performance is reduced.

Compatibility

This table shows the Adaptive Server Enterprise system functions and their status in Sybase IQ:

Table 28. Status of ASE system functions in Sybase IQ

Function	Status
col_length	Implemented
col_name	Implemented
db_id	Implemented
db_name	Implemented
index_col	Implemented
object_id	Implemented
object_name	Implemented
proc_role	Always returns 0
show_role	Always returns NULL
tsequal	Not implemented
user_id	Implemented
user_name	Implemented
suser_id	Implemented
suser_name	Implemented
datalength	Implemented
curunreservedpgs	Not implemented
data_pgs	Not implemented
host_id	Not implemented
host_name	Not implemented
lct_admin	Not implemented
reserved_pgs	Not implemented
rowcnt	Not implemented
used_pgs	Not implemented

Function	Status
<code>valid_name</code>	Not implemented
<code>valid_user</code>	Not implemented

Notes

- Some of the system functions are implemented in Sybase IQ as system stored procedures.
- The `db_id`, `db_name`, `datalength`, `suser_id`, and `suser_name` functions are implemented as built-in functions.

Connection Properties

Retrieve the value of a specific connection property or the values of all connection properties.

Connection properties apply to an individual connection. For descriptions of all connection properties, see *SQL Anywhere Server - Database Administration > Configuring your database > Connection, database, and database server properties > Connection properties*. The connection properties **QueryBypassedCosted**, **QueryBypassedOptimized**, **QueryDescribedOptimizer**, and **StatementPostAnnotatesSimple** apply to SQL Anywhere objects but not to IQ tables.

Note: This reference points to SQL Anywhere documentation.

Examples

Use the `connection_property` system function to retrieve the value of a connection property. The following statement returns the number of pages that have been read from file by the current connection:

```
select connection_property ( 'DiskRead' )
```

Use the `sa_conn_properties` system procedure to retrieve the values of all connection properties.

```
call sa_conn_properties
```

A separate row appears for each connection, for each property.

See also

- *PROPERTY Function [System]* on page 259
- *PROPERTY_NAME Function [System]* on page 261
- *PROPERTY_NUMBER Function [System]* on page 261
- *CONNECTION_PROPERTY Function [System]* on page 142

Properties Available for the Server

Retrieve the value of a specific server property or the values of all server properties.

Server properties apply across the server as a whole. For descriptions of all server properties, see *SQL Anywhere Server - Database Administration > Configuring your database > Connection, database, and database server properties > Database server properties*.

The Server Edition property returns the SQL Anywhere edition, not the Sybase IQ edition. To show Sybase IQ license information, use the **sp_iqlmconfig** system procedure.

The Server Edition property returns the SQL Anywhere edition, not the Sybase IQ edition. To show Sybase IQ license information, use the **sp_iqlmconfig** system procedure.

Examples

Use the `property` system function to retrieve the value of a server property. The following statement returns the number of cache pages being used to hold the main heap:

```
select property ( 'MainHeapPages' ) from iq_dummy
```

Use the `sa_eng_properties` system procedure to retrieve the values of all server properties.

```
call sa_eng_properties
```

See also

- *sp_iqlmconfig Procedure* on page 452
- *PROPERTY Function [System]* on page 259
- *PROPERTY_NAME Function [System]* on page 261
- *PROPERTY_NUMBER Function [System]* on page 261
- *CONNECTION_PROPERTY Function [System]* on page 142

Properties Available for each Database

You can retrieve the value of a specific database property or the values of all database properties. Database properties apply to an entire database.

For descriptions of all of the database properties, see *SQL Anywhere Server - Database Administration > Configuring your database > Connection, database, and database server properties > Database server properties*. The server properties **QueryBypassedCosted**, **QueryBypassedOptimized**, **QueryDescribedOptimizer**, and **StatementPostAnnotatesSimple** are updated only for queries against SQL Anywhere tables.

Examples

Use the `db_property` system function to retrieve the value of a database property. The following statement returns the page size of the current database:

```
select db_property ( 'PageSize' )
```

Use the `sa_db_properties` system procedure to retrieve the values of all database properties:

```
call sa_db_properties
```

See also

- *PROPERTY* Function [System] on page 259
- *PROPERTY_NAME* Function [System] on page 261
- *PROPERTY_NUMBER* Function [System] on page 261
- *CONNECTION_PROPERTY* Function [System] on page 142

SQL and Java User-Defined Functions

There are two mechanisms for creating user-defined functions in Sybase IQ. You can use the SQL language to write the function, or you can use Java.

Note: User-defined functions are processed by SQL Anywhere. They do not take advantage of the performance features of Sybase IQ. Queries that include user-defined functions run at least 10 times slower than queries without them.

In very few cases, differences in semantics between SQL Anywhere and Sybase IQ can produce different results for a query if it is issued in a user-defined function. For example, Sybase IQ treats the CHAR and VARCHAR data types as distinct and different, while SQL Anywhere treats CHAR data as if it were VARCHAR.

User-Defined Functions in SQL

You can implement your own functions in SQL using the **CREATE FUNCTION** statement. The **RETURN** statement inside the **CREATE FUNCTION** statement determines the data type of the function.

Once you have created a SQL user-defined function, you can use it anywhere a built-in function of the same data type is used.

Note: Avoid using the CONTAINS predicate in a view that has a user-defined function, because the CONTAINS criteria is ignored. Use the LIKE predicate instead, or issue the query outside of a view.

For more information on creating SQL functions, see *System Administration Guide: Volume 2* > Using Procedures and Batches.

User-Defined Functions in Java

Although SQL functions are useful, Java classes provide a more powerful and flexible way of implementing user-defined functions, with the additional advantage that you can move them from the database server to a client application if desired.

Any *class method* of an installed Java class can be used as a user-defined function anywhere a built-in function of the same data type is used.

Instance methods are tied to particular instances of a class, and so have different behavior from standard user-defined functions.

For more information on creating Java classes, and on class methods, see *SQL Anywhere Server - Programming > Java in the database*.

Time Series and Forecasting Functions

Time series and forecasting functions are designed exclusively for financial time series analysis. Time series and forecasting functions are documented in the Time Series Guide.

Note: Time series and forecasting capability is available only with RAP – The Trading Edition Enterprise.

Miscellaneous Functions

Miscellaneous functions perform operations on arithmetic, string, or date/time expressions, including the return values of other functions.

This table lists the miscellaneous functions and their parameters.

Table 29. Miscellaneous functions

Miscellaneous functions	Parameters
ARGN	(<i>integer-expr</i> , <i>expression</i> [, ...])
COALESCE	(<i>expression</i> , <i>expression</i> [, <i>expression</i> ...])
IFNULL	(<i>expression1</i> , <i>expression2</i> [, <i>expression3</i>])
ISNULL	(<i>expression</i> , <i>expression</i> [, <i>expression</i> ...])
NULLIF	(<i>expression1</i> , <i>expression2</i>)
NUMBER	(*)
ROWID	(<i>table-name</i>)

Compatibility

Adaptive Server Enterprise supports only the COALESCE, ISNULL, and NULLIF functions.

Alphabetical List of Functions

This section describes each SQL function individually.

The function type, for example, Numeric or String, is indicated in brackets next to the function name.

Some of the results in the examples have been rounded or truncated.

The actual values of database object IDs, such as the object ID of a table or the column ID of a column, might differ from the values shown in the examples.

ABS Function [Numeric]

Returns the absolute value of a numeric expression.

Syntax

```
ABS ( numeric-expression )
```

Parameters

Parameter	Description
numeric-expression	The number whose absolute value is to be returned.

Returns

An absolute value of the numeric expression.

Numeric-expression data type	Returns
INT	INT
FLOAT	FLOAT
DOUBLE	DOUBLE
NUMERIC	NUMERIC

Example

The following statement returns the value 66:

```
SELECT ABS( -66 ) FROM iq_dummy
```

Standards and compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

ACOS Function [Numeric]

Returns the arc-cosine, in radians, of a numeric expression.

Syntax

```
ACOS ( numeric-expression )
```

*Parameters***Table 30. Parameters**

Parameter	Description
numeric-expression	The cosine of the angle.

Returns

DOUBLE

Example

The following statement returns the value 1.023945:

```
SELECT ACOS( 0.52 ) FROM iq_dummy
```

Standards and compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *COS Function [Numeric]* on page 149
- *ATAN2 Function [Numeric]* on page 128
- *ATAN Function [Numeric]* on page 127
- *ASIN Function [Numeric]* on page 126
- *COT Function [Numeric]* on page 150
- *SIN Function [Numeric]* on page 293
- *TAN Function [Numeric]* on page 318

ARGN Function [Miscellaneous]

Returns a selected argument from a list of arguments.

Syntax

```
ARGN ( integer-expression, expression [ , ... ] )
```

Parameters

Table 31. Parameters

Parameter	Description
<i>integer-expression</i>	The position of an argument within a list of expressions.
<i>expression</i>	<i>n</i> expression of any data type passed into the function. All supplied expressions must be of the same data type.

Returns

Using the value of the *integer-expression* as *n*, returns the *n*th argument (starting at 1) from the remaining list of arguments.

Example

The following statement returns the value 6:

```
SELECT ARGN( 6, 1,2,3,4,5,6 ) FROM iq_dummy
```

Usage

Using the value of *integer-expression* as *n* returns the *n*th argument (starting at 1) from the remaining list of arguments. While the expressions can be of any data type, they must all be of the same data type. The integer expression must be from one to the number of expressions in the list or NULL is returned. Multiple expressions are separated by a comma.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

ASCII Function [String]

Returns the integer ASCII value of the first byte in a string-expression.

Syntax

```
ASCII ( string-expression )
```

*Parameters***Table 32. Parameters**

Parameter	Description
string-expression	The string.

Returns

SMALLINT

Example

The following statement returns the value 90, when the collation sequence is set to the default ISO_BINENG:

```
SELECT ASCII( 'Z' ) FROM iq_dummy
```

Usage

If the string is empty, **ASCII** returns zero. Literal strings must be enclosed in quotes.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

ASIN Function [Numeric]

Returns the arc-sine, in radians, of a number.

Syntax

```
ASIN ( numeric-expression )
```

*Parameters***Table 33. Parameters**

Parameter	Description
numeric-expression	The sine of the angle.

Returns

DOUBLE

Example

The following statement returns the value 0.546850.

```
SELECT ASIN( 0.52 ) FROM iq_dummy
```


Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *COS Function [Numeric]* on page 149
- *ATAN2 Function [Numeric]* on page 128
- *ATAN Function [Numeric]* on page 127
- *ACOS Function [Numeric]* on page 124
- *COT Function [Numeric]* on page 150
- *SIN Function [Numeric]* on page 293
- *TAN Function [Numeric]* on page 318

ATAN Function [Numeric]

Returns the arc-tangent, in radians, of a number.

Syntax

```
ATAN ( numeric-expression )
```

*Parameters***Table 34. Parameters**

Parameter	Description
numeric-expression	The tangent of the angle.

Returns

DOUBLE

Example

The following statement returns the value 0.479519:

```
SELECT ATAN( 0.52 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *COS Function [Numeric]* on page 149
- *ATAN2 Function [Numeric]* on page 128

SQL Functions

- *ASIN Function [Numeric]* on page 126
- *ACOS Function [Numeric]* on page 124
- *COT Function [Numeric]* on page 150
- *SIN Function [Numeric]* on page 293
- *TAN Function [Numeric]* on page 318

ATAN2 Function [Numeric]

Returns the arc-tangent, in radians, of the ratio of two numbers.

Syntax

```
ATAN2 ( numeric-expression1, numeric-expression2 )
```

Parameters

Parameter	Description
<i>numeric-expression1</i>	The numerator in the ratio whose arc tangent is calculated.
<i>numeric-expression2</i>	The denominator in the ratio whose arc-tangent is calculated.

Returns

DOUBLE

Example

The following statement returns the value 0.00866644968879073143:

```
SELECT ATAN2( 0.52, 060 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—ATAN2 is not supported by Adaptive Server Enterprise.

See also

- *COS Function [Numeric]* on page 149
- *ATAN Function [Numeric]* on page 127
- *ASIN Function [Numeric]* on page 126
- *ACOS Function [Numeric]* on page 124
- *COT Function [Numeric]* on page 150
- *SIN Function [Numeric]* on page 293
- *TAN Function [Numeric]* on page 318

AVG Function [Aggregate]

Computes the average of a numeric expression for a set of rows, or computes the average of a set of unique values.

Syntax

```
AVG ( numeric-expression | DISTINCT column-name )
```

Parameters

Table 35. Parameters

Parameter	Description
numeric-expression	The value whose average is calculated over a set of rows.
DISTINCT column-name	Computes the average of the unique values in column-name. This is of limited usefulness, but is included for completeness

Returns

Returns the NULL value for a group containing no rows.

Returns DOUBLE if the argument is DOUBLE, otherwise NUMERIC.

Example

The following statement returns the value 49988.6:

```
SELECT AVG ( salary ) FROM Employees
```

Usage

This average does not include rows where *numeric-expression* is the NULL value. Returns the NULL value for a group containing no rows.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *COUNT Function [Aggregate]* on page 153
- *SUM Function [Aggregate]* on page 315

BFILE Function [Data extraction]

Extracts individual LONG BINARY and LONG VARCHAR cells to individual operating system files on the server.

Usage

The IQ data extraction facility includes the **BFILE** function, which allows you to extract individual LONG BINARY and LONG VARCHAR cells to individual operating system files on the server. You can use **BFILE** with or without the data extraction facility.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

BIGINTTOHEX Function [Data Type Conversion]

Returns the hexadecimal equivalent in VARCHAR (16) of a decimal integer.

Syntax

```
BIGINTTOHEX ( integer-expression )
```

Parameters

Parameter	Description
integer-expression	The integer to be converted to hexadecimal.

Examples

Returns the value 00000000000000009:

```
SELECT BIGINTTOHEX(9) FROM iq_dummy
```

Returns the value FFFFFFFFFFFFFFFF7:

```
SELECT BIGINTTOHEX (-9) FROM iq_dummy
```

Usage

BIGINTTOHEX accepts an integer expression that evaluates to BIGINT and returns the hexadecimal equivalent. Returned values are left appended with zeros up to a maximum of 16 digits. All types of unscaled integer data types are accepted as integer expressions.

Conversion is done automatically, if required. Constants are truncated, only if the fraction values are zero. A column cannot be truncated, if the column is declared with a positive scale value. If conversion fails, Sybase IQ returns an error unless the **CONVERSION_ERROR** option is OFF. In that case, the result is NULL.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *HEXTOBIGINT Function [Data Type Conversion]* on page 198
- *HEXTOINT Function [Data Type Conversion]* on page 199
- *INTTOHEX Function [Data Type Conversion]* on page 209

BIT_LENGTH Function [String]

Returns an unsigned 64-bit value containing the bit length of the column parameter.

Syntax

```
BIT_LENGTH( column-name )
```

*Parameters***Table 36. Parameters**

Parameter	Description
column-name	The name of a column

Returns

INT

Usage

The return value of a NULL argument is NULL.

The **BIT_LENGTH** function supports all Sybase IQ data types.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by SQL Anywhere or Adaptive Server Enterprise.

See also

- *BYTE_LENGTH Function [String]* on page 132
- *CHAR_LENGTH Function [String]* on page 137

SQL Functions

- *COL_LENGTH* Function [System] on page 141
- *DATALENGTH* Function [System] on page 155
- *LEN* Function [String] on page 222
- *LENGTH* Function [String] on page 223
- *OBJECT_NAME* Function [System] on page 249
- *OCTET_LENGTH* Function [String] on page 250
- *STR_REPLACE* Function [String] on page 308

BYTE_LENGTH Function [String]

Returns the number of bytes in a string.

Syntax

```
BYTE_LENGTH ( string-expression )
```

Parameters

Parameters	Description
string-expression	The string whose length is to be calculated.

Returns

INT

Example

Returns the value 12:

```
SELECT BYTE_LENGTH( 'Test Message' ) FROM iq_dummy
```

Usage

Trailing white space characters are included in the length returned.

The return value of a NULL string is NULL.

If the string is in a multibyte character set, the **BYTE_LENGTH** value differs from the number of characters returned by **CHAR_LENGTH**.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *BIT_LENGTH* Function [String] on page 131
- *CHAR_LENGTH* Function [String] on page 137
- *COL_LENGTH* Function [System] on page 141
- *DATALENGTH* Function [System] on page 155
- *LEN* Function [String] on page 222
- *LENGTH* Function [String] on page 223
- *OBJECT_NAME* Function [System] on page 249
- *OCTET_LENGTH* Function [String] on page 250
- *STR_REPLACE* Function [String] on page 308

BYTE_LENGTH64 Function

BYTE_LENGTH64 returns an unsigned 64-bit value containing the byte length of the LONG BINARY column parameter.

Usage

BYTE_LENGTH64 also supports the LONG VARCHAR data type and LONG BINARY and LONG VARCHAR variables of any data size.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

BYTE_SUBSTR64 and BYTE_SUBSTR Functions

BYTE_SUBSTR64 and **BYTE_SUBSTR** return the long binary byte substring of the LONG BINARY column parameter.

Usage

The **BYTE_SUBSTR64** and **BYTE_SUBSTR** functions also support the LONG VARCHAR data type and LONG BINARY and LONG VARCHAR variables of any data size.

CHAR_LENGTH64 also supports LONG VARCHAR variables of any data size.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

CAST Function [Data Type Conversion]

Returns the value of an expression converted to a supplied data type.

Syntax

```
CAST ( expression AS data type )
```

Parameters

Parameters	Description
expression	The expression to be converted.
data type	The target data type.

Returns

The specified data type.

Examples

The following function ensures a string is used as a date:

```
CAST( '2000-10-31' AS DATE )
```

The value of the expression **1 + 2** is calculated, and the result cast into a single-character string, the length the data server assigns:

```
CAST( 1 + 2 AS CHAR )
```

You can use the **CAST** function to shorten strings:

```
SELECT CAST( lname AS CHAR(5) ) FROM Customers
```

Usage

If you do not indicate a length for character string types, Sybase IQ chooses an appropriate length. If neither precision nor scale is specified for a DECIMAL conversion, the database server selects appropriate values.

If neither precision nor scale is specified for the explicit conversion of NULL to NUMERIC, the default is NUMERIC(1,0). For example,

```
SELECT CAST( NULL AS NUMERIC ) A,
       CAST( NULL AS NUMERIC(15,2) ) B
```

is described as:

```
A NUMERIC(1,0)
B NUMERIC(15,2)
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *CONVERT Function [Data Type Conversion]* on page 144
- *HOURS Function [Date and Time]* on page 201
- *MINUTES Function [Date and Time]* on page 234

- *MONTHS Function [Date and Time]* on page 237
- *REPLACE Function [String]* on page 279
- *SECOND Function [Date and Time]* on page 290
- *WEEKS Function [Date and Time]* on page 340
- *YEAR Function [Date and Time]* on page 345
- *YEARS Function [Date and Time]* on page 346
- *DAYS Function [Date and Time]* on page 175

CEIL Function [Numeric]

Returns the smallest integer greater than or equal to the specified expression.

CEIL is as synonym for **CEILING**.

Syntax

```
CEIL ( numeric-expression )
```

Parameters

Parameters	Description
expression	A column, variable, or expression with a data type that is either exact numeric, approximate numeric, money, or any type that can be implicitly converted to one of these types. For other data types, CEIL generates an error. The return value has the same data type as the value supplied.

Usage

For a given expression, the **CEIL** function takes one argument. For example, **CEIL (-123.45)** returns -123. **CEIL (123.45)** returns 124.

See *System Administration Guide: Volume 1 > International Language and Character Sets*.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *FLOOR Function [Numeric]* on page 192
- *CEILING Function [Numeric]* on page 136

CEILING Function [Numeric]

Returns the ceiling (smallest integer not less than) of a number.

CEIL is as synonym for **CEILING**.

Syntax

```
CEILING ( numeric-expression )
```

Parameters

Parameter	Description
numeric-expression	The number whose ceiling is to be calculated.

Returns

DOUBLE

Examples

The following statement returns the value 60.00000:

```
SELECT CEILING( 59.84567 ) FROM iq_dummy
```

The following statement returns the value 123:

```
SELECT CEILING( 123 ) FROM iq_dummy
```

The following statement returns the value 124.00:

```
SELECT CEILING( 123.45 ) FROM iq_dummy
```

The following statement returns the value -123.00:

```
SELECT CEILING( -123.45 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *FLOOR Function [Numeric]* on page 192
- *CEIL Function [Numeric]* on page 135

CHAR function [String]

Returns the character with the ASCII value of a number.

Syntax

```
CHAR ( integer-expression )
```

Parameters

Parameter	Description
integer-expression	The number to be converted to an ASCII character. The number must be in the range 0 to 255, inclusive.

Returns

VARCHAR

Examples

The following statement returns the value “Y”:

```
SELECT CHAR( 89 ) FROM iq_dummy
```

The following statement returns the value “S”:

```
SELECT CHAR( 83 ) FROM iq_dummy
```

Usage

The character in the current database character set corresponding to the supplied numeric expression modulo 256 is returned.

CHAR returns NULL for integer expressions with values greater than 255 or less than zero.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

CHAR_LENGTH Function [String]

Returns the number of characters in a string.

Syntax

```
CHAR_LENGTH ( string-expression )
```

Parameters

Parameter	Description
string-expression	The string whose length is to be calculated.

Returns

INT

Usage

Trailing white space characters are included in the length returned.

The return value of a NULL string is NULL.

If the string is in a multibyte character set, the **CHAR_LENGTH** value may be less than the **BYTE_LENGTH** value.

CHAR_LENGTH64 also supports LONG VARCHAR variables of any data size.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data. See *Unstructured Data Analytics in Sybase IQ > Function Support*.

Example

The following statement returns the value 8:

```
SELECT CHAR_LENGTH( 'Chemical' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *BIT_LENGTH Function [String]* on page 131
- *BYTE_LENGTH Function [String]* on page 132
- *COL_LENGTH Function [System]* on page 141
- *DATALength Function [System]* on page 155
- *LEN Function [String]* on page 222
- *LENGTH Function [String]* on page 223
- *OBJECT_NAME Function [System]* on page 249
- *OCTET_LENGTH Function [String]* on page 250
- *STR_REPLACE Function [String]* on page 308

CHAR_LENGTH64 Function

The **CHAR_LENGTH64** function returns an unsigned 64-bit value containing the character length of the LONG VARCHAR column parameter, including the trailing blanks.

Usage

CHAR_LENGTH64 also supports LONG VARCHAR variables of any data size.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data. See *Unstructured Data Analytics in Sybase IQ > Function Support*.

CHARINDEX Function [String]

Returns the position of the first occurrence of a specified string in another string.

Syntax

```
CHARINDEX ( string-expression1, string-expression2 )
```

Parameters

Parameter	Description
string-expression1	The string for which you are searching. This string is limited to 255 bytes.
string-expression2	The string to be searched. The position of the first character in the string being searched is 1.

Returns

INT

Example

The statement:

```
SELECT Surname, GivenName
FROM Employees
WHERE CHARINDEX('K', Surname) = 1
```

returns the following values:

Surname	GivenName
Klobucher	James
Kuo	Felicia
Kelly	Moira

Usage

All the positions or offsets, returned or specified, in the **CHARINDEX** function are always character offsets and may be different from the byte offset for multibyte data.

If the string being searched contains more than one instance of the specified string, **CHARINDEX** returns the position of the first instance.

If the string being searched does not contain the specified string, **CHARINDEX** returns zero (0).

Searching for a zero-length string returns 1.

If any of the arguments is NULL, the result is NULL.

CHARINDEX returns a 32 bit signed integer position for CHAR and VARCHAR columns.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data. See *Unstructured Data Analytics in Sybase IQ > Function Support*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *SUBSTRING Function [String]* on page 313

COALESCE Function [Miscellaneous]

Returns the first non-NULL expression from a list.

Syntax

```
COALESCE ( expression, expression [ , ... ] )
```

Parameters

Table 37. Parameters

Parameter	Description
expression	Any expression.

Returns

ANY

Example

The following statement returns the value 34:

```
SELECT COALESCE( NULL, 34, 13, 0 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *ISNULL Function [Miscellaneous]* on page 212

COL_LENGTH Function [System]

Returns the defined length of a column.

Syntax

```
COL_LENGTH ( table-name , column-name )
```

Parameter	Description
table-name	The table name.
column-name	The column name.

Example

Returns the column length 35:

```
SELECT COL_LENGTH ( 'CUSTOMERS', 'ADDRESS' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *BIT_LENGTH Function [String]* on page 131
- *BYTE_LENGTH Function [String]* on page 132
- *CHAR_LENGTH Function [String]* on page 137
- *DATALENGTH Function [System]* on page 155
- *LEN Function [String]* on page 222
- *LENGTH Function [String]* on page 223
- *OBJECT_NAME Function [System]* on page 249
- *OCTET_LENGTH Function [String]* on page 250
- *STR_REPLACE Function [String]* on page 308

COL_NAME Function [System]

Returns the column name.

Syntax

```
COL_NAME ( table-id , column-id [ , database-id ] )
```

Parameters

Table 38. Parameters

Parameter	Description
table-id	The object ID of the table.
column-id	The column ID of the column.
database-id	The database ID.

Examples

The following statement returns the column name `lname`. The object ID of the `Customers` table is 100209, as returned by the **OBJECT_ID** function. The column ID is stored in the `column_id` column of the `syscolumn` system table. The database ID of the `iqdemo` database is 0, as returned by the **DB_ID** function.

```
SELECT COL_NAME( 100209, 3, 0 ) FROM iq_dummy
```

The following statement returns the column name **city**.

```
SELECT COL_NAME ( 100209, 5 )FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *DB_ID Function [System]* on page 176
- *DB_NAME Function [System]* on page 177
- *DB_PROPERTY Function [System]* on page 178
- *NEXT_DATABASE Function [System]* on page 242
- *OBJECT_ID Function [System]* on page 248
- *OBJECT_NAME Function [System]* on page 249

CONNECTION_PROPERTY Function [System]

Returns the value of a given connection property as a string.

Syntax

```
CONNECTION_PROPERTY ( { integer-expression1 | string-expression }
                    ... [ , integer-expression2 ] )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Parameter	Description
integer-expression1	In most cases, it is more convenient to supply a string expression as the first argument. If you do supply integer-expression1, it is the connection property ID. You can determine this using the PROPERTY_NUMBER function.
string-expression	The connection property name. You must specify either the property ID or the property name.
integer-expression2	The connection ID of the current database connection. The current connection is used if this argument is omitted.

Returns

VARCHAR

Example

The following statement returns the number of prepared statements being maintained, for example, 4:

```
SELECT connection_property( 'PrepStmt' )FROM iq_dummy
```

Usage

The current connection is used if the second argument is omitted.

Standards and Compatibility

- Vendor extension to ISO/ANSI SQL grammar.
- Compatible with Adaptive Server Enterprise.

See also

- *PROPERTY Function [System]* on page 259
- *PROPERTY_NAME Function [System]* on page 261
- *PROPERTY_NUMBER Function [System]* on page 261
- *Properties Available for the Server* on page 120
- *Properties Available for each Database* on page 120
- *Connection Properties* on page 119
- *sp_iqshowpsexec Procedure* on page 484
- *sp_iqcontext Procedure* on page 391

CONVERT Function [Data Type Conversion]

Returns an expression converted to a supplied data type.

Syntax

```
CONVERT ( data-type, expression [ , format-style ] )
```

Parameters

Parameter	Description
data-type	The data type to which the expression is converted.
expression	The expression to be converted.
format-style	For converting strings to date or time data types and vice versa, format-style is a style code number that describes the date format string to be used.

If no *format-style* argument is provided, the database option settings are used.

Table 39. CONVERT format style code output

Without century (yy)	With century (yyyy)	Output
-	0 or 100	mmm dd yyyy hh:nnAM (or PM)
1	101	mm/dd/yy[yy]
2	102	[yy]yy.mm.dd
3	103	dd/mm/yy[yy]
4	104	dd.mm.yy[yy]
5	105	dd-mm-yy[yy]
6	106	dd mmm yy[yy]
7	107	mmm dd, yy[yy]
8	108	hh:nn:ss
-	9 or 109	mmm dd yyyy hh:nn:ss:sssAM (or PM)
10	110	mm-dd-yy[yy]

Without century (yy)	With century (yyyy)	Output
11	111	[yy]yy/mm/dd
12	112	[yy]yymmdd
-	13 or 113	dd mmm yyyy hh:nn:ss:sss (24 hour clock, Europe default + milliseconds, 4-digit year)
14	114	hh:nn:ss (24 hour clock)
-	20 or 120	yyyy-mm-dd hh:nn:ss (24-hour clock, ODBC canonical, 4-digit year)
-	21 or 121	yyyy-mm-dd hh:nn:ss:sss (24 hour clock, ODBC canonical with milliseconds, 4-digit year)
36	136	hh:nn:ss.ssssssAM (or PM)
37	137	hh:nn:ss.ssssss
38	138	mmm dd yy[yy] hh:nn:ss.ssssssAM (or PM)
39	139	mmm dd yy[yy] hh:nn:ss.ssssss
40	140	[yy]yy-mm-dd hh:nn:ss.ssssss
-	365	yyyyjjj (as a string or integer, where jjj is the Julian day number from 1 to 366 within the year)

Abbreviations and values for date parts in the **CONVERT** format style table:

Abbreviation	Date part	Values
hh	hour	0 – 23
nn	minute	0 – 59
ss	second	0 – 59
sss	millisecond	0 – 999
sssss	microsecond	0 – 999999
mmm	month	Jan to Dec
dd	day	1 – 31
yyyy	year	0001 – 9999
mm	month	1 – 12

SQL Functions

Returns

The data type specified.

Examples

The following statements illustrate the use of format styles:

```
SELECT CONVERT( CHAR( 20 ), order_date, 104 )  
FROM sales_order
```

order_date
16.03.1993
20.03.1993
23.03.1993
25.03.1993
...

```
SELECT CONVERT( CHAR( 20 ), order_date, 7 )  
FROM sales_order
```

order_date
mar 16, 93
mar 20, 93
mar 23, 93
mar 25, 93
...

```
SELECT order_datetime, CONVERT(CHAR(30), order_datetime, 40)  
order_datetime40,  
CONVERT(CHAR(30), order_datetime, 140) order_datetime140  
FROM sales_order;
```

order_datetime	order_datetime40	order_datetime140
03/05/2009 01:03:05.123456	09-03-05 01:03:05.123456	2009-03-05 01:03:05.123456
03/05/2009 13:05:07.654321	09-03-05 13:05:07.654321	2009-03-05 13:05:07.654321

```
SELECT CONVERT(CHAR(50), DATETIME('2009-11-03  
11:10:42.033189'), 136) FROM iq_dummy returns 11:10:42.033189AM
```

```
SELECT CONVERT(CHAR(50), NOW(), 137) FROM iq_dummy returns
14:54:48.794122
```

The following statements illustrate the use of the format style 365, which converts data of type DATE and DATETIME to and from either string or integer type data:

```
CREATE TABLE tab
  (date_col DATE, int_col INT, char7_col CHAR(7));
INSERT INTO tab (date_col, int_col, char7_col)
  VALUES ('Dec 17, 2004', 2004352, '2004352');
```

```
SELECT CONVERT(VARCHAR(8), tab.date_col, 365) FROM tab; returns
'2004352'
```

```
SELECT CONVERT(INT, tab.date_col, 365) from tab; returns 2004352
```

```
SELECT CONVERT(DATE, tab.int_col, 365) FROM TAB; returns
2004-12-17
```

```
SELECT CONVERT(DATE, tab.char7_col, 365) FROM tab; returns
2004-12-17
```

The following statement illustrates conversion to an integer, and returns the value 5.

```
SELECT CONVERT( integer, 5.2 ) FROM iq_dummy
```

Usage

The result data type of a **CONVERT** function is a LONG VARCHAR. If you use **CONVERT** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **CONVERT** to the correct data type and size.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise and SQL Anywhere, except for format style 365, which is a Sybase IQ-only extension.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *HOURS Function [Date and Time]* on page 201
- *MINUTES Function [Date and Time]* on page 234
- *MONTHS Function [Date and Time]* on page 237
- *REPLACE Function [String]* on page 279
- *SECOND Function [Date and Time]* on page 290
- *WEEKS Function [Date and Time]* on page 340
- *YEAR Function [Date and Time]* on page 345
- *YEARS Function [Date and Time]* on page 346

CORR Function [Aggregate]

Returns the correlation coefficient of a set of number pairs.

Syntax 1

```
CORR (dependent-expression, independent-expression)
```

Syntax 2

```
CORR (dependent-expression, independent-expression)
```

```
OVER (window-spec)
```

Parameters

Table 40. Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent-expression.
independent-expression	The variable that influences the outcome.

Returns

DOUBLE

Example

The following example performs a correlation to discover whether age is associated with income level. This function returns the value 0.440227:

```
SELECT CORR( Salary, ( YEAR( NOW( ) ) - YEAR( BirthDate ) ) ) FROM Employees;
```

Usage

The **CORR** function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then CORR returns NULL.

dependent-expression and independent-expression are both numeric. The function is applied to the set of (dependent-expression, independent-expression) after eliminating the pairs for which either dependent-expression or independent-expression is NULL. The following computation is made:

$$\text{COVAR_POP}(y, x) / (\text{STDDEV_POP}(x) * \text{STDDEV_POP}(y))$$

where x represents the dependent-expression and y represents the independent-expression.

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL foundation feature outside of core SQL.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

COS Function [Numeric]

Returns the cosine of a number, expressed in radians.

Syntax

```
COS ( numeric-expression )
```

Parameters

Table 41. Parameters

Parameter	Description
numeric-expression	The angle, in radians.

Returns

This function converts its argument to **DOUBLE**, performs the computation in double-precision floating point, and returns a **DOUBLE** as the result. If the parameter is **NULL**, the result is **NULL**.

Example

The following statement returns the value 0.86781:

```
SELECT COS( 0.52 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *ATAN2 Function [Numeric]* on page 128
- *ATAN Function [Numeric]* on page 127
- *ASIN Function [Numeric]* on page 126

SQL Functions

- *ACOS Function [Numeric]* on page 124
- *COT Function [Numeric]* on page 150
- *SIN Function [Numeric]* on page 293
- *TAN Function [Numeric]* on page 318

COT Function [Numeric]

Returns the cotangent of a number, expressed in radians.

Syntax

```
COT ( numeric-expression )
```

Parameters

Table 42. Parameters

Parameter	Description
numeric-expression	The angle, in radians.

Returns

This function converts its argument to DOUBLE, performs the computation in double-precision floating point, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

Example

The following statement returns the value 1.74653:

```
SELECT COT( 0.52 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *COS Function [Numeric]* on page 149
- *ATAN2 Function [Numeric]* on page 128
- *ATAN Function [Numeric]* on page 127
- *ASIN Function [Numeric]* on page 126
- *ACOS Function [Numeric]* on page 124
- *SIN Function [Numeric]* on page 293
- *TAN Function [Numeric]* on page 318

COVAR_POP Function [Aggregate]

Returns the population covariance of a set of number pairs.

Syntax 1

```
COVAR_POP (dependent-expression, independent-expression)
```

Syntax 2

```
COVAR_POP (dependent-expression, independent-expression)
```

```
OVER (window-spec)
```

Parameters

Table 43. Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Example

The following example measures the strength of association between employee age and salary. This function returns the value 73785.840059:

```
SELECT COVAR_POP( Salary, ( YEAR( NOW( ) ) - YEAR( BirthDate ) ) )
FROM Employees;
```

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then **COVAR_POP** returns NULL.

dependent-expression and independent-expression are both numeric. The function is applied to the set of (dependent-expression, independent-expression) after eliminating the pairs for which either dependent-expression or independent-expression is NULL. The following computation is made:

$$(\text{SUM}(x*y) - \text{SUM}(x) * \text{SUM}(y) / n) / n$$

where x represents the dependent-expression and y represents the independent-expression.

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions.*

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL foundation feature outside of core SQL.
- Sybase—Compatible with SQL Anywhere.

COVAR_SAMP Function [Aggregate]

Returns the sample covariance of a set of number pairs.

Syntax 1

COVAR_SAMP (*dependent-expression*, *independent-expression*)

Syntax 2

COVAR_SAMP (*dependent-expression*, *independent-expression*)

OVER (*window-spec*)

Parameters

Table 44. Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Example

The following example measures the strength of association between employee age and salary. This function returns the value 74782.946005:

```
SELECT COVAR_SAMP( Salary, ( 2008 - YEAR( BirthDate ) ) ) FROM Employees ;
```

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then **COVAR_SAMP** returns NULL.

Both dependent-expression and independent-expression are numeric. The function is applied to the set of (dependent-expression, independent-expression) after eliminating the pairs for which either dependent-expression or independent-expression is NULL.

$$(\text{SUM}(x*y) - \text{SUM}(x) * \text{SUM}(y) / n) / (n-1)$$

where x represents the dependent-expression and y represents the independent-expression.

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions.*

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL foundation feature outside of core SQL.
- Sybase—Compatible with SQL Anywhere.

COUNT Function [Aggregate]

Counts the number of rows in a group, depending on the specified parameters.

Syntax

```
COUNT ( * | expression | DISTINCT column-name )
```

Parameters

Parameter	Description
*	Returns the number of rows in each group.
expression	Returns the number of rows in each group where expression is not the NULL value.
DISTINCT column-name	Returns the number of different values in column-name. Rows where the value is the NULL value are not included in the count.

Note: When the query results are displayed, the * is not displayed in the column header, and appears as:

```
Count ( )
```

Returns

INT

Example

Returns each unique city, and the number of rows with that city value:

```
SELECT city , Count(*)
FROM Employees
GROUP BY city
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *AVG Function [Aggregate]* on page 129
- *SUM Function [Aggregate]* on page 315
- *Windowing Aggregate Function Usage* on page 103

CUME_DIST Function [Analytical]

The **CUME_DIST** function is a rank analytical function that calculates the relative position of one value among a group of rows. It returns a decimal value between 0 and 1.

Syntax

```
CUME_DIST () OVER (window-spec)
```

Returns

A DOUBLE value between 0 and 1

Example

The following example returns a result set that provides a cumulative distribution of the salaries of employees who live in California:

```
SELECT DepartmentID, Surname, Salary,CUME_DIST() OVER (PARTITION BY
DepartmentIDORDER BY Salary DESC) "Rank"FROM Employees WHERE State IN
('CA');
```

The returned result set is:

Table 45. CUME_DIST result set

DepartmentID	Surname	Salary	Rank
200	Savarino	72,300.000	0.333333

DepartmentID	Surname	Salary	Rank
200	Clark	45,000.000	0.666667
200	Overbey	39,300.000	1.000000

Usage

Sybase IQ calculates the cumulative distribution of a value of x in a set S of size N using: $\text{CUME_DIST}(x) = \text{number of values in } S \text{ coming before and including } x \text{ in the specified order} / N$

Composite sort-keys are not currently allowed in the **CUME_DIST** function. You can use composite sort-keys with any of the other rank functions.

You can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement. The *window-spec* must contain the **ORDER BY** clause, and cannot contain a **ROWS** or **RANGE** clause.

Note: DISTINCT is not supported.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL feature T612.
- Sybase—Compatible with SQL Anywhere.

DATALENGTH Function [System]

Returns the length of the expression in bytes.

Syntax

```
DATALENGTH ( expression )
```

Parameters

Parameter	Description
expression	The expression is usually a column name. If the expression is a string constant, it must be enclosed in quotes.

Returns

UNSIGNED INT

*Usage***Table 46. DATALENGTH return values**

Data type	DATALENGTH
SMALLINT	2
INTEGER	4
DOUBLE	8
CHAR	Length of the data
BINARY	Length of the data

Example

Returns the value 35, the longest string in the `company_name` column:

```
SELECT MAX( DATALENGTH( company_name ) )
FROM Customers
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *BIT_LENGTH Function [String]* on page 131
- *BYTE_LENGTH Function [String]* on page 132
- *CHAR_LENGTH Function [String]* on page 137
- *COL_LENGTH Function [System]* on page 141
- *LEN Function [String]* on page 222
- *LENGTH Function [String]* on page 223
- *OBJECT_NAME Function [System]* on page 249
- *OCTET_LENGTH Function [String]* on page 250
- *STR_REPLACE Function [String]* on page 308

DATE Function [Date and Time]

Converts the expression into a date, and removes any hours, minutes, or seconds.

Syntax

```
DATE ( expression )
```

*Parameters***Table 47. Parameters**

Parameter	Description
expression	The value to be converted to date format. The expression is usually a string.

Returns

DATE

Example

The following statement returns the value 1988-11-26 as a date.

```
SELECT DATE( '1988-11-26 21:20:53' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

DATEADD Function [Date and Time]

Returns the date produced by adding the specified number of the specified date parts to a date.

Syntax

```
DATEADD ( date-part , numeric-expression , date-expression )
```

Parameters

Parameter	Description
date part	The date part to be added to the date.
numeric-expression	The number of date parts to be added to the date. The numeric-expression can be any numeric type; the value is truncated to an integer. The maximum microsecond in numeric-expression is 2147483647, that is, 35:47.483647 (35 mins 47 secs 483647 mcs).
date-expression	The date to be modified.

Returns

TIMESTAMP

Example

The following statement returns the value 1995-11-02 00:00:00.000:

```
SELECT DATEADD( MONTH, 102, '1987/05/02' ) FROM iq_dummy
```

The following statement returns the value 2009-11-10 14:57:52.722016:

```
SELECT DATEADD(MICROSECOND, 15, '2009-11-10  
14:57:52.722001') FROM iq_dummy
```

The following statement returns the value 1985-05-02 00:00:00.123456:

```
SELECT DATEADD(MICROSECOND, 123456, '1985/05/02')  
FROM iq_dummy
```

The following statement returns the value 1985-05-01 23:59:59.876544:

```
SELECT DATEADD(MICROSECOND, -123456, '1985/05/02')  
FROM iq_dummy
```

The following statement returns the value 2009-11-03 11:10:42.033192:

```
SELECT DATEADD(MCS, 2, '2009-11-03 11:10:42.033190')  
FROM iq_dummy
```

Usage

DATEADD is a Transact-SQL compatible data manipulation function.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *DATECEILING Function [Date and Time]* on page 158
- *DATEDIFF Function [Date and Time]* on page 161
- *DATEFLOOR Function [Date and Time]* on page 164
- *DATEPART Function [Date and Time]* on page 169
- *DATENAME Function [Date and Time]* on page 168
- *DATEROUND Function [Date and Time]* on page 170
- *Date Parts* on page 110

DATECEILING Function [Date and Time]

Calculates a new date, time, or datetime value by increasing the provided value up to the nearest larger value of the specified granularity.

Syntax

```
DATECEILING ( date-part, datetime-expression [, multiple -expression]
```


Parameters

Parameter	Description
date-part	The date part to be added to the date.
datetime-expression	The date, time, or date-time expression containing the value you are evaluating.
multiple-expression	(Optional). A nonzero positive integer value expression specifying how many multiples of the units specified by the date-part parameter to use within the calculation. For example, you can use multiple-expression to specify that you want to regularize your data to 200-microsecond intervals or 10-minute intervals. If multiple-expression evaluates to zero, evaluates to a negative number, is an explicit NULL constant, or is not a valid value for the specified date-part, Sybase IQ generates an error. If multiple-expression evaluates to a NULL, then the function result is NULL.

Examples

This statement returns the value August 13, 2009 10:40.00.000AM:

```
SELECT DATECEILING( MI, 'August 13, 2009, 10:32.00.132AM', 10) FROM iq_dummy
```

This statement returns the value August 13, 2009 10:32.35.456800 AM:

```
SELECT DATECEILING( US, 'August 13, 2009, 10:32.35.456789AM', 200 ) FROM iq_dummy
```

This statement returns the value August 13, 2009 10:32.35.600000 AM:

```
SELECT DATECEILING( US, 'August 13, 2009, 10:32.35.456789AM', 200000 ) FROM iq_dummy
```

This statement returns the value August 13, 2009 10:32.35.456789 AM:

```
SELECT DATECEILING( US, 'August 13, 2009, 10:32.35.456789AM') FROM iq_dummy
```

Usage

This function calculates a new date, time, or datetime value by increasing the provided value up to the nearest larger value with the specified granularity. If you include the optional *multiple-expression* parameter, then the function increases the date and time up to the nearest specified multiple of the specified granularity.

SQL Functions

The data type of the calculated date and time matches the data type of the *multiple-expression* parameter.

The following date parts are not compatible with **DATECEILING**:

- DayofYear
- WeekDay
- CalYearofWeek
- CalWeekofYear
- CalDayofWeek

If you specify a *multiple-expression* for the microsecond, millisecond, second, minute, or hour date parts, Sybase IQ assumes that the multiple applies from the start of the next larger unit of granularity:

- Multiples of microsecond start from the current second
- Multiples of millisecond start from the current second
- Multiples of second start from the current minute
- Multiples of minute start from the current hour
- Multiples of hour start from the current day

For example, if you specify a multiple of two minutes, Sybase IQ applies two-minute intervals starting at the current hour.

For the microsecond, millisecond, second, minute, and hour date parts, specify a *multiple-expression* value that divides evenly into the range of the specified date part:

- For hours, the valid *multiple-expression* values are: 1, 2, 3, 4, 6, 8, 12, 24
- For seconds and minutes, the valid *multiple-expression* values are: 1, 2, 3, 4, 5, 6, 10, 12, 15, 20, 30, 60
- For milliseconds, the valid *multiple-expression* values are: 1, 2, 4, 5, 8, 10, 20, 25, 40, 50, 100, 125, 200, 250, 500, 1000
- For microseconds, the valid *multiple-expression* values are:

1	40	400	4000	40000
2	50	500	5000	50000
4	64	625	6250	62500
5	80	800	8000	100000
8	100	1000	10000	125000
10	125	1250	12500	200000
16	160	1600	15625	250000
20	200	2000	20000	500000

25	250	2500	25000	1000000
32	320	3125	31250	

If you specify a *multiple-expression* for the day, week, month, quarter, or year date parts, Sybase IQ assumes the intervals started at the smallest date value (0001-01-01), smallest time value (00:00:00.000000), or smallest date-time value (0001-01-01.00:00:00.000000). For example, if you specify a multiple of 10 days, Sybase IQ calculates 10-day intervals starting at 0001-01-01.

For the day, week, month, quarter, or year date parts, you need not specify a multiple that divides evenly into the next larger unit of time granularity.

If Sybase IQ rounds to a multiple of the week date part, the date value is always Sunday.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

See also

- *DATEADD Function [Date and Time]* on page 157
- *DATEDIFF Function [Date and Time]* on page 161
- *DATEFLOOR Function [Date and Time]* on page 164
- *DATEPART Function [Date and Time]* on page 169
- *DATENAME Function [Date and Time]* on page 168
- *DATEROUND Function [Date and Time]* on page 170
- *Date Parts* on page 110

DATEDIFF Function [Date and Time]

Returns the interval between two dates.

Syntax

```
DATEDIFF ( date-part, date-expression1, date-expression2 )
```

Parameters

Table 48. Parameters

Parameter	Description
date-part	Specifies the date part in which the interval is to be measured.

Parameter	Description
date-expression1	The starting date for the interval. This value is subtracted from date-expression2 to return the number of date parts between the two arguments.
date-expression2	The ending date for the interval. <i>date-expression1</i> is subtracted from this value to return the number of date parts between the two arguments.

Returns

INT

Examples

The following statement returns 1:

```
SELECT DATEDIFF( HOUR, '4:00AM', '5:50AM' )
FROM iq_dummy
```

The following statement returns 102:

```
SELECT DATEDIFF( MONTH, '1987/05/02', '1995/11/15' )
FROM iq_dummy
```

The following statement returns 0:

```
SELECT DATEDIFF( DAY, '00:00', '23:59' ) FROM iq_dummy
```

The following statement returns 4:

```
SELECT DATEDIFF( DAY, '1999/07/19 00:00', '1999/07/23
23:59' ) FROM iq_dummy
```

The following statement returns 0:

```
SELECT DATEDIFF( MONTH, '1999/07/19', '1999/07/23' )
FROM iq_dummy
```

The following statement returns 1:

```
SELECT DATEDIFF( MONTH, '1999/07/19', '1999/08/23' )
FROM iq_dummy
```

The following statement returns 4:

```
SELECT DATEDIFF(MCS, '2009-11-03 11:10:42.033185',
'2009-11-03 11:10:42.033189') FROM iq_dummy
```

The following statement returns 15:

```
SELECT DATEDIFF(MICROSECOND, '2009-11-10
14:57:52.722001', '2009-11-10 14:57:52.722016')
FROM iq_dummy
```

The following statement returns 1,500,000:

```
SELECT DATEDIFF(MCS, '2000/07/07/07 07:07:06.277777',
'2000/07/07/07 07:07:07.777777') FROM iq_dummy
```

Usage

This function calculates the number of date parts between two specified dates. The result is a signed integer value equal to **(date2 - date1)**, in date parts.

DATEDIFF results are truncated, not rounded, when the result is not an even multiple of the date part.

When you use **day** as the date part, **DATEDIFF** returns the number of midnights between the two times specified, including the second date, but not the first. For example, the following statement returns the value 5. Midnight of the first day 2003/08/03 is not included in the result. Midnight of the second day *is* included, even though the time specified is before midnight.

```
SELECT DATEDIFF( DAY, '2003/08/03 14:00', '2003/08/08 14:00' ) FROM
iq_dummy
```

When you use **month** as the date part, **DATEDIFF** returns the number of first-of-the-months between two dates, including the second date but not the first. For example, both of the following statements return the value 9:

```
SELECT DATEDIFF( MONTH, '2003/02/01', '2003/11/15' ) FROM iq_dummy;
SELECT DATEDIFF( MONTH, '2003/02/01', '2003/11/01' ) FROM iq_dummy;
```

The first date 2003/02/01 is a first-of-month, but is not included in the result of either query. The second date 2003/11/01 in the second query is also a first-of-month and *is* included in the result.

When you use **week** as the date part, **DATEDIFF** returns the number of Sundays between the two dates, including the second date but not the first. For example, in the month 2003/08, the dates of the Sundays are 03, 10, 17, 24, and 31. The following query returns the value 4:

```
SELECT DATEDIFF( week, '2003/08/03', '2003/08/31' ) FROM iq_dummy;
```

The first Sunday (2003/08/03) is not included in the result.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *DATEADD Function [Date and Time]* on page 157
- *DATECEILING Function [Date and Time]* on page 158
- *DATEFLOOR Function [Date and Time]* on page 164
- *DATEPART Function [Date and Time]* on page 169
- *DATENAME Function [Date and Time]* on page 168
- *DATEROUND Function [Date and Time]* on page 170
- *Date Parts* on page 110

DATEFLOOR Function [Date and Time]

Calculates a new date, time, or datetime value by reducing the provided value down to the nearest lower value of the specified multiple with the specified granularity.

Syntax

```
DATEFLOOR ( date-part, datetime-expression [, multiple-expression])
```

Parameters

Parameter	Description
date part	The date part to be added to the date.
datetime-expression	The date, time, or date-time expression containing the value you are evaluating.
multiple-expression	(Optional). A nonzero positive integer value expression specifying how many multiples of the units specified by date-part to use within the calculation. For example, you can use multiple-expression to specify that you want to regularize your data to 200-microsecond intervals or 10-minute intervals If multiple-expression evaluates to zero, evaluates to a negative number, is an explicit NULL constant, or is not a valid value for the specified date-part, Sybase IQ generates an error. If multiple-expression evaluates to a NULL, then the function result is NULL.

Examples

- This statement returns the value August 13, 2009 10:35:00.000AM:

```
SELECT DATEFLOOR( MINUTE, 'August 13, 2009 10:35:22.123AM') FROM iq_dummy
```
- This statement returns the value August 13, 2009 10:32:35.456600 AM:

```
SELECT DATEFLOOR( US, 'August 13, 2009, 10:32:35.456789AM', 200 ) FROM iq_dummy
```
- This statement returns the value August 13, 2009 10:32:35.400000 AM:

```
SELECT DATEFLOOR( US, 'August 13, 2009, 10:32:35.456789AM', 200000 ) FROM iq_dummy
```
- This statement returns the value August 13, 2009 10:32:35.456789 AM:

```
SELECT DATEFLOOR( US, 'August 13, 2009, 10:32:35.456789AM') FROM iq_dummy
```

Usage

This function calculates a new date, time, or datetime value by reducing the provided value down to the nearest lower value with the specified granularity. If you include the optional *multiple-expression* parameter, then the function reduces the date and time down to the nearest specified multiple of the specified granularity.

The data type of the calculated date and time matches the data type of the *multiple-expression* parameter.

The following date parts are not compatible with **DATEFLOOR**:

- DayofYear
- WeekDay
- CalYearofWeek
- CalWeekofYear
- CalDayofWeek

If you specify a *multiple-expression* for the microsecond, millisecond, second, minute, or hour date parts, Sybase IQ assumes that the multiple applies from the start of the next larger unit of granularity:

- Multiples of microsecond start from the current second
- Multiples of millisecond start from the current second
- Multiples of second start from the current minute
- Multiples of minute start from the current hour
- Multiples of hour start from the current day

For example, if you specify a multiple of two minutes, Sybase IQ applies two minute intervals starting at the current hour.

For the microsecond, millisecond, second, minute, and hour date parts, specify a *multiple-expression* value that divides evenly into the range of the specified date part:

- For hours, the valid *multiple-expression* values are: 1, 2, 3, 4, 6, 8, 12, 24
- For seconds and minutes, the valid *multiple-expression* values are: 1, 2, 3, 4, 5, 6, 10, 12, 15, 20, 30, 60
- For milliseconds, the valid *multiple-expression* values are: 1, 2, 4, 5, 8, 10, 20, 25, 40, 50, 100, 125, 200, 250, 500, 1000
- For microseconds, the valid *multiple-expression* values are:

1	40	400	4000	40000
2	50	500	5000	50000
4	64	625	6250	62500
5	80	800	8000	100000

8	100	1000	10000	125000
10	125	1250	12500	200000
16	160	1600	15625	250000
20	200	2000	20000	500000
25	250	2500	25000	1000000
32	320	3125	31250	

If you specify a *multiple-expression* for the day, week, month, quarter, or year date parts, Sybase IQ assumes the intervals started at the smallest date value (0001-01-01), smallest time value (00:00:00.000000), or smallest date-time value (0001-01-01.00:00:00.000000). For example, if you specify a multiple of 10 days, then Sybase IQ calculates 10-day intervals starting at 0001-01-01.

For the day, week, month, quarter, or year date parts, you need not specify a multiple that divides evenly into the next larger unit of time granularity.

If Sybase IQ rounds to a multiple of the week date part, the date value is always Sunday.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

See also

- *DATEADD Function [Date and Time]* on page 157
- *DATECEILING Function [Date and Time]* on page 158
- *DATEDIFF Function [Date and Time]* on page 161
- *DATEPART Function [Date and Time]* on page 169
- *DATENAME Function [Date and Time]* on page 168
- *DATEROUND Function [Date and Time]* on page 170
- *Date Parts* on page 110

DATEFORMAT Function [Date and Time]

Returns a string representing a date expression in the specified format.

Syntax

```
DATEFORMAT ( datetime-expression, string-expression )
```


*Parameters***Table 49. Parameters**

Parameter	Description
datetime-expression	The date/time to be converted. Must be a date, time, timestamp, or character string.
string-expression	The format of the converted date.

Returns

VARCHAR

Example

The following statement returns string values like “Jan 01, 1989”:

```
SELECT DATEFORMAT( start_date, 'Mmm dd, yyyy' ) from Employees;
```

The following statement returns the string “Feb 19, 1987”:

```
SELECT DATEFORMAT( CAST ( '1987/02/19' AS DATE ), 'Mmm Dd, yyyy' )
FROM iq_dummy
```

Usage

The *datetime-expression* to convert must be a date, time, or timestamp data type, but can also be a CHAR or VARCHAR character string. If the date is a character string, Sybase IQ implicitly converts the character string to date, time, or timestamp data type, so an explicit cast, as in the example above, is unnecessary.

Any allowable date format can be used for *string-expression*. Date format strings cannot contain any multibyte characters. Only single-byte characters are allowed in a date/time/datetime format string, even when the collation order of the database is a multibyte collation order like 932JPN.

If '?' represents a multibyte character, then the following query fails:

```
SELECT DATEFORMAT ( start_date, 'yy?') FROM Employees;
```

Instead, move the multibyte character outside of the date format string using the concatenation operator:

```
SELECT DATEFORMAT (start_date, 'yy') + '?' FROM Employees;
```

To set the format used for dates retrieved from the database, see *Reference: Statements and Options > Database Options > Alphabetical List of Options > DATE_FORMAT Option*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

DATENAME Function [Date and Time]

Returns the name of the specified part (such as the month “June”) of a date/time value, as a character string.

Syntax

```
DATENAME ( date-part, date-expression )
```

Parameters

Table 50. Parameters

Parameter	Description
date-part	The date part to be named.
date-expression	The date for which the date part name is to be returned. The date must contain the requested date-part.

Returns

VARCHAR

Example

The following statement returns the value May:

```
SELECT DATENAME( MONTH , '1987/05/02' )
FROM iq_dummy
```

The following statement returns the value 722,001:

```
SELECT DATENAME(MICROSECOND, '2009-11-10
14:57:52.722001') FROM iq_dummy
```

The following statement returns the value 777,777:

```
SELECT DATENAME(MICROSECOND, '2000/07/07
07:07:07.777777') FROM iq_dummy
```

The following statement returns the value 33,189:

```
SELECT DATENAME(MCS, '2009-11-03 11:10:42.033189')
FROM iq_dummy
```

Usage

DATENAME returns a character string, even if the result is numeric, such as 23, for the day.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.

- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *DATEADD Function [Date and Time]* on page 157
- *DATECEILING Function [Date and Time]* on page 158
- *DATEDIFF Function [Date and Time]* on page 161
- *DATEFLOOR Function [Date and Time]* on page 164
- *DATEPART Function [Date and Time]* on page 169
- *DATEROUND Function [Date and Time]* on page 170
- *Date Parts* on page 110

DATEPART Function [Date and Time]

Returns an integer value for the specified part of a date/time value.

Syntax

```
DATEPART ( date-part, date-expression )
```

Parameters

Table 51. Parameters

Parameter	Description
date-part	The date part to be returned.
date-expression	The date for which the part is to be returned. The date must contain the date-part field.

Returns

INT

Example

The following statement returns the value 5:

```
SELECT DATEPART( MONTH, '1987/05/02' )
FROM iq_dummy
```

The following statement returns the value 722,001:

```
SELECT DATEPART(MICROSECOND, '2009-11-10
14:57:52.722001') FROM iq_dummy
```

The following statement returns the value 777,777:

```
SELECT DATEPART(MICROSECOND, '2000/07/07
07:07:07.777777') FROM iq_dummy
```

The following statement returns the value 33,189:

SQL Functions

```
SELECT DATEPART(MCS, '2009-11-03 11:10:42.033189')  
FROM iq_dummy
```

Usage

The **DATE**, **TIME**, and **DTTM** indexes do not support some date parts (Calyearofweek, Calweekofyear, Caldayofweek, Dayofyear, Millisecond, Microsecond).

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *DATEADD Function [Date and Time]* on page 157
- *DATECEILING Function [Date and Time]* on page 158
- *DATEDIFF Function [Date and Time]* on page 161
- *DATEFLOOR Function [Date and Time]* on page 164
- *DATENAME Function [Date and Time]* on page 168
- *DATEROUND Function [Date and Time]* on page 170
- *Date Parts* on page 110

DATEROUND Function [Date and Time]

Calculates a new date, time, or datetime value by rounding the provided value up or down to the nearest multiple of the specified value with the specified granularity.

Syntax

```
DATEROUND (date-part, datetime-expression [,multiple-expression] )
```

Parameters

Parameter	Description
date part	The date part to be returned.
datetime-expression	he date, time, or date-time expression containing the value you are evaluating.

Parameter	Description
multiple-expression	<p>(Optional). A nonzero positive integer value expression specifying how many multiples of the units specified by date-part to use within the calculation. For example, you can use multiple-expression to specify that you want to regularize your data to 200-microsecond intervals or 10-minute intervals.</p> <p>If <i>multiple-expression</i> evaluates to zero, evaluates to a negative number, is an explicit NULL constant, or is not a valid value for the specified <i>date-part</i>, then Sybase IQ generates an error. If <i>multiple-expression</i> evaluates to a NULL, then the function result is NULL.</p>

Examples

This statement returns the value August 13, 2009, 10:30.000AM:

```
SELECT DATEROUND( MI, 'August 13, 2009 10:33.123AM', 10) FROM
iq_dummy
```

This statement returns the value August 13, 2009 10:32:35.456600 AM:

```
SELECT DATEROUND( US, 'August 13, 2009, 10:32:35.456500AM', 200 )
FROM iq_dummy
```

This statement returns the value August 13, 2009 10:32:35.456789 AM:

```
SELECT DATEROUND( US, 'August 13, 2009, 10:32:35.456789AM') FROM
iq_dummy
```

This statement returns the value August 13, 2009 10:32:35.456400 AM:

```
SELECT DATEROUND( US, 'August 13, 2009, 10:32:35.456499AM', 200 )
FROM iq_dummy
```

Usage

This function calculates a new date, time, or datetime value by rounding the provided value up or down to the nearest value with the specified granularity. If you include the optional *multiple-expression* parameter, then the function rounds the date and time to the nearest specified multiple of the specified granularity.

The data type of the calculated date and time matches the data type of the *multiple-expression* parameter.

The following date parts are not compatible with **DATEROUND**:

- DayofYear
- WeekDay

SQL Functions

- CalYearofWeek
- CalWeekofYear
- CalDayofWeek

If you specify a *multiple-expression* for the microsecond, millisecond, second, minute, or hour date parts, Sybase IQ assumes that the multiple applies from the start of the next larger unit of granularity:

- Multiples of microsecond start from the current second
- Multiples of millisecond start from the current second
- Multiples of second start from the current minute
- Multiples of minute start from the current hour
- Multiples of hour start from the current day

For example, if you specify a multiple of two minutes, Sybase IQ applies two minute intervals starting at the current hour.

For the microsecond, millisecond, second, minute, and hour date parts, specify a *multiple-expression* value that divides evenly into the range of the specified date part:

- For hours, the valid *multiple-expression* values are: 1, 2, 3, 4, 6, 8, 12, 24
- For seconds and minutes, the valid *multiple-expression* values are: 1, 2, 3, 4, 5, 6, 10, 12, 15, 20, 30, 60
- For milliseconds, the valid *multiple-expression* values are: 1, 2, 4, 5, 8, 10, 20, 25, 40, 50, 100, 125, 200, 250, 500, 1000
- For microseconds, the valid *multiple-expression* values are:

1	40	400	4000	40000
2	50	500	5000	50000
4	64	625	6250	62500
5	80	800	8000	100000
8	100	1000	10000	125000
10	125	1250	12500	200000
16	160	1600	15625	250000
20	200	2000	20000	500000
25	250	2500	25000	1000000
32	320	3125	31250	

If you specify a *multiple-expression* for the day, week, month, quarter, or year date parts, Sybase IQ assumes the intervals started at the smallest date value (0001-01-01), smallest time value (00:00:00.000000), or smallest date-time value (0001-01-01.00:00:00.000000). For

example, if you specify a multiple of 10 days, then Sybase IQ calculates 10-day intervals starting at 0001-01-01.

For the day, week, month, quarter, or year date parts, you need not specify a multiple that divides evenly into the next larger unit of time granularity.

If Sybase IQ rounds to a multiple of the week date part, then the date value is always Sunday.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

See also

- *DATEADD Function [Date and Time]* on page 157
- *DATECEILING Function [Date and Time]* on page 158
- *DATEDIFF Function [Date and Time]* on page 161
- *DATEFLOOR Function [Date and Time]* on page 164
- *DATEPART Function [Date and Time]* on page 169
- *DATENAME Function [Date and Time]* on page 168
- *Date Parts* on page 110

DATETIME Function [Date and Time]

Converts an expression into a timestamp.

Syntax

```
DATETIME ( expression )
```

Parameters

Table 52. Parameters

Parameter	Description
expression	The expression to be converted. The expression is usually a string. Conversion errors may be reported.

Returns

TIMESTAMP

Example

The following statement returns a timestamp with value 1998-09-09 12:12:12.000:

```
SELECT DATETIME( '1998-09-09 12:12:12.000' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

DAY Function [Date and Time]

Returns an integer from 1 to 31 corresponding to the day of the month of the date specified.

Syntax

```
DAY ( date-expression )
```

*Parameters***Table 53. Parameters**

Parameter	Description
date-expression	The date.

Returns

SMALLINT

Example

The following statement returns the value 12:

```
SELECT DAY( '2001-09-12' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

DAYNAME Function [Date and Time]

Returns the name of the day of the week from the specified date.

Syntax

```
DAYNAME ( date-expression )
```

*Parameters***Table 54. Parameters**

Parameter	Description
date-expression	The date.

Returns
VARCHAR

Example

The following statement returns the value Saturday:

```
SELECT DAYNAME ( '1987/05/02' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

DAYS Function [Date and Time]

Returns the number of days since an arbitrary starting date, returns the number of days between two specified dates, or adds the specified *integer-expression* number of days to a given date.

DAYS ignores hours, minutes, and seconds.

Syntax

```
DAYS ( datetime-expression )  
| ( datetime-expression, datetime-expression )  
| ( datetime-expression, integer-expression )
```

Parameters

Table 55.

Parameter	Description
<i>datetime-expression</i>	A date and time.
<i>integer-expression</i>	The number of days to be added to the <i>datetime-expression</i> . If the <i>integer-expression</i> is negative, the appropriate number of days are subtracted from the date/time. If you supply an <i>integer-expression</i> , the <i>datetime-expression</i> must be explicitly cast as a date.

Returns

INT when you specify two *datetime* expressions.

TIMESTAMP when the second argument you specify is an integer.

Examples

The following statement returns the integer value 729948:

SQL Functions

```
SELECT DAYS( '1998-07-13 06:07:12' ) FROM iq_dummy
```

The following statement returns the integer value -366, which is the difference between the two dates:

```
SELECT DAYS( '1998-07-13 06:07:12',  
'1997-07-12 10:07:12' ) FROM iq_dummy
```

The following statement returns the value 1999-07-14:

```
SELECT DAYS( CAST('1998-07-13' AS DATE ), 366 )  
FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133

DB_ID Function [System]

Returns the database ID number.

Syntax

```
DB_ID ( [ database-name ] )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Table 56. Parameters

Parameter	Description
database-name	A string expression containing the database name. If database-name is a string constant, it must be enclosed in quotes. If no database-name is supplied, the ID number of the current database is returned

Returns

INT

Examples

Returns the value 0, if iqdemo is the only running database:

```
SELECT DB_ID( 'iqdemo' ) FROM iq_dummy
```

Returns the value 0, if executed against the only running database:

```
SELECT DB_ID() FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *COL_NAME Function [System]* on page 141
- *DB_NAME Function [System]* on page 177
- *DB_PROPERTY Function [System]* on page 178
- *NEXT_DATABASE Function [System]* on page 242
- *OBJECT_ID Function [System]* on page 248
- *OBJECT_NAME Function [System]* on page 249

DB_NAME Function [System]

Returns the database name.

Syntax

```
DB_NAME ( [ database-id ] )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Table 57. Parameters

Parameter	Description
database-id	The ID of the database. database-id must be a numeric expression

Returns

VARCHAR

Example

Returns the database name **iqdemo**, when executed against the demo database:

```
SELECT DB_NAME( 0 ) FROM iq_dummy
```

Usage

If no *database-id* is supplied, the name of the current database is returned.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *COL_NAME Function [System]* on page 141
- *DB_ID Function [System]* on page 176
- *DB_PROPERTY Function [System]* on page 178
- *NEXT_DATABASE Function [System]* on page 242
- *OBJECT_ID Function [System]* on page 248
- *OBJECT_NAME Function [System]* on page 249

DB_PROPERTY Function [System]

Returns the value of the given property.

Syntax

```
DB_PROPERTY ( { property-id | property-name }
[ , { database-id | database-name } ] )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Table 58. Parameters

Parameter	Description
property-id	The database property ID.
property-name	The database property name.
database-id	The database ID number, as returned by DB_ID. Typically, the database name is used.
database-name	The name of the database, as returned by DB_NAME.

Returns

VARCHAR

Example

The following statement returns the page size of the current database, in bytes.

```
SELECT DB_PROPERTY( 'PAGESIZE' ) FROM iq_dummy
```

Usage

Returns a string. The current database is used if the second argument is omitted.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *COL_NAME Function [System]* on page 141
- *DB_ID Function [System]* on page 176
- *DB_NAME Function [System]* on page 177
- *NEXT_DATABASE Function [System]* on page 242
- *OBJECT_ID Function [System]* on page 248
- *OBJECT_NAME Function [System]* on page 249

DEGREES Function [Numeric]

Converts a number from radians to degrees.

Syntax

```
DEGREES ( numeric-expression )
```

Parameters

Table 59. Parameters

Parameter	Description
numeric-expression	An angle in radians.

Returns

Returns the degrees of the angle given by *numeric-expression*.

DOUBLE

Example

The following statement returns the value 29.793805:

```
SELECT DEGREES( 0.52 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

DENSE_RANK Function [Analytical]

Ranks items in a group.

Syntax

```
DENSE_RANK ( ) OVER ( ORDER BY expression [ ASC | DESC ] )
```

Parameters

Table 60. Parameters

Parameter	Description
expression	A sort specification that can be any valid expression involving a column reference, aggregates, or expressions invoking these items.

Returns

INTEGER

Example

The following statement illustrates the use of the **DENSE_RANK** function:

```
SELECT s_suppkey, DENSE_RANK(
OVER ( ORDER BY ( SUM(s_acctBal) DESC )
AS rank_dense FROM supplier GROUP BY s_suppkey;

s_suppkey      sum_acctBal      rank_dense
supplier#011   200,000          1
supplier#002   200,000          1
supplier#013   123,000          2
supplier#004   110,000          3
supplier#035   110,000          3
supplier#006   50,000           4
supplier#021   10,000           5
```

Usage

DENSE_RANK is a rank analytical function. The dense rank of row R is defined as the number of rows preceding and including R that are distinct within the groups specified in the **OVER** clause or distinct over the entire result set. The difference between **DENSE_RANK** and **RANK** is that **DENSE_RANK** leaves no gap in the ranking sequence when there is a tie. **RANK** leaves a gap when there is a tie.

DENSE_RANK requires an **OVER (ORDER BY)** clause. The **ORDER BY** clause specifies the parameter on which ranking is performed and the order in which the rows are sorted in each group. This **ORDER BY** clause is used only within the **OVER** clause and is *not* an **ORDER BY** for the **SELECT**. No aggregation functions in the rank query are allowed to specify **DISTINCT**.

The **OVER** clause indicates that the function operates on a query result set. The result set is the rows that are returned after the **FROM**, **WHERE**, **GROUP BY**, and **HAVING** clauses have all been evaluated. The **OVER** clause defines the data set of the rows to include in the computation of the rank analytical function.

The **ASC** or **DESC** parameter specifies the ordering sequence ascending or descending. Ascending order is the default.

DENSE_RANK is allowed only in the select list of a **SELECT** or **INSERT** statement or in the **ORDER BY** clause of the **SELECT** statement. **DENSE_RANK** can be in a view or a union. The **DENSE_RANK** function cannot be used in a subquery, a **HAVING** clause, or in the select list of an **UPDATE** or **DELETE** statement. Only one rank analytical function is allowed per query.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

See also

- *RANK Function [Analytical]* on page 265

DIFFERENCE Function [String]

Compares two strings, evaluates the similarity between them, and returns a value from 0 to 4.

The best match is 4.

Syntax

```
DIFFERENCE ( string-expression1, string-expression2 )
```

Parameters

Table 61. Parameters

Parameter	Description
<i>string-expression1</i>	The first string to compare.
<i>string-expression2</i>	The second string to compare.

Returns

SMALLINT

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

SQL Functions

- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *SOUNDEX Function [String]* on page 299

DIFFERENCE Function Examples

Use the examples as reference for **DIFFERENCE** function usage.

The following statement returns the value 4:

```
SELECT DIFFERENCE( 'Smith', 'Smith' ) FROM iq_dummy
```

The following statement returns the value 4:

```
SELECT DIFFERENCE( 'Smith', 'Smyth' ) FROM iq_dummy
```

The following statement returns the value 3:

```
SELECT DIFFERENCE( 'Smith', 'Sweeney' ) FROM iq_dummy
```

The following statement returns the value 2:

```
SELECT DIFFERENCE( 'Smith', 'Jones' ) FROM iq_dummy
```

The following statement returns the value 1:

```
SELECT DIFFERENCE( 'Smith', 'Rubin' ) FROM iq_dummy
```

The following statement returns the value 0:

```
SELECT DIFFERENCE( 'Smith', 'Wilkins' ) FROM iq_dummy
```

DOW Function [Date and Time]

Returns a number from 1 to 7 representing the day of the week of the specified date, with Sunday=1, Monday=2, and so on.

Syntax

```
DOW ( date-expression )
```

Parameters

Table 62. Parameters

Parameter	Description
date-expression	The date.

Returns

SMALLINT

Example

The following statement returns the value 5:


```
SELECT DOW( '1998-07-09' ) FROM iq_dummy
```

Usage

See *Reference: Statements and Options > Alphabetical List of Options >*

DATE_FIRST_DAY_OF_WEEK Option if you need Monday (or another day) to be the first day of the week.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

ERRORMSG Function [Miscellaneous]

Provides the error message for the current error, or for a specified SQLSTATE or SQLCODE value.

Syntax

```
ERRORMSG ( [ sqlstate | sqlcode ] )
```

sqlstate: string

sqlcode: integer

Parameters

Table 63. Parameters

Parameter	Definition
sqlstate	The SQLSTATE value for which the error message is to be returned.
sqlcode	The SQLCODE value for which the error message is to be returned.

Returns

A string containing the error message.

VARCHAR

Example

The following statement returns the error message for SQLCODE -813:

```
select errormsg( -813 )
```

Return Value

A string containing the error message. If no argument is supplied, the error message for the current state is supplied. Any substitutions (such as table names and column names) are made.

SQL Functions

If an argument is supplied, the error message for the supplied SQLSTATE or SQLCODE is returned, with no substitutions. Table names and column names are supplied as placeholders ('???').

The **ERRORMSG** function returns SQL Anywhere and Sybase IQ error messages.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

EVENT_CONDITION Function [System]

Specifies when an event handler is triggered.

To define an event and its associated handler, use the **CREATE EVENT** statement.

See *Reference: Statements and Options > SQL Statements > CREATE EVENT Statement*.

Syntax

```
EVENT_CONDITION ( condition-name )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Table 64. Parameters

Parameter	Definition
condition-name	The condition triggering the event. The possible values are preset in the database, and are case-insensitive. Each condition is valid only for certain event types.

Table 65. Valid conditions for events

Condition name	Units	Valid for	Comment
DBFreePercent	N/A	DBDiskSpace	DBDiskSpace shows free space in the system database file (.db file), not the IQ store.
DBFreeSpace	Megabytes	DBDiskSpace	
DBSize	Megabytes	GrowDB	
ErrorNumber	N/A	RAISERROR	

Condition name	Units	Valid for	Comment
IdleTime	Seconds	ServerIdle	
Interval	Seconds	All	Time since handler last executed.
LogFreePercent	N/A	LogDiskSpace	
LogFreeSpace	Megabytes	LogDiskSpace	
LogSize	Megabytes	GrowLog	
RemainingValues	Integer	GlobalAutoincrement	The number of remaining values.
TempFreePercent	N/A	TempDiskSpace	TempDiskSpace shows free space in the system temporary file (pointed to by TEMP or IQTMP15 environment variable), not the IQ temporary store.
TempFreeSpace	Megabytes	TempDiskSpace	
TempSize	Megabytes	GrowTemp	

Returns

INT

Example

The following event definition uses the **EVENT_CONDITION** function:

```
create event LogNotifier
type LogDiskSpace
where event_condition( 'LogFreePercent' ) < 50
handler
begin
    message 'LogNotifier message'
end
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *EVENT_PARAMETER Function [System]* on page 186

EVENT_CONDITION_NAME Function [System]

Can be used to list the possible parameters for **EVENT_CONDITION**.

To define an event and its associated handler, use the **CREATE EVENT** statement.

See *Reference: Statements and Options > SQL Statements > CREATE EVENT Statement*.

Syntax

```
EVENT_CONDITION_NAME ( integer )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Parameter	Description
integer	Must be greater than or equal to zero.

Returns

VARCHAR

Usage

You can use **EVENT_CONDITION_NAME** to obtain a list of all **EVENT_CONDITION** arguments by looping over integers until the function returns NULL.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

EVENT_PARAMETER Function [System]

Provides context information for event handlers.

To define an event and its associated handler, use the **CREATE EVENT** statement.

See *Reference: Statements and Options > SQL Statements > CREATE EVENT statement*.

Syntax

```
EVENT_PARAMETER ( context-name )
```

```
context-name :
'ConnectionID'
'User'
'EventName'
'Executions'
'IQDBMainSpaceName'
'NumActive'
```

'TableName'
condition-name

Note: CIS functional compensation performance considerations apply.

Parameters

Table 66. Parameters

Parameter	Description
context-name	One of the preset strings. The strings are case-insensitive, and carry the following information:
ConnectionId	The connection ID, as returned by <code>connection_property('id')</code>
User	The user ID for the user that caused the event to be triggered.
EventName	The name of the event that has been triggered.
Executions	The number of times the event handler has been executed.
NumActive	The number of active instances of an event handler. This is useful if you want to limit an event handler so that only one instance executes at any given time.
TableName	The name of the table, for use with Remaining-Values.

In addition, you can access any of the valid *condition-name* arguments to the **EVENT_CONDITION** function from the **EVENT_PARAMETER** function.

Returns

VARCHAR

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *EVENT_CONDITION* Function [System] on page 184

EXP Function [Numeric]

Returns the exponential function, e to the power of a number.

Syntax

```
EXP ( numeric-expression )
```

*Parameters***Table 67. Parameters**

Parameter	Description
numeric-expression	The exponent.

Returns

DOUBLE

Example

The following statement returns the value 3269017.3724721107:

```
SELECT EXP( 15 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

EXP_WEIGHTED_AVG Function [Aggregate]

Calculates an exponential weighted moving average.

Weightings determine the relative importance of each quantity that makes up the average.

Syntax

```
EXP_WEIGHTED_AVG ( expression, period-expression )
```

```
OVER ( window-spec )
```

Parameters

Parameter	Description
expression	A numeric expression for which a weighted value is being computed.
period-expression	A numeric expression specifying the period for which the average is to be computed.

Usage

Similar to the **WEIGHTED_AVG** function, the weights in **EXP_WEIGHTED_AVG** decrease over time. However, weights in **WEIGHTED_AVG** decrease arithmetically, whereas weights in **EXP_WEIGHTED_AVG** decrease exponentially. Exponential weighting applies more weight to the most recent values, and decreases the weight for older values while still applying some weight.

Sybase IQ calculates the exponential moving average using:

$$S * C + (1 - S) * PEMA$$

In the calculation above, Sybase IQ applies the smoothing factor by multiplying the current closing price (C) by the smoothing constant (S) added to the product of the previous day's exponential moving average value (PEMA) and 1 minus the smoothing factor.

Sybase IQ calculates the exponential moving average over the entire period specified by the **OVER** clause. *period-expression* specifies the moving range of the exponential moving average.

You can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement. The *window-spec* must contain an **ORDER BY** statement and cannot contain a frame specification.

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause. DISTINCT is not supported.

Example

The following example returns an exponential weighted average of salaries for employees in Florida with the salary of recently hired employees contributing the most weight to the average. There are three rows used in the weighting:

```
SELECT DepartmentID, Surname, Salary, EXP_WEIGHTED_AVG(Salary, 3)
OVER (ORDER BY YEAR(StartDate) DESC) as "W_AVG" FROM Employees WHERE
State IN ('FL') ORDER BY StartDate DESC
```

The returned result set is:

Table 68. EXP_WEIGHTED_AVG result set

DepartmentID	Surname	Salary	W_AVG
400	Evans	68,940.000	34,470.000000
300	Litton	58,930.000	46,700.000000
200	Sterling	64,900.000	55,800.000000
200	Kelly	87,500.000	71,650.000000
400	Charlton	28,300.000	49,975.000000

DepartmentID	Surname	Salary	W_AVG
100	Lull	87,900.000	68,937.500000
100	Gowda	59,840.000	60,621.875000
400	Francis	53,870.000	61,403.750000

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

See also

- *WEIGHTED_AVG Function [Aggregate]* on page 341
- *Windowing Aggregate Function Usage* on page 103

FIRST_VALUE Function [Aggregate]

Returns the first value from a set of values.

Syntax

```
FIRST_VALUE (expression [IGNORE NULLS | RESPECT NULLS])  
OVER (window-spec)
```

Parameters

Parameter	Description
expression	The expression on which to determine the first value in an ordered set.

Returns

Data type of the argument.

Usage

FIRST_VALUE returns the first value in a set of values, which is usually an ordered set. If the first value in the set is null, then the function returns NULL unless you specify **IGNORE NULLS**. If you specify **IGNORE NULLS**, then **FIRST_VALUE** returns the first non-null value in the set, or NULL if all values are null.

The data type of the returned value is the same as that of the input value.

You cannot use **FIRST_VALUE** or any other analytic function for expression. That is, you cannot nest analytic functions, but you can use other built-in function expressions for expression.

If the window-spec does not contain an **ORDER BY** expression, or if the **ORDER BY** expression is not precise enough to guarantee a unique ordering, then the result is arbitrary. If there is no window-spec, then the result is arbitrary.

You can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Note: DISTINCT is not supported.

Example

The following example returns the relationship, expressed as a percentage, between each employee's salary and that of the most recently hired employee in the same department:

```
SELECT DepartmentID, EmployeeID,
100 * Salary / ( FIRST_VALUE( Salary ) OVER (
PARTITION BY DepartmentID ORDER BY Year(StartDate) DESC ) )
AS percentage
FROM Employees order by DepartmentID DESC;
```

The returned result set is:

Table 69. FIRST_VALUE result set

DepartmentID	EmployeeID	Percentage
500	1,658	100.000000000000000000000000000000
500	1,570	138.842709713689113761394
500	1,615	110.428462434244870095972
500	1,013	109.585190539292454724330
500	750	137.734409508894510701521
500	921	167.449704854836766654619
500	868	113.239368750752921334778
500	703	222.867927558928643135365
500	191	119.664297474199895594908
400	1,684	100.000000000000000000000000000000
400	1,740	76.128652163477274215016
400	1,751	76.353400685155687446813
400	1,607	133.758100765890593292456
400	1,507	77.996465120338650199655

DepartmentID	EmployeeID	Percentage
400	1,576	150.428767810774836893669

In this example, employee 1658 is the first row for department 500, indicating that employee 1658 is the most recent hire in that department, and therefore receives a percentage of 100%. Percentages for the remaining employees in department 500 are calculated relative to that of employee 1658. For example, employee 1570 earns approximately 139% of what employee 1658 earns.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

FLOOR Function [Numeric]

Returns the floor of (largest integer not greater than) a number.

Syntax

```
FLOOR ( numeric-expression )
```

Parameters

Table 70. Parameters

Parameter	Description
numeric-expression	The number, usually a float.

Returns

DOUBLE

Examples

The following statement returns the value 123.00:

```
SELECT FLOOR ( 123 ) FROM iq_dummy
```

The following statement returns the value 123:

```
SELECT FLOOR ( 123.45 ) FROM iq_dummy
```

The following statement returns the value -124.00.

```
SELECT FLOOR ( -123.45 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *CEILING Function [Numeric]* on page 136
- *CEIL Function [Numeric]* on page 135

GETDATE Function [Date and Time]

Returns the current date and time.

Syntax

```
GETDATE ( )
```

Returns

TIMESTAMP

Example

The following statement returns the system date and time.

```
SELECT GETDATE( ) FROM iq_dummy
```

Usage

GETDATE is a Transact-SQL compatible data manipulation function.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

GRAPHICAL_PLAN Function [String]

Returns the graphical query plan to Interactive SQL in an XML format string.

Syntax

```
GRAPHICAL_PLAN ( string-expression
[ , statistics-level
[ , cursor-type
[ , update-status ] ] )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Parameter	Description
string-expression	SQL statement for which the plan is to be generated. string-expression is generally a SELECT statement, but it can also be an UPDATE or DELETE , INSERT SELECT , or SELECT INTO statement.
statistics-level	An integer. Statistics-level can be: <ul style="list-style-type: none"> • 0 – Optimizer estimates only (default). • 2 – Detailed statistics including node statistics. • 3 – Detailed statistics.
cursor-type	A cursor type, expressed as a string. Possible values are: asensitive, insensitive, sensitive, or key-set-driven. If cursor-type is not specified, asensitive is used by default.
update-status	A string parameter accepting one of the following values indicating how the optimizer should treat the given cursor: READ-ONLY – The cursor is read-only. READ-WRITE (default) – The cursor can be read or written to. READ-WRITE (default) – The cursor can be read or written to.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **GRAPHICAL_PLAN** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **GRAPHICAL_PLAN** to the correct data type and size.

Usage

If you do not provide an argument to the **GRAPHICAL_PLAN** function, the query plan is returned to you from the cache. If there is no query plan in the cache, then this message appears:

```
plan not available
```

The behavior of **GRAPHICAL_PLAN** function is controlled by database options `QUERY_PLAN_TEXT_ACCESS` and `QUERY_PLAN_TEXT_CACHING`. If `QUERY_PLAN_TEXT_ACCESS` is OFF (the default), then this message appears:

```
Plan not available. The database option QUERY_PLAN_TEXT_ACCESS is OFF
```

If a user needs access to the plan, the DBA must set option `QUERY_PLAN_TEXT_ACCESS` ON for that user.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_TEXT_ACCESS Option*. See *Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_TEXT_CACHING Option*.

If `QUERY_PLAN_TEXT_ACCESS` is ON, and the query plan for the string expression is available in the cache maintained on the server, the query plan from the cache is returned to you.

If the query plan is not available in the cache and you are authorized to view plans on the client, then a query plan with optimizer estimates (query plan with `NOEXEC` option ON) is generated and appears on the Interactive SQL client plan window.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > NOEXEC Option*.

When a user requests a query plan that has not yet been executed, the query plan is not available in the cache. Instead, a query plan with optimizer estimates is returned without `QUERY_PLAN_AFTER_RUN` statistics.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_AFTER_RUN Option*.

Query plans for stored procedures are not accessible using the **GRAPHICAL_PLAN** function.

Users can view the query plan for cursors opened for Sybase IQ queries. A cursor is declared and opened using **DECLARE CURSOR** and **OPEN CURSOR**. To obtain the query plan for the most recently opened cursor, use:

```
SELECT GRAPHICAL_PLAN ( );
```

With the `QUERY_PLAN_AFTER_RUN` option OFF, the plan appears after **OPEN CURSOR** or **CLOSE CURSOR**. However, if `QUERY_PLAN_AFTER_RUN` is ON, **CLOSE CURSOR** must be executed before you request the plan.

For information on viewing the query optimizer's execution plan for a SQL statement in the Plan Viewer window in Interactive SQL, see *SQL Anywhere Server - Database Administration > Administering your database > SQL Anywhere graphical administration tools > Using Interactive SQL > Using the Plan Viewer to view graphical plans*.

Examples

The following example passes a **SELECT** statement as a string parameter and returns the plan for executing the query. It saves the plan in the file `gplan.xml`.

Note: If you use the **OUTPUT** statement's **HEXADECIMAL** clause set to **ASIS** to get formatted plan output, the values of characters are written without any escaping, even if the value contains control characters. **ASIS** is useful for text that contains formatting characters such as tabs or carriage returns.

```
SELECT GRAPHICAL_PLAN ('SELECT * FROM Employees');OUTPUT to 'C:\gplan.xml' HEXADECIMAL ASIS quote '';
```

The following example returns the query plan from the cache, if available:

```
SELECT GRAPHICAL_PLAN ( );
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *HTML_PLAN Function [String]* on page 204

GROUPING Function [Aggregate]

Identifies whether a column in a **ROLLUP** or **CUBE** operation result set is NULL because it is part of a subtotal row, or NULL because of the underlying data.

Syntax

```
GROUPING ( group-by-expression )
```

Parameters

Parameter	Description
group-by-expression	An expression appearing as a grouping column in the result set of a query that uses a GROUP BY clause with the ROLLUP or CUBE keyword. The function identifies subtotal rows added to the result set by a ROLLUP or CUBE operation.

Currently, Sybase IQ does not support the **PERCENTILE_CONT** or **PERCENTILE_DISC** functions with **GROUP BY CUBE** operations.

Returns

Value	Description
1	Indicates that group-by-expression is NULL because it is part of a subtotal row. The column is not a prefix column for that row.
0	Indicates that group-by-expression is a prefix column of a subtotal row.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103

GROUP_MEMBER Function [System]

Identifies whether the user belongs to the specified group.

Syntax

```
GROUP_MEMBER ( group-name-string-expression [ , user-name-string-expression ] )
```

Parameters

Parameter	Description
<i>group-name-string-expression</i>	Identifies the group to be considered.
<i>user-name-string-expression</i>	Identifies the user to be considered. If not supplied, then the current user name is assumed.

*Return Values***Table 71. Return Values**

Value	Description
0	Returns 0 if the group does not exist, if the user does not exist, or if the user does not belong to the specified group.

Value	Description
1	Returns an integer other than 0 if the user is a member of the specified group.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

HEXTOBIGINT Function [Data Type Conversion]

Returns the BIGINT equivalent of a hexadecimal string.

Syntax

```
HEXTOBIGINT ( hexadecimal-string )
```

Parameters

Parameter	Description
hexadecimal-string	The hexadecimal value to be converted to a big integer (BIGINT). Input can be in the following forms, with either a lowercase or uppercase “0x” in the prefix, or no prefix: <i>0xhex-string</i> <i>0Xhex-string</i> <i>hex-string</i>

Examples

The following statements return the value 4294967287:

```
SELECT HEXTOBIGINT ( '0xffffffff7' ) FROM iq_dummy
SELECT HEXTOBIGINT ( '0Xffffffff7' ) FROM iq_dummy
SELECT HEXTOBIGINT ( 'ffffffff7' ) FROM iq_dummy
```

Usage

The **HEXTOBIGINT** function accepts hexadecimal integers and returns the BIGINT equivalent. Hexadecimal integers can be provided as CHAR and VARCHAR value expressions, as well as BINARY and VARBINARY expressions.

The **HEXTOBIGINT** function accepts a valid hexadecimal string, with or without a “0x” or “0X” prefix, enclosed in single quotes.

Input of fewer than 16 digits is assumed to be left-padded with zeros.

For data type conversion failure on input, Sybase IQ returns an error unless the `CONVERSION_ERROR` option is set to `OFF`. When `CONVERSION_ERROR` is `OFF`, invalid hexadecimal input returns `NULL`.

See Reference: Statements and Options > Database Options > Alphabetical List of Options > `CONVERSION_ERROR` Option [TSQL].

An error is returned if a `BINARY` or `VARBINARY` value exceeds 8 bytes and a `CHAR` or `VARCHAR` value exceeds 16 characters, with the exception of the value being appended with `'0x.'`

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *BIGINTTOHEX Function [Data Type Conversion]* on page 130
- *HEXTOINT Function [Data Type Conversion]* on page 199
- *INTTOHEX Function [Data Type Conversion]* on page 209

HEXTOINT Function [Data Type Conversion]

Returns the unsigned `BIGINT` equivalent of a hexadecimal string.

Syntax

```
HEXTOINT ( hexadecimal-string )
```

Parameters

Parameters	Description
hexadecimal-string	The string to be converted to an integer. Input can be in the following forms, with either a lowercase or uppercase “x” in the prefix, or no prefix: <i>0xhex-string</i> <i>0Xhex-string</i> <i>hex-string</i>

Returns

The `HEXTOINT` function returns the platform-independent SQL `INTEGER` equivalent of the hexadecimal string. The hexadecimal value represents a negative integer if the 8th digit from the right is one of the digits 8-9 and the uppercase or lowercase letters A-F and the previous leading digits are all uppercase or lowercase letter F. The following is not a valid use of

SQL Functions

HEXTOINT since the argument represents a positive integer value that cannot be represented as a signed 32-bit integer:

```
SELECT HEXTOINT( '0x0080000001' );
```

INT

Examples

The following statements return the value 420:

```
SELECT HEXTOINT ( '0x1A4' ) FROM iq_dummy
```

```
SELECT HEXTOINT ( '0X1A4' ) FROM iq_dummy
```

```
SELECT HEXTOINT ( '1A4' ) FROM iq_dummy
```

Usage

For invalid hexadecimal input, Sybase IQ returns an error unless the `CONVERSION_ERROR` option is OFF. When `CONVERSION_ERROR` is OFF, invalid hexadecimal input returns NULL.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > CONVERSION_ERROR Option [TSQL]*.

The database option `ASE_FUNCTION_BEHAVIOR` specifies that output of Sybase IQ functions, including **INTTOHEX** and **HEXTOINT**, is consistent with the output of Adaptive Server Enterprise functions.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > ASE_FUNCTION_BEHAVIOR Option*.

When the `ASE_FUNCTION_BEHAVIOR` option is ON:

- Sybase IQ **HEXTOINT** assumes input is a hexadecimal string of 8 characters; if the length is less than 8 characters long, the string is left padded with zeros.
- Sybase IQ **HEXTOINT** accepts a maximum of 16 characters prefixed with 0x (a total of 18 characters); use caution, as a large input value can result in an integer value that overflows the 32-bit signed integer output size.
- The data type of the output of the Sybase IQ **HEXTOINT** function is assumed to be a 32-bit signed integer.
- Sybase IQ **HEXTOINT** accepts a 32-bit hexadecimal integer as a signed representation.
- For more than 8 hexadecimal characters, Sybase IQ **HEXTOINT** considers only relevant characters.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *BIGINTTOHEX Function [Data Type Conversion]* on page 130
- *HEXTOBIGINT Function [Data Type Conversion]* on page 198
- *INTTOHEX Function [Data Type Conversion]* on page 209

HOURL Function [Date and Time]

Returns a number from 0 to 23 corresponding to the hour component of the specified date/time.

Syntax

```
HOURL ( datetime-expression )
```

Table 72. Parameters

Parameter	Definition
<i>datetime-expression</i>	The date/time.

Returns

SMALLINT

Example

The following statement returns the value 21:

```
SELECT HOURL( '1998-07-09 21:12:13' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

HOURS Function [Date and Time]

Returns the number of hours since an arbitrary starting date and time, the number of whole hours between two specified times, or adds the specified integer-expression number of hours to a time.

Syntax

```
HOURS ( datetime-expression
| datetime-expression, datetime-expression
| datetime-expression, integer-expression )
```

*Parameters***Table 73. Parameters**

Parameter	Description
datetime-expression	A date and time.
integer-expression	The number of hours to be added to the <i>datetime-expression</i> . If <i>integer-expression</i> is negative, the appropriate number of hours are subtracted from the date/time. If you supply an integer expression, the <i>datetime-expression</i> must be explicitly cast as a datetime data type.

Returns

INT

Examples

The following statement returns the value 17518758:

```
SELECT HOURS( '1998-07-13 06:07:12' ) FROM iq_dummy
```

The following statement returns the value 4, to signify the difference between the two times:

```
SELECT HOURS( '1999-07-13 06:07:12',  
             '1999-07-13 10:07:12' ) FROM iq_dummy
```

The following statement returns the datetime value 1999-05-13 02:05:07.000:

```
SELECT HOURS( CAST( '1999-05-12 21:05:07'  
AS DATETIME ), 5 ) FROM iq_dummy
```

Usage

The second syntax returns the number of whole hours from the first date/time to the second date/time. The number might be negative.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *CONVERT Function [Data Type Conversion]* on page 144
- *MINUTES Function [Date and Time]* on page 234
- *MONTHS Function [Date and Time]* on page 237
- *REPLACE Function [String]* on page 279

- *SECOND Function [Date and Time]* on page 290
- *WEEKS Function [Date and Time]* on page 340
- *YEAR Function [Date and Time]* on page 345
- *YEARS Function [Date and Time]* on page 346

HTML_DECODE Function [HTTP]

Decodes special character entities that appear in HTML literal strings.

Syntax

```
HTML_DECODE ( string )
```

Note: CIS functional compensation performance considerations apply.

The **HTML_DECODE** function is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > HTML_DECODE function [HTTP]*.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **HTML_DECODE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **HTML_DECODE** to the correct data type and size.

HTML_ENCODE Function [HTTP]

Encodes special characters within strings to be inserted into HTML documents.

Syntax

```
HTML_ENCODE ( string )
```

Note: CIS functional compensation performance considerations apply.

The **HTML_ENCODE** function is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > HTML_ENCODE function [HTTP]*.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **HTML_ENCODE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **HTML_ENCODE** to the correct data type and size.

HTML_PLAN Function [String]

Returns query plans in an HTML format string.

Syntax

```
HTML_PLAN ( string-expression )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Parameter	Description
string-expression	SQL statement for which the plan is to be generated. It is primarily a SELECT statement but can be an UPDATE or DELETE statement.

If you do not provide an argument to the **HTML_PLAN** function, the query plan is returned to you from the cache. If there is no query plan in the cache, this message appears:

```
No plan available
```

The behavior of the **HTML_PLAN** function is controlled by database options **QUERY_PLAN_TEXT_ACCESS** and **QUERY_PLAN_TEXT_CACHING**. If **QUERY_PLAN_TEXT_ACCESS** is OFF (the default), this message appears:

```
Plan not available. The database option QUERY_PLAN_TEXT_ACCESS is OFF
```

If **QUERY_PLAN_TEXT_ACCESS** is ON, and the query plan for the string expression is available in the cache maintained on the server, the query plan from the cache is returned to you.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_TEXT_ACCESS Option*. See *Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_TEXT_CACHING Option*.

The **HTML_PLAN** function can be used to return query plans to Interactive SQL using **SELECT**, **UPDATE**, **DELETE**, **INSERT SELECT**, and **SELECT INTO**.

Users can view the query plan for cursors opened for Sybase IQ queries. To obtain the query plan for the most recently opened cursor, use:

```
SELECT HTML_PLAN ( );
```

With **QUERY_PLAN_AFTER_RUN** option OFF, the plan appears after **OPEN CURSOR** or **CLOSE CURSOR**. However, if **QUERY_PLAN_AFTER_RUN** is ON, **CLOSE CURSOR** must be executed before you request the plan.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > QUERY_PLAN_AFTER_RUN Option*.

For information on viewing the query optimizer's execution plan for a SQL statement in the Plan Viewer window in Interactive SQL, see *SQL Anywhere Server - Database Administration > Administering your database > SQL Anywhere graphical administration tools > Using Interactive SQL > Using the Plan Viewer to view graphical plans*.

When you request an **HTML_PLAN** for a SQL Anywhere query or for an OMNI/CIS decomposed query, the following message is returned:

```
No plan. HTML_PLAN function is not supported for this type of
statement or database.
```

Examples

The following example passes a **SELECT** statement as a string parameter and returns the HTML plan for executing the query. It saves the plan in the file `hplan.html`.

```
SELECT HTML_PLAN ('SELECT * FROM Employees');
OUTPUT to 'C:\hplan.html' HEXADECIMAL ASIS QUOTE '';
```

The **OUTPUT TO** clause **HEXADECIMAL ASIS** is useful for text that contains formatting characters such as tabs or carriage returns. When set to **ASIS**, values are written as is, without any escaping, even if the values contain control characters.

The following example returns the HTML query plan from the cache, if available.

```
SELECT HTML_PLAN ( );
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *GRAPHICAL_PLAN Function [String]* on page 193

HTTP_DECODE Function [HTTP]

Decodes special characters within strings for use with HTTP.

Syntax

```
HTTP_DECODE ( string )
```

Note: CIS functional compensation performance considerations apply.

The `HTTP_DECODE` function is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > HTTP_DECODE function [HTTP]*.

HTTP_ENCODE Function [HTTP]

Encodes special characters in strings for use with HTTP.

Syntax

```
HTTP_ENCODE ( string )
```

Note: CIS functional compensation performance considerations apply.

The `HTTP_ENCODE` function is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > HTTP_ENCODE function [HTTP]*.

HTTP_HEADER Function [HTTP]

Gets the value of an HTTP header.

Syntax

```
HTTP_HEADER ( field-name )
```

Note: CIS functional compensation performance considerations apply.

The `HTTP_HEADER` function is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > HTTP_HEADER function [HTTP]*.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use `HTTP_HEADER` in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set `HTTP_HEADER` to the correct data type and size.

HTTP_VARIABLE function [HTTP]

Gets the value of an HTTP variable.

Syntax

```
HTTP_VARIABLE ( var-name [ [ , instance ] , header-field )
```

Note: CIS functional compensation performance considerations apply.

The `HTTP_VARIABLE` function is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > HTTP_VARIABLE function [HTTP]*.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **HTTP_VARIABLE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **HTTP_VARIABLE** to the correct data type and size.

IFNULL Function [Miscellaneous]

Returns the first nonnull expression, or NULL.

If the first expression is the NULL value, then the value of the second expression is returned. If the first expression is not NULL, the value of the third expression is returned. If the first expression is not NULL and there is no third expression, then the NULL value is returned.

Syntax

```
IFNULL ( expression1, expression2 [ , expression3 ] )
```

Parameters

Table 74. Parameters

Parameter	Description
expression1	The expression to be evaluated. Its value determines whether <i>expression2</i> or <i>expression3</i> is returned.
expression2	The return value if <i>expression1</i> is NULL
expression3	The return value if <i>expression1</i> is not NULL.

Returns

The data type returned depends on the data type of *expression-2* and *expression-3*.

Examples

The following statement returns the value -66:

```
SELECT IFNULL( NULL, -66 ) FROM iq_dummy
```

The following statement returns NULL, because the first expression is not NULL and there is no third expression:

```
SELECT IFNULL( -66, -66 ) FROM iq_dummy
```

Standards and compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

INDEX_COL Function [System]

Returns the name of the indexed column.

Syntax

```
INDEX_COL ( table-name, index-id, key_# [ , user-id ] )
```

Parameters

Parameter	Definition
table-name	A table name.
index-id	The index ID of an index of <i>table-name</i> .
key_#	A key in the index specified by <i>index-id</i> . This parameter specifies the column number in the index. For a single column index, <i>key_#</i> is equal to 0. For a multicolumn index, <i>key_#</i> is equal to 0 for the first column, 1 for the second column, and so on.
user-id	The user ID of the owner of <i>table-name</i> . If <i>user-id</i> is not specified, this value defaults to the caller's user ID.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *OBJECT_ID Function [System]* on page 248

INSERTSTR Function [String]

Inserts a string into another string at a specified position.

Syntax

```
INSERTSTR ( numeric-expression, string-expression1, string-expression2 )
```

*Parameters***Table 75. Parameters**

Parameter	Definition
numeric-expression	The position after which <i>string-expression2</i> is to be inserted. Use zero to insert a string at the beginning.
string-expression1	The string into which <i>string-expression2</i> is to be inserted.
string-expression2	The string to be inserted.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **INSERTSTR** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **INSERTSTR** to the correct data type and size.

Example

The following statement returns the value “backoffice”:

```
SELECT INSERTSTR( 0, 'office ', 'back' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported in Adaptive Server Enterprise. The STUFF function is equivalent and is supported in both Adaptive Server Enterprise and Sybase IQ.

INTTOHEX Function [Data Type Conversion]

Returns the hexadecimal equivalent of a decimal integer.

Syntax

```
INTTOHEX ( integer-expression )
```

*Parameters***Table 76. Parameters**

Parameter	Description
integer-expression	The integer to be converted to hexadecimal.

Returns

VARCHAR

Examples

The following statement returns the value 3B9ACA00:

```
SELECT INTTOHEX( 1000000000 ) FROM iq_dummy
```

The following statement returns the value 00000002540BE400:

```
SELECT INTTOHEX ( 10000000000 ) FROM iq_dummy
```

Usage

If data conversion of input to **INTTOHEX** conversion fails, Sybase IQ returns an error, unless the `CONVERSION_ERROR` option is OFF. In that case, the result is NULL.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > CONVERSION_ERROR Option [TSQL]*.

The database option `ASE_FUNCTION_BEHAVIOR` specifies that output of Sybase IQ functions, including **INTTOHEX** and **HEXTPOINT**, be consistent with the output of Adaptive Server Enterprise functions. The default value of `ASE_FUNCTION_BEHAVIOR` is OFF.

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > ASE_FUNCTION_BEHAVIOR Option*.

When the `ASE_FUNCTION_BEHAVIOR` option is disabled (the value is OFF):

- The output of **INTTOHEX** is compatible with SQL Anywhere.
- Depending on the input, the output of **INTTOHEX** can be 8 digits or 16 digits and is left padded with zeros; the return data type is VARCHAR.
- The output of **INTTOHEX** does not have a '0x' or '0X' prefix.
- The input to **INTTOHEX** can be up to a 64-bit integer.

When the `ASE_FUNCTION_BEHAVIOR` option is enabled (the value is ON):

- The output of **INTTOHEX** is compatible with Adaptive Server Enterprise.
- The output of **INTTOHEX** is always 8 digits and is left-padded with zeros; the return data type is VARCHAR.
- The output of **INTTOHEX** does not have a '0x' or '0X' prefix.
- Sybase IQ **INTTOHEX** assumes input is a 32-bit signed integer; a larger value can overflow and a conversion error can result. For example, the statement:

```
SELECT INTTOHEX( 1000000000 ) FROM iq_dummy
```

returns the value 3B9ACA00. But the statement:

```
SELECT INTTOHEX( 10000000000 ) FROM iq_dummy
```

results in a conversion error.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *BIGINTTOHEX Function [Data Type Conversion]* on page 130
- *HEXTOBIGINT Function [Data Type Conversion]* on page 198
- *HEXTOINT Function [Data Type Conversion]* on page 199

ISDATE Function [Date and Time]

Tests whether a string argument can be converted to a date.

If a conversion is possible, the function returns 1; otherwise, it returns 0. If the argument is null, 0 is returned.

Syntax

```
ISDATE ( string )
```

*Parameters***Table 77. Parameters**

Parameter	Description
string	The string to be analyzed to determine whether the string represents a valid date.

Returns

INT

Example

The following example tests whether the `birth_date` column holds valid dates, returning invalid dates as NULL, and valid dates in date format.

```
select birth_date from MyData;
-----
1990/32/89
0101/32/89
1990/12/09

select
  case when isdate(birth_date)=0 then NULL
  else cast(birth_date as date)
  end
from MyData;
-----
```

```
(NULL)
(NULL)
1990-12-09
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase:
 - SQL Anywhere uses ISO 8601 date interchange formats.
 - Not supported by Adaptive Server Enterprise.

ISNULL Function [Miscellaneous]

Returns the value of the first non-NULL expression in the parameter list.

At least two expressions must be passed to the function.

Syntax

```
ISNULL ( expression, expression [ ..., expression ] )
```

Parameters

Table 78. Parameters

Parameter	Description
expression	An expression to be tested against NULL.

Returns

The return type for this function depends on the expressions specified. That is, when the database server evaluates the function, it first searches for a data type in which all the expressions can be compared. When found, the database server compares the expressions and then returns the result in the type used for the comparison. If the database server cannot find a common comparison type, an error is returned.

Example

The following statement returns the value -66:

```
SELECT ISNULL( NULL , -66, 55, 45, NULL, 16 ) FROM iq_dummy
```

Usage

The **ISNULL** function is the same as the **COALESCE** function.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *COALESCE Function [Miscellaneous]* on page 140

ISNUMERIC Function [Miscellaneous]

Tests whether a string argument can be converted to a numeric.

If a conversion is possible, the function returns 1; otherwise, it returns 0. If the argument is null, 0 is returned.

Syntax

```
ISNUMERIC ( string )
```

Note: CIS functional compensation performance considerations apply.

*Parameters***Table 79. Parameters**

Parameter	Description
string	The string to be analyzed to determine whether the string represents a valid numeric value

Returns

INT

Usage

For optimal performance, avoid using **ISNUMERIC** in predicates, where it is processed by the SQL Anywhere portion of the product and cannot take advantage of the performance features of Sybase IQ.

Example

The following example tests whether the `height_in_cms` column holds valid numeric data, returning invalid numeric data as NULL, and valid numeric data in `int` format.

```
data height_in_cms
-----
asde
asde
180
156

select case
  when isnumeric(height_in_cms)=0
  then NULL
  else cast(height_in_cms as int)
end
from MyData
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

LAG Function [Analytical]

An inter-row function that returns the value of an attribute in a previous row in the table or table partition.

Syntax

```
LAG (value_expr) [, offset [, default]] OVER ([PARTITION BY window partition] ORDER BY window ordering)
```

Parameters

Parameter	Description
<i>value_expr</i>	Table column or expression defining the offset data to return from the table.
<i>offset</i>	The number of rows above the current row, expressed as a non-negative exact numeric literal, or as a SQL variable with exact numeric data. The permitted range is 0 to 231.
<i>default</i>	The value to return if the <i>offset</i> value goes beyond the scope of the cardinality of the table or partition.
<i>window partition</i>	(Optional) One or more value expressions separated by commas indicating how you want to divide the set of result rows.
<i>window ordering</i>	Defines the expressions for sorting rows within window partitions, if specified, or within the result set if you did not specify a window partition.

Usage

The **LAG** function requires an **OVER (ORDER BY)** window specification. The window partitioning clause in the **OVER (ORDER BY)** clause is optional. The **OVER (ORDER BY)** clause must not contain a window frame **ROWS/RANGE** specification.

You cannot define an analytic expression in *value_expr*. That is, you cannot nest analytic functions, but you can use other built-in function expressions for *value_expr*.

You must enter a non-negative numeric data type for *offset*. Entering **0** returns the current row. Entering a negative number generates an error.

The default value of *default* is **NULL**. The data type of *default* must be implicitly convertible to the data type of the *value_expr* value or else Sybase IQ generates a conversion error.

Example

The following example returns salary data from the Employees table, partitions the data by department ID, and orders the data according to employee start date. The **LAG** function returns the salary from the previous row (a physical offset of one row) and displays it under the **LAG (Salary)** column:

```
SELECT DepartmentID dID, StartDate, Salary, LAG(Salary, 1)
OVER(PARTITION BY dID ORDER BY StartDate) FROM Employees ORDER BY
1,2;
```

The returned result set is:

dID	StartDate	Salary	Lag(Salary)
100	1984-08-28	45,700.000	NULL
100	1985-01-01	62,000.000	45,700.000
100	1985-06-17	57,490.000	62,000.000
100	1986-06-07	72,995.000	57,490.000
100	1986-07-01	48,023.690	72,995.000
...			
200	1985-02-03	38,500.000	NULL
200	1985-12-06	54,800.000	38,500.000
200	1987-02-19	39,300.000	54,800.000
200	1987-07-10	49,500.000	39,300.000
200	1988-10-04	54,600.000	49,500.000
200	1988-11-12	39,800.000	54,600.000
...			

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

See also

- *LEAD Function [Analytical]* on page 219

LAST_VALUE Function [Aggregate]

Returns the last value from a set of values.

Syntax

```
LAST_VALUE (expression [IGNORE NULLS | RESPECT NULLS])
```

```
OVER (window-spec)
```

Parameters

Parameter	Definition
expression	The expression on which to determine the last value in an ordered set

Returns

Data type of the argument.

Usage

LAST_VALUE returns the last value in a set of values, which is usually an ordered set. If the last value in the set is null, then the function returns NULL unless you specify IGNORE NULLS. If you specify IGNORE NULLS, then **LAST_VALUE** returns the last non-null value in the set, or NULL if all values are null.

The data type of the returned value is the same as that of the input value.

You cannot use **LAST_VALUE** or any other analytic function for expression. That is, you cannot nest analytic functions, but you can use other built-in function expressions for expression.

If the window-spec does not contain an **ORDER BY** expression, or if the **ORDER BY** expression is not precise enough to guarantee a unique ordering, then the result is arbitrary. If there is no window-spec, then the result is arbitrary.

You can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Note: DISTINCT is not supported.

Example

The following example returns the salary of each employee, plus the name of the employee with the highest salary in their department:

```
SELECT GivenName + ' ' + Surname AS employee_name,
       Salary, DepartmentID,
       LAST_VALUE( employee_name ) OVER Salary_Window AS
highest_paid
FROM Employees
WINDOW Salary_Window AS ( PARTITION BY DepartmentID ORDER BY Salary
RANGE BETWEEN UNBOUNDED PRECEDING
AND UNBOUNDED FOLLOWING )
ORDER BY DepartmentID DESC;
```

The returned result set is:

Table 80. LAST_VALUE result set

employee_name	Salary	DepartmentID	highest_paid
Michael Lynch	24,903.000	500	Jose Martinez
Joseph Barker	27,290.000	500	Jose Martinez
Sheila Romero	27,500.000	500	Jose Martinez
Felicia Kuo	28,200.000	500	Jose Martinez
Jeannette Bertrand	29,800.000	500	Jose Martinez
Jane Braun	34,300.000	500	Jose Martinez
Anthony Rebeiro	34,576.000	500	Jose Martinez
Charles Crowley	41,700.000	500	Jose Martinez
Jose Martinez	55,500.800	500	Jose Martinez
Doug Charlton	28,300.000	400	Scott Evans
Elizabeth Lambert	29,384.000	400	Scott Evans
Joyce Butterfield	34,011.000	400	Scott Evans
Robert Nielsen	34,889.000	400	Scott Evans
Alex Ahmed	34,992.000	400	Scott Evans
Ruth Wetherby	35,745.000	400	Scott Evans
...

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

LCASE Function [String]

Converts all characters in a string to lowercase.

Syntax

```
LCASE ( string-expression )
```

*Parameters***Table 81. Parameters**

Parameter	Description
string-expression	The string to be converted to lowercase.

Returns

CHAR

NCHAR

LONG VARCHAR

VARCHAR

NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **LCASE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **LCASE** to the correct data type and size.

Example

The following statement returns the value “lower case”:

```
SELECT LCASE( 'LOWER Case' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—**LCASE** is not supported in Adaptive Server Enterprise; you can use **LOWER** to get the same functionality.

See also

- *LEFT Function [String]* on page 220
- *LOWER Function [String]* on page 228
- *REPLACE Function [String]* on page 279
- *REVERSE Function [String]* on page 283
- *RIGHT Function [String]* on page 284
- *UCASE Function [String]* on page 331
- *UPPER Function [String]* on page 332

LEAD Function [Analytical]

An inter-row function that returns the value of an attribute in a subsequent row in the table or table partition.

Syntax

```
LEAD (value_expr) [, offset [, default]] OVER ([PARTITION BY window partition] ORDER BY window ordering)
```

Parameters

Parameter	Description
<i>value_expr</i>	Table column or expression defining the offset data to return from the table.
<i>offset</i>	The number of rows below the current row, expressed as a non-negative exact numeric literal, or as a SQL variable with exact numeric data. The permitted range is 0 to 231.
<i>default</i>	The value to return if the <i>offset</i> value goes beyond the scope of the table or partition.
<i>window partition</i>	(Optional) One or more value expressions separated by commas indicating how you want to divide the set of result rows.
<i>window ordering</i>	Defines the expressions for sorting rows within window partitions, if specified, or within the result set if you did not specify a window partition.

Usage

The **LEAD** function requires an **OVER (ORDER BY)** window specification. The window partitioning clause in the **OVER (ORDER BY)** clause is optional. The **OVER (ORDER BY)** clause must not contain a window frame **ROWS/RANGE** specification.

You cannot define an analytic expression in *value_expr*. That is, you cannot nest analytic functions, but you can use other built-in function expressions for *value_expr*.

You must enter a non-negative numeric data type for *offset*. Entering **0** returns the current row. Entering a negative number generates an error.

The default value of *default* is **NULL**. The data type of *default* must be implicitly convertible to the data type of the *value_expr* value or else Sybase IQ generates a conversion error.

Example

The following example returns salary data from the Employees table, partitions the data by department ID, and orders the data according to employee start date. The **LEAD** function returns the salary from the next row (a physical offset of one row) and displays it under the **LEAD (Salary)** column:

```
SELECT DepartmentID dID, StartDate, Salary, LEAD(Salary, 1)
OVER(PARTITION BY dID ORDER BY StartDate) FROM Employees ORDER BY
1,2;
```

The returned result set is:

dID	StartDate	Salary	Lead(Salary)
100	1984-08-28	45,700.000	62,000.000
100	1985-01-01	62,000.000	57,490.000
100	1985-06-17	57,490.000	72,995.000
100	1986-06-07	72,995.000	48,023.690
...			
100	1990-08-19	54,900.000	NULL
200	1985-02-03	38,500.000	39,300.000
200	1987-02-19	39,300.000	49,500.000
200	1987-07-10	49,500.000	54,600.000
200	1988-11-28	46,200.000	34,892.000
200	1989-06-01	34,892.000	87,500.000
...			
200	1993-08-12	47,653.000	NULL

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

See also

- *LAG Function [Analytical]* on page 214

LEFT Function [String]

Returns a specified number of characters from the beginning of a string.

Syntax

```
LEFT ( string-expression, numeric-expression )
```

Parameters

Table 82. Parameters

Parameter	Description
string-expression	The string.

Parameter	Description
numeric-expression	The number of characters to return.

Returns

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **LEFT** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **LEFT** to the correct data type and size.

Example

The following statement returns the value “choco”:

```
SELECT LEFT( 'chocolate', 5 ) FROM iq_dummy
```

Usage

If the string contains multibyte characters, and the proper collation is being used, the number of bytes returned may be greater than the specified number of characters.

Note: The result data type of a **LEFT** function is a LONG VARCHAR. If you use **LEFT** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics option license or use **CAST** and set **LEFT** to the correct data type and size.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *LCASE Function [String]* on page 217
- *LOWER Function [String]* on page 228
- *REPLACE Function [String]* on page 279
- *REVERSE Function [String]* on page 283
- *RIGHT Function [String]* on page 284
- *UCASE Function [String]* on page 331
- *UPPER Function [String]* on page 332

LEN Function [String]

Takes one argument as an input of type `BINARY` or `STRING` and returns the number of characters, as defined by the database's collation sequence, of a specified string expression, excluding trailing blanks.

The result may differ from the string's byte length for multi-byte character sets.

`BINARY` and `VARBINARY` are also allowed, in which case `LEN()` returns the number of bytes of the input.

LEN is an alias of **LENGTH** function

Syntax

```
LEN ( string_expr )
```

Parameters

Table 83. Parameters

Parameters	Description
string_expr	The string expression to be evaluated.

Example

The following example returns the value 3152:

```
select len(Photo) from Productswhere ID = 500
```

Usage

This function is the equivalent of **CHAR_LENGTH** (*string_expression*).

Permissions

Any user can execute **LEN**.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.

See also

- *BIT_LENGTH Function [String]* on page 131
- *BYTE_LENGTH Function [String]* on page 132
- *CHAR_LENGTH Function [String]* on page 137
- *COL_LENGTH Function [System]* on page 141
- *DATALENGTH Function [System]* on page 155
- *LENGTH Function [String]* on page 223

- *OBJECT_NAME* Function [System] on page 249
- *OCTET_LENGTH* Function [String] on page 250
- *STR_REPLACE* Function [String] on page 308

LENGTH Function [String]

Returns the number of characters in the specified string.

Syntax

```
LENGTH ( string-expression )
```

Parameters

Table 84. Parameters

Parameter	Description
string-expression	The string.

Returns

INT

Example

The following statement returns the value 9:

```
SELECT LENGTH( 'chocolate' ) FROM iq_dummy
```

Usage

If the string contains multibyte characters, and the proper collation is being used, **LENGTH** returns the number of characters, not the number of bytes. If the string is of **BINARY** data type, the **LENGTH** function behaves as **BYTE_LENGTH**.

The **LENGTH** function is the same as the **CHAR_LENGTH** function.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise. Use the **CHAR_LENGTH** function instead.

See also

- *BIT_LENGTH* Function [String] on page 131
- *BYTE_LENGTH* Function [String] on page 132
- *CHAR_LENGTH* Function [String] on page 137
- *COL_LENGTH* Function [System] on page 141
- *DATALENGTH* Function [System] on page 155

SQL Functions

- *LEN Function [String]* on page 222
- *OBJECT_NAME Function [System]* on page 249
- *OCTET_LENGTH Function [String]* on page 250
- *STR_REPLACE Function [String]* on page 308

LIST function [Aggregate]

Returns a comma-delimited list of values.

Syntax

```
LIST(  
[DISTINCT] string-expression  
[, 'delimiter-string']  
[ORDER BY order-by-expression [ ASC | DESC ], ... ] )
```

Note: CIS functional compensation performance considerations apply.

The **LIST** function [Aggregate] is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > LIST function [Aggregate]*.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **LIST** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **LIST** to the correct data type and size.

LN Function [Numeric]

Returns the natural logarithm of the specified expression.

Syntax

```
LN ( numeric-expression )
```

Parameters

Parameter	Description
expression	Is a column, variable, or expression with a data type that is either exact numeric, approximate numeric, money, or any type that can be implicitly converted to one of these types. For other data types, the LN function generates an error. The return value is of DOUBLE data type.

Usage

LN takes one argument. For example, **LN (20)** returns 2.995732.

The **LN** function is an alias of the **LOG** function.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise. Use the LOG function instead.

See also

- *LOG Function [Numeric]* on page 226
- *LOG10 Function [Numeric]* on page 227

LOCATE Function [String]

Returns the position of one string within another.

Syntax

```
LOCATE ( string-expression1, string-expression2
[ , numeric-expression ] )
```

Parameters

Parameter	Description
string-expression1	The string to be searched.
string-expression2	The string for which you are searching. This string is limited to 255 bytes.
numeric-expression	The character position at which to begin the search in the string. The first character is position 1. If the starting offset is negative, LOCATE returns the last matching string offset, rather than the first. A negative offset indicates how much of the end of the string to exclude from the search. The number of bytes excluded is calculated as $(-1 * \text{offset}) - 1$.

The *numeric-expression* is a 32 bit signed integer for CHAR, VARCHAR, and BINARY columns.

Returns

INT

Examples

The following statement returns the value 8:

```
SELECT LOCATE( 'office party this week - rsvp as soon as possible',
'party', 2 ) FROM iq_dummy
```

SQL Functions

In the second example, the *numeric-expression* starting offset for the search is a negative number.

```
CREATE TABLE t1(name VARCHAR(20), dirname VARCHAR(60));
INSERT INTO t1 VALUES('m1000', 'c:\test\functions\locate.sql');
INSERT INTO t1 VALUES('m1001', 'd:\test\functions\trim.sql');
COMMIT;

SELECT LOCATE(dirname, '\', -1), dirname FROM t1;
```

The result is:

```
18 c:\test\functions\locate.sql
18 d:\test\functions\trim.sql
```

Usage

If *numeric-expression* is specified, the search starts at that offset into the string being searched.

If *numeric-expression* is not specified, **LOCATE** returns only the position of the first instance of the specified string.

The first string can be a long string (longer than 255 bytes), but the second is limited to 255 bytes. If a long string is given as the second argument, the function returns a NULL value.

If any of the arguments is NULL, the result is NULL.

Searching for a zero-length string returns 1.

If the string does not contain the specified string, the **LOCATE** function returns zero (0).

All the positions or offsets, returned or specified, in the **LOCATE** function are always character offsets and may be different from the byte offset for multibyte data.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *PATINDEX Function [String]* on page 250
- *LIKE Conditions* on page 39

LOG Function [Numeric]

Returns the natural logarithm of a number.

LN is an alias of **LOG**.

Syntax

```
LOG ( numeric-expression )
```

*Parameters***Table 85. Parameters**

Parameter	Description
numeric-expression	The number.

Returns

This function converts its argument to DOUBLE, performs the computation in double-precision floating point, and returns a DOUBLE as the result. If the parameter is NULL, the result is NULL.

Example

The following statement returns the value 3.912023:

```
SELECT LOG( 50 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *LN Function [Numeric]* on page 224
- *LOG10 Function [Numeric]* on page 227

LOG10 Function [Numeric]

Returns the base 10 logarithm of a number.

Syntax

```
LOG10 ( numeric-expression )
```

*Parameters***Table 86. Parameters**

Parameter	Description
numeric-expression	The number.

Returns

This function converts its argument to DOUBLE, and performs the computation in double-precision floating point. If the parameter is NULL, the result is NULL.

Example

The following statement returns the value 1.698970.

```
SELECT LOG10( 50 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *LN Function [Numeric]* on page 224
- *LOG Function [Numeric]* on page 226

LOWER Function [String]

Converts all characters in a string to lowercase.

Syntax

```
LOWER ( string-expression )
```

Parameters

Table 87. Parameters

Parameter	Description
string-expression	The string to be converted.

Returns

CHAR

NCHAR

LONG VARCHAR

VARCHAR

NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use LOWER in a SELECT INTO statement, you must have an Unstructured Data Analytics Option license or use CAST and set LOWER to the correct data type and size.

Example

The following statement returns the value “lower case”:

```
SELECT LOWER( 'LOWER Case' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *LCASE Function [String]* on page 217
- *LEFT Function [String]* on page 220
- *REPLACE Function [String]* on page 279
- *REVERSE Function [String]* on page 283
- *RIGHT Function [String]* on page 284
- *UCASE Function [String]* on page 331
- *UPPER Function [String]* on page 332

LTRIM Function [String]

Removes leading blanks from a string.

Syntax

```
LTRIM ( string-expression )
```

Parameters

Parameter	Description
string-expression	The string to be trimmed.

Returns

VARCHAR

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **LTRIM** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **LTRIM** to the correct data type and size.

Example

The following statement returns the value “Test Message” with all leading blanks removed:

SQL Functions

```
SELECT LTRIM( '      Test Message' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *RTRIM Function [String]* on page 289

MAX Function [Aggregate]

Returns the maximum *expression* value found in each group of rows.

Syntax

```
MAX ( expression  
| DISTINCT column-name )
```

Parameters

Parameter	Description
<i>expression</i>	The expression for which the maximum value is to be calculated. This is commonly a column name.
DISTINCT <i>column-name</i>	Returns the same as MAX (<i>expression</i>), and is included for completeness.

Returns

The same data type as the argument.

Example

The following statement returns the value 138948.000, representing the maximum salary in the Employees table:

```
SELECT MAX ( Salary )  
FROM Employees
```

Usage

Rows where *expression* is NULL are ignored. Returns NULL for a group containing no rows.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103
- *MIN Function [Aggregate]* on page 232

MEDIAN Function [Aggregate]

Returns the median of an expression.

Syntax 1

```
MEDIAN([ALL | DISTINCT] expression)
```

Syntax 2

```
MEDIAN([ALL | DISTINCT] expression)
```

```
OVER (window-spec)
```

Parameters

Parameter	Description
<i>expression</i>	A numeric expression for which a median value is to be computed.

Usage

The median is the number separating the higher half of a sample, a population, or a probability distribution, from the lower half.

The data type of the returned value is the same as that of the input value. NULLs are ignored in the calculation of the median value. You can use the optional keyword **DISTINCT** to eliminate duplicate values before the aggregate function is applied. **ALL**, which performs the operation on all rows, is the default.

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Note: The *window-spec* cannot contain a **ROW**, **RANGE** or **ORDER BY** specification; *window-spec* can only specify a **PARTITION** clause. **DISTINCT** is not supported if a **WINDOW** clause is used.

Example

The following query returns the median salary for each department in Florida:

```
SELECT DepartmentID, Surname, Salary,
MEDIAN(Salary) OVER (PARTITION BY DepartmentID) "Median"
FROM Employees
WHERE State IN ('FL')
```

The returned result is:

Table 88. MEDIAN result set

DepartmentID	Surname	Salary	Median
100	Lull	87,900.000	73,870.000
100	Gowda	59,840.000	73,870.000
200	Sterling	64,900.000	76,200.000
200	Kelly	87,500.000	76,200.000
300	Litton	58,930.000	58,930.000
400	Francis	53,870.000	38,70.000
400	Charlton	28,300.000	53,870.000
400	Evans	68,940.000	53,870.000

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

See also

- *Windowing Aggregate Function Usage* on page 103

MIN Function [Aggregate]

Returns the minimum expression value found in each group of rows.

Syntax

```
MIN ( expression
| DISTINCT column-name )
```

Parameters

Parameter	Description
expression	The expression for which the minimum value is to be calculated. This is commonly a column name.
DISTINCT column-name	Returns the same as MIN (expression), and is included for completeness.

Returns

The same data type as the argument.

Example

The following statement returns the value 24903.000, representing the minimum salary in the Employees table:

```
SELECT MIN ( Salary )
FROM Employees
```

Usage

Rows where *expression* is NULL are ignored. Returns NULL for a group containing no rows.

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103
- *MAX Function [Aggregate]* on page 230

MINUTE Function [Date and Time]

Returns a number from 0 to 59 corresponding to the minute component of the specified date/time value.

Syntax

```
MINUTE ( datetime-expression )
```

Parameters

Parameter	Description
<i>datetime-expression</i>	The date/time value.

Returns

SMALLINT

Example

The following statement returns the value 22:

```
SELECT MINUTE( '1998-07-13 12:22:34' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

MINUTES Function [Date and Time]

Returns the number of minutes since an arbitrary date and time, the number of whole minutes between two specified times, or adds the specified integer-expression number of minutes to a time.

Syntax

```
MINUTES ( datetime-expression
| datetime-expression, datetime-expression
| datetime-expression, integer-expression )
```

Parameters

Parameter	Description
<i>datetime-expression</i>	A date and time.
<i>integer-expression</i>	The number of minutes to be added to the <i>datetime-expression</i> . If <i>integer-expression</i> is negative, the appropriate number of minutes are subtracted from the date/time. If you supply an integer expression, the <i>datetime-expression</i> must be explicitly cast as a <code>datetime</code> data type

Returns

INT

TIMESTAMP

Examples

Returns the value 1051125487:

```
SELECT MINUTES( '1998-07-13 06:07:12' ) FROM iq_dummy
```

Returns the value 240, to signify the difference between the two times:

```
SELECT MINUTES( '1999-07-13 06:07:12',
'1999-07-13 10:07:12' ) FROM iq_dummy
```

Returns the datetime value 1999-05-12 21:10:07.000:

```
SELECT MINUTES( CAST( '1999-05-12 21:05:07'
AS DATETIME ), 5) FROM iq_dummy
```

Usage

The second syntax returns the number of whole minutes from the first date/time to the second date/time. The number might be negative.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported in Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *CONVERT Function [Data Type Conversion]* on page 144
- *HOURS Function [Date and Time]* on page 201
- *MONTHS Function [Date and Time]* on page 237
- *REPLACE Function [String]* on page 279
- *SECOND Function [Date and Time]* on page 290
- *WEEKS Function [Date and Time]* on page 340
- *YEAR Function [Date and Time]* on page 345
- *YEARS Function [Date and Time]* on page 346

MOD Function [Numeric]

Returns the remainder when one whole number is divided by another.

Syntax

```
MOD ( dividend, divisor )
```

Parameters

Parameters	Description
<i>dividend</i>	The dividend, or numerator of the division.
<i>divisor</i>	The divisor, or denominator of the division.

Returns

SMALLINT

INT

NUMERIC

Example

The following statement returns the value 2:

```
SELECT MOD( 5, 3 ) FROM iq_dummy
```

Usage

Division involving a negative *dividend* gives a negative or zero result. The sign of the *divisor* has no effect.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported in Adaptive Server Enterprise. The % operator is used as a modulo operator in Adaptive Server Enterprise.

See also

- *REMAINDER Function [Numeric]* on page 278

MONTH Function [Date and Time]

Returns a number from 1 to 12 corresponding to the month of the given date.

Syntax

```
MONTH ( date-expression )
```

Parameters

Parameters	Description
date-expression	A date/time value.

Returns

SMALLINT

Example

The following statement returns the value 7:

```
SELECT MONTH( '1998-07-13' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

MONTHNAME Function [Date and Time]

Returns the name of the month from the specified date expression.

Syntax

```
MONTHNAME ( date-expression )
```

Parameters

Parameter	Description
date-expression	The datetime value.

Returns

VARCHAR

Example

The following statement returns the value **September**, when the DATE_ORDER option is set to the default value of *ymd*.

```
SELECT MONTHNAME( '1998-09-05' ) FROM iq_dummy
```

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > DATE_ORDER Option*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

MONTHS Function [Date and Time]

Returns the number of months since an arbitrary starting date/time or the number of months between two specified date/times, or adds the specified integer-expression number of months to a date/time.

Syntax

```
MONTHS ( date-expression
| date-expression, datetime-expression
| date-expression, integer-expression )
```

Parameters

Parameter	Description
date-expression	A date and time.
integer-expression	The number of months to be added to the <i>date-expression</i> . If <i>integer-expression</i> is negative, the appropriate number of months are subtracted from the date/time value. If you supply an integer expression, the <i>date-expression</i> must be explicitly cast as a datetime data type.

Returns

INT

TIMESTAMP

Examples

The following statement returns the value 23982:

```
SELECT MONTHS( '1998-07-13 06:07:12' ) FROM iq_dummy
```

The following statement returns the value 2, to signify the difference between the two dates:

```
SELECT MONTHS( '1999-07-13 06:07:12',  
              '1999-09-13 10:07:12' ) FROM iq_dummy
```

The following statement returns the datetime value 1999-10-12 21:05:07.000:

```
SELECT MONTHS( CAST( '1999-05-12 21:05:07'  
AS DATETIME ), 5) FROM iq_dummy
```

Usage

The first syntax returns the number of months since an arbitrary starting date. This number is often useful for determining whether two date/time expressions are in the same month in the same year.

```
MONTHS( invoice_sent ) = MONTHS( payment_received )
```

Comparing the **MONTH** function would incorrectly include a payment made 12 months after the invoice was sent.

The second syntax returns the number of months from the first date to the second date. The number might be negative. It is calculated from the number of the first days of the month between the two dates. Hours, minutes and seconds are ignored.

The third syntax adds *integer-expression* months to the given date. If the new date is past the end of the month (such as **MONTHS** ('1992-01-31', 1)) the result is set to the last day of the month. If *integer-expression* is negative, the appropriate number of months are subtracted from the date. Hours, minutes and seconds are ignored.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *CONVERT Function [Data Type Conversion]* on page 144
- *HOURS Function [Date and Time]* on page 201
- *MINUTES Function [Date and Time]* on page 234

- *REPLACE Function [String]* on page 279
- *SECOND Function [Date and Time]* on page 290
- *WEEKS Function [Date and Time]* on page 340
- *YEAR Function [Date and Time]* on page 345
- *YEARS Function [Date and Time]* on page 346

NEWID Function [Miscellaneous]

Generates a UUID (Universally Unique Identifier) value.

The returned UUID value is a binary. A UUID is the same as a GUID (Globally Unique Identifier).

Syntax

```
NEWID ( )
```

Parameters

There are no parameters associated with **NEWID()**.

Returns

UNIQUEIDENTIFIER

Example

The following statement creates the table `t1` and then updates the table, setting the value of the column `uid_col` to a unique identifier generated by the **NEWID** function, if the current value of the column is `NULL`.

```
CREATE TABLE t1 (uid_col int);
UPDATE t1
  SET uid_col = NEWID()
  WHERE uid_col IS NULL
```

If you execute the following statement,

```
SELECT NEWID()
```

the unique identifier is returned as a `BINARY(16)`. For example, the value might be `0xd3749fe09cf446e399913bc6434f1f08`. You can convert this string into a readable format using the **UIDTOSTR()** function.

Usage

The **NEWID()** function generates a unique identifier value.

UUIDs can be used to uniquely identify objects in a database. The values are generated such that a value produced on one computer does not match that produced on another, hence they can also be used as keys in replication and synchronization environments.

The **NEWID** function is supported only in the following positions:

SQL Functions

- **SELECT** list of a top level query block
- **SET** clause of an **UPDATE** statement
- **VALUES** clause of **INSERT...VALUES**

You can use a value generated by the **NEWID** function as a column default value in a Sybase IQ table.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Binary Data Types* on page 78
- *STRTOUUID Function [String]* on page 311
- *UUIDTOSTR Function [String]* on page 334
- *Character Data Types* on page 69
- *Binary Data Types* on page 689

NEXT_CONNECTION Function [System]

Returns the next connection number, or the first connection if the parameter is NULL.

Syntax

```
NEXT_CONNECTION ( { connection-id }, { database-id } )
```

Note: CIS functional compensation performance considerations apply.

Returns

INT

Parameters

Parameter	Description
connection-id	An integer, usually returned from a previous call to NEXT_CONNECTION . If <i>connection-id</i> is NULL, NEXT_CONNECTION returns the most recent connection ID.
database-id	An integer representing one of the databases on the current server. If you supply no <i>database-id</i> , the current database is used. If you supply NULL, then NEXT_CONNECTION returns the next connection regardless of database.

Usage

You can use **NEXT_CONNECTION** to enumerate the connections to a database. To get the first connection, pass `NULL`; to get each subsequent connection, pass the previous return value. The function returns `NULL` when there are no more connections.

NEXT_CONNECTION can be used to enumerate the connections to a database. Connection IDs are generally created in monotonically increasing order. This function returns the next connection ID in reverse order.

To get the connection ID value for the most recent connection, enter `NULL` as the *connection-id*. To get the subsequent connection, enter the previous return value. The function returns `NULL` when there are no more connections in the order.

NEXT_CONNECTION is useful if you want to disconnect all the connections created before a specific time. However, because **NEXT_CONNECTION** returns the connection IDs in reverse order, connections made after the function is started are not returned. If you want to ensure that all connections are disconnected, prevent new connections from being created before you run **NEXT_CONNECTION**.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

Examples

The following statement returns an identifier for the first connection on the current database. The identifier is an integer value like 10.

```
SELECT NEXT_CONNECTION( NULL );
```

The following statement returns a value like 5.

```
SELECT NEXT_CONNECTION( 10 );
```

The following call returns the next connection ID in reverse order from the specified *connection-id* on the current database.

```
SELECT NEXT_CONNECTION( connection-id );
```

The following call returns the next connection ID in reverse order from the specified *connection-id* (regardless of database).

```
SELECT NEXT_CONNECTION( connection-id, NULL );
```

The following call returns the next connection ID in reverse order from the specified *connection-id* on the specified database.

```
SELECT NEXT_CONNECTION( connection-id, database-id );
```

The following call returns the first (earliest) connection (regardless of database).

```
SELECT NEXT_CONNECTION( NULL, NULL );
```

SQL Functions

The following call returns the first (earliest) connection on the specified database.

```
SELECT NEXT_CONNECTION( NULL, database-id );
```

NEXT_DATABASE Function [System]

Returns the next database ID number, or the first database if the parameter is NULL.

Syntax

```
NEXT_DATABASE ( { NULL | database-id } )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Parameter	Description
database-id	An integer that specifies the ID number of the database.

Returns

INT

Examples

The following statement returns the value 0, the first database value:

```
SELECT NEXT_DATABASE( NULL ) FROM iq_dummy
```

The following statement returns NULL, indicating that there are no more databases on the server:

```
SELECT NEXT_DATABASE( 0 ) FROM iq_dummy
```

Usage

You can use **NEXT_DATABASE** to enumerate the databases running on a database server. To get the first database, pass NULL; to get each subsequent database, pass the previous return value. The function returns NULL when there are no more databases.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *COL_NAME Function [System]* on page 141
- *DB_ID Function [System]* on page 176
- *DB_NAME Function [System]* on page 177
- *DB_PROPERTY Function [System]* on page 178

- *OBJECT_ID Function [System]* on page 248
- *OBJECT_NAME Function [System]* on page 249

NEXT_HTTP_HEADER Function [HTTP]

Gets the next HTTP header name.

Syntax

```
NEXT_HTTP_HEADER ( header-name )
```

Note: CIS functional compensation performance considerations apply.

The `NEXT_HTTP_HEADER` function is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > NEXT_HTTP_HEADER function [HTTP]*.

Note: This reference points to SQL Anywhere documentation.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use `NEXT_HTTP_HEADER` in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set `NEXT_HTTP_HEADER` to the correct data type and size.

NEXT_HTTP_VARIABLE Function [HTTP]

Get the next HTTP variable name.

Syntax

```
NEXT_HTTP_VARIABLE ( var-name )
```

Note: CIS functional compensation performance considerations apply.

The `NEXT_HTTP_VARIABLE` function is a SQL Anywhere function. See *SQL Anywhere Server - SQL Reference > SQL functions > Functions > NEXT_HTTP_VARIABLE function [HTTP]*.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use `NEXT_HTTP_VARIABLE` in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set `NEXT_HTTP_VARIABLE` to the correct data type and size.

NOW Function [Date and Time]

Returns the current date and time. This is the historical syntax for **CURRENT_TIMESTAMP**.

Syntax

```
NOW ( * )
```

Returns

TIMESTAMP

Example

The following statement returns the current date and time.

```
SELECT NOW( * ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

NTILE Function [Analytical]

Distributes query results into a specified number of buckets and assigns the bucket number to each row in the bucket.

Syntax

```
NTILE ( expression1 )  
OVER ( ORDER BY expression2 [ ASC | DESC ] )
```

Parameters

Parameter	Description
expression1	A constant integer from 1 to 32767, which specifies the number of buckets.
expression2	A sort specification that can be any valid expression involving a column reference, aggregates, or expressions invoking these items.

Example

The following example uses the **NTILE** function to determine the sale status of car dealers. The dealers are divided into four groups based on the number of cars each dealer sold. The dealers with `ntile = 1` are in the top 25% for car sales.

```
SELECT dealer_name, sales,  
NTILE(4) OVER ( ORDER BY sales DESC )  
FROM carSales;
```

dealer_name	sales	ntile
Boston	1000	1
Worcester	950	1
Providence	950	1
SF	940	1
Lowell	900	2
Seattle	900	2
Natick	870	2
New Haven	850	2
Portland	800	3
Houston	780	3
Hartford	780	3
Dublin	750	3
Austin	650	4
Dallas	640	4
Dover	600	4

To find the top 10% of car dealers by sales, you specify **NTILE(10)** in the example **SELECT** statement. Similarly, to find the top 50% of car dealers by sales, specify **NTILE(2)**.

Usage

NTILE is a rank analytical function that distributes query results into a specified number of buckets and assigns the bucket number to each row in the bucket. You can divide a result set into one-hundredths (percentiles), tenths (deciles), fourths (quartiles), or other numbers of groupings.

NTILE requires an **OVER (ORDER BY)** clause. The **ORDER BY** clause specifies the parameter on which ranking is performed and the order in which the rows are sorted in each group. This **ORDER BY** clause is used only within the **OVER** clause and is *not* an **ORDER BY** for the **SELECT**. No aggregation functions in the rank query are allowed to specify **DISTINCT**.

The **OVER** clause indicates that the function operates on a query result set. The result set is the rows that are returned after the **FROM**, **WHERE**, **GROUP BY**, and **HAVING** clauses have all been evaluated. The **OVER** clause defines the data set of the rows to include in the computation of the rank analytical function.

The **ASC** or **DESC** parameter specifies the ordering sequence ascending or descending. Ascending order is the default.

NTILE is allowed only in the select list of a **SELECT** or **INSERT** statement or in the **ORDER BY** clause of the **SELECT** statement. **NTILE** can be in a view or a union. The **NTILE** function cannot be used in a subquery, a **HAVING** clause, or in the select list of an **UPDATE** or **DELETE** statement. Only one **NTILE** function is allowed per query.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere

See also

- *PERCENTILE_CONT* Function [Analytical] on page 254
- *PERCENTILE_DISC* Function [Analytical] on page 256
- *YEAR* Function [Date and Time] on page 345

NULLIF Function [Miscellaneous]

Provides an abbreviated **CASE** expression by comparing expressions.

Syntax

```
NULLIF ( expression1, expression2 )
```

Parameters

Parameter	Description
expression1	An expression to be compared.
expression2	An expression to be compared.

Returns

Data type of the first argument.

Examples

The following statement returns a:

```
SELECT NULLIF( 'a', 'b' ) FROM iq_dummy
```

The following statement returns NULL:

```
SELECT NULLIF( 'a', 'a' ) FROM iq_dummy
```

Usage

NULLIF compares the values of the two expressions.

If the first expression equals the second expression, **NULLIF** returns NULL.

If the first expression does not equal the second expression, or if the second expression is NULL, **NULLIF** returns the first expression.

The **NULLIF** function provides a short way to write some **CASE** expressions. **NULLIF** is equivalent to:

```
CASE WHEN expression1 = expression2 THEN NULL  
ELSE expression1 END
```

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *CASE Expressions* on page 28
- *NULLIF Function for Abbreviated CASE Expressions* on page 29

NUMBER Function [Miscellaneous]

Generates numbers starting at 1 for each successive row in the results of the query.

Syntax

```
NUMBER ( * )
```

Returns

INT

Usage

Use the **NUMBER** function only in a select list or a **SET** clause of an **UPDATE** statement. For example, the following statement updates each row of the `seq_id` column with a number 1 greater than the previous row. The number is applied in the order specified by the **ORDER BY** clause.

```
update empl
set seq_id = number(*)
order by empl_id
```

In an **UPDATE** statement, if the **NUMBER(*)** function is used in the **SET** clause and the **FROM** clause specifies a one-to-many join, **NUMBER(*)** generates unique numbers that increase, but may not increment sequentially due to row elimination.

NUMBER can also be used to generate primary keys when using the **INSERT** from **SELECT** statement, although using **IDENTITY/AUTOINCREMENT** is a preferred mechanism for generating sequential primary keys.

Note: A syntax error is generated if you use **NUMBER** in a **DELETE** statement, **WHERE** clause, **HAVING** clause, **ORDER BY** clause, subquery, query involving aggregation, any constraint, **GROUP BY**, **DISTINCT**, a query containing **UNION ALL**, or a derived table.

Example

The following statement returns this numbered list:

number(*)
1
2
3
4

number(*)
5

```
SELECT NUMBER( * )
FROM Departments
WHERE DepartmentID > 10
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

OBJECT_ID Function [System]

Returns the object ID.

Syntax

```
OBJECT_ID ( object-name )
```

Parameters

Table 89. Parameters

Parameter	Description
object-name	The name of the object.

Examples

The following statement returns the object ID 100209 of the *Customers* table:

```
SELECT OBJECT_ID ('CUSTOMERS') FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *COL_NAME Function [System]* on page 141
- *DB_ID Function [System]* on page 176
- *DB_NAME Function [System]* on page 177
- *DB_PROPERTY Function [System]* on page 178
- *NEXT_DATABASE Function [System]* on page 242
- *OBJECT_NAME Function [System]* on page 249
- *INDEX_COL Function [System]* on page 208

OBJECT_NAME Function [System]

Returns the object name.

Syntax

```
OBJECT_NAME ( object-id [ , database-id ] )
```

*Parameters***Table 90. Parameters**

Parameter	Description
object-id	The object ID.
database-id	The database ID.

Examples

The following statement returns the name “customer:”

```
SELECT OBJECT_NAME ( 100209 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *BIT_LENGTH Function [String]* on page 131
- *BYTE_LENGTH Function [String]* on page 132
- *CHAR_LENGTH Function [String]* on page 137
- *COL_LENGTH Function [System]* on page 141
- *DATALENGTH Function [System]* on page 155
- *LEN Function [String]* on page 222
- *LENGTH Function [String]* on page 223
- *OCTET_LENGTH Function [String]* on page 250
- *STR_REPLACE Function [String]* on page 308
- *COL_NAME Function [System]* on page 141
- *DB_ID Function [System]* on page 176
- *DB_NAME Function [System]* on page 177
- *DB_PROPERTY Function [System]* on page 178
- *NEXT_DATABASE Function [System]* on page 242
- *OBJECT_ID Function [System]* on page 248

OCTET_LENGTH Function [String]

Returns an unsigned 64-bit value containing the byte length of the column parameter.

Syntax

```
OCTET_LENGTH( column-name )
```

Parameters

Parameter	Description
column-name	The name of a column.

Usage

The return value of a NULL argument is NULL.

The **OCTET_LENGTH** function supports all Sybase IQ data types.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

Standards and Compatibility

- Sybase—Not supported by SQL Anywhere or Adaptive Server Enterprise.

See also

- *BIT_LENGTH Function [String]* on page 131
- *BYTE_LENGTH Function [String]* on page 132
- *CHAR_LENGTH Function [String]* on page 137
- *COL_LENGTH Function [System]* on page 141
- *DATALENGTH Function [System]* on page 155
- *LEN Function [String]* on page 222
- *LENGTH Function [String]* on page 223
- *OBJECT_NAME Function [System]* on page 249
- *STR_REPLACE Function [String]* on page 308

PATINDEX Function [String]

Returns the starting position of the first occurrence of a specified pattern.

Syntax

```
PATINDEX ( '%pattern%', string-expression )
```

Parameters

Parameter	Description
pattern	The pattern for which you are searching. This string is limited to 126 bytes for patterns with wildcards. If the leading percent wildcard is omitted, PATINDEX returns one (1) if the pattern occurs at the beginning of the string, and zero if not. If <i>pattern</i> starts with a percent wildcard, then the two leading percent wildcards are treated as one. Patterns without wildcards (percent % or underscore _) can be up to 255 bytes in length.
string-expression	The string to be searched for the pattern.

Returns

INT

Examples

The following statement returns the value 2:

```
SELECT PATINDEX( '%hoco%', 'chocolate' ) FROM iq_dummy
```

The following statement returns the value 11:

```
SELECT PATINDEX( '%4_5_', '0a1A 2a3A 4a5A' ) FROM iq_dummy
```

Usage

PATINDEX returns the starting position of the first occurrence of the pattern. If the string being searched contains more than one instance of the string pattern, **PATINDEX** returns only the position of the first instance.

The pattern uses the same wildcards as the **LIKE** comparison. This table lists the pattern wildcards.

Table 91. PATINDEX pattern wildcards

Wildcard	Matches
_ (underscore)	Any one character
% (percent)	Any string of zero or more characters
[]	Any single character in the specified range or set
[^]	Any single character not in the specified range or set

SQL Functions

If the pattern is not found, **PATINDEX** returns zero (0).

Searching for a pattern longer than 126 bytes returns NULL.

Searching for a zero-length string returns 1.

If any of the arguments is NULL, the result is zero (0).

All the positions or offsets, returned or specified, in the **PATINDEX** function are always character offsets and may be different from the byte offset for multibyte data.

PATINDEX returns a 32 bit unsigned integer position for CHAR and VARCHAR columns.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *LIKE Conditions* on page 39
- *LOCATE Function [String]* on page 225

PERCENT_RANK Function [Analytical]

Computes the (fractional) position of one row returned from a query with respect to the other rows returned by the query, as defined by the **ORDER BY** clause.

Returns a decimal value between 0 and 1.

Syntax

```
PERCENT_RANK ( ) OVER ( ORDER BY expression [ ASC | DESC ] )
```

Parameters

Parameter	Description
expression	A sort specification that can be any valid expression involving a column reference, aggregates, or expressions invoking these items.

Returns

The **PERCENT_RANK** function returns a DOUBLE value between 0 and 1.

Example

The following statement illustrates the use of the **PERCENT_RANK** function:

```
SELECT s_suppkey, SUM(s_acctBal) AS sum_acctBal,
PERCENT_RANK() OVER ( ORDER BY SUM(s_acctBal) DESC )
AS percent_rank_all FROM supplier GROUP BY s_suppkey;
```

s_suppkey	sum_acctBal	percent_rank_all
supplier#011	200000	0
supplier#002	200000	0
supplier#013	123000	0.3333
supplier#004	110000	0.5
supplier#035	110000	0.5
supplier#006	50000	0.8333
supplier#021	10000	1

Usage

PERCENT_RANK is a rank analytical function. The percent rank of a row *R* is defined as the rank of a row in the groups specified in the **OVER** clause minus one divided by the number of total rows in the groups specified in the **OVER** clause minus one. **PERCENT_RANK** returns a value between 0 and 1. The first row has a percent rank of zero.

The **PERCENT_RANK** of a row is calculated as

$$(R_x - 1) / (N_{totalRow} - 1)$$

where *R_x* is the rank position of a row in the group and *N_{totalRow}* is the total number of rows in the group specified by the **OVER** clause.

PERCENT_RANK requires an **OVER (ORDER BY)** clause. The **ORDER BY** clause specifies the parameter on which ranking is performed and the order in which the rows are sorted in each group. This **ORDER BY** clause is used only within the **OVER** clause and is *not* an **ORDER BY** for the **SELECT**. No aggregation functions in the rank query are allowed to specify **DISTINCT**.

The **OVER** clause indicates that the function operates on a query result set. The result set is the rows that are returned after the **FROM**, **WHERE**, **GROUP BY**, and **HAVING** clauses have all been evaluated. The **OVER** clause defines the data set of the rows to include in the computation of the rank analytical function.

The ASC or DESC parameter specifies the ordering sequence ascending or descending. Ascending order is the default.

PERCENT_RANK is allowed only in the select list of a **SELECT** or **INSERT** statement or in the **ORDER BY** clause of the **SELECT** statement. **PERCENT_RANK** can be in a view or a union. The **PERCENT_RANK** function cannot be used in a subquery, a **HAVING** clause, or in the select list of an **UPDATE** or **DELETE** statement. Only one rank analytical function is allowed per query.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

PERCENTILE_CONT Function [Analytical]

Given a percentile, returns the value that corresponds to that percentile. Assumes a continuous distribution data model.

Note: If you are simply trying to compute a percentile, use the **NTILE** function instead, with a value of 100.

Syntax

```
PERCENTILE_CONT ( expression1 )
WITHIN GROUP ( ORDER BY expression2 [ ASC | DESC ] )
```

Parameters

Parameter	Description
<i>expression1</i>	A constant of numeric data type and range from 0 to 1 (inclusive). If the argument is NULL, a “wrong argument for percentile” error is returned. If the argument value is less than 0 or greater than 1, a “data value out of range” error is returned.
<i>expression2</i>	A sort specification that must be a single expression involving a column reference. Multiple expressions are not allowed and no rank analytical functions, set functions, or subqueries are allowed in this sort expression.

Example

The following example uses the **PERCENTILE_CONT** function to determine the 10th percentile value for car sales in a region.

The following data set is used in the example:

sales	region	dealer_name
900	Northeast	Boston
800	Northeast	Worcester
800	Northeast	Providence
700	Northeast	Lowell
540	Northeast	Natick
500	Northeast	New Haven
450	Northeast	Hartford
800	Northwest	SF
600	Northwest	Seattle
500	Northwest	Portland
400	Northwest	Dublin
500	South	Houston
400	South	Austin

300	South	Dallas
200	South	Dover

The following **SELECT** statement contains the **PERCENTILE_CONT** function:

```
SELECT region, PERCENTILE_CONT(0.1)
WITHIN GROUP ( ORDER BY sales DESC )
FROM carSales GROUP BY region;
```

The result of the **SELECT** statement lists the 10th percentile value for car sales in a region:

region	percentile_cont
Northeast	840
Northwest	740
South	470

Usage

The inverse distribution analytical functions return a k-th percentile value, which can be used to help establish a threshold acceptance value for a set of data. The function **PERCENTILE_CONT** takes a percentile value as the function argument, and operates on a group of data specified in the **WITHIN GROUP** clause, or operates on the entire data set. The function returns one value per group. If the **GROUP BY** column from the query is not present, the result is a single row. The data type of the results is the same as the data type of its **ORDER BY** item specified in the **WITHIN GROUP** clause. The data type of the **ORDER BY** expression for **PERCENTILE_CONT** must be numeric.

PERCENTILE_CONT requires a **WITHIN GROUP (ORDER BY)** clause.

The **ORDER BY** clause, which must be present, specifies the expression on which the percentile function is performed and the order in which the rows are sorted in each group. For the **PERCENTILE_CONT** function, the data type of this expression must be numeric. This **ORDER BY** clause is used only within the **WITHIN GROUP** clause and is *not* an **ORDER BY** for the **SELECT**.

The **WITHIN GROUP** clause distributes the query result into an ordered data set from which the function calculates a result. The **WITHIN GROUP** clause must contain a single sort item. If the **WITHIN GROUP** clause contains more or less than one sort item, an error is reported.

The ASC or DESC parameter specifies the ordering sequence ascending or descending. Ascending order is the default.

The **PERCENTILE_CONT** function is allowed in a subquery, a **HAVING** clause, a view, or a union. **PERCENTILE_CONT** can be used anywhere the simple nonanalytical aggregate functions are used. The **PERCENTILE_CONT** function ignores the NULL value in the data set.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

See also

- *NTILE Function [Analytical]* on page 244
- *PERCENTILE_DISC Function [Analytical]* on page 256

PERCENTILE_DISC Function [Analytical]

Given a percentile, returns the value that corresponds to that percentile. Assumes a discrete distribution data model.

Note: If you are simply trying to compute a percentile, use the **NTILE** function instead, with a value of 100.

Syntax

```
PERCENTILE_DISC ( expression1 )
WITHIN GROUP ( ORDER BY expression2 [ ASC | DESC ] )
```

Parameters

Parameter	Description
expression1	A constant of numeric data type and range from 0 to 1 (inclusive). If the argument is NULL, then a “wrong argument for percentile” error is returned. If the argument value is less than 0 or greater than 1, then a “data value out of range” error is returned.
expression2	A sort specification that must be a single expression involving a column reference. Multiple expressions are not allowed and no rank analytical functions, set functions, or subqueries are allowed in this sort expression.

Example

The following example uses the **PERCENTILE_DISC** function to determine the 10th percentile value for car sales in a region.

The following data set is used in the example:

```
sales      region      dealer_name
900        Northeast   Boston
800        Northeast   Worcester
800        Northeast   Providence
700        Northeast   Lowell
540        Northeast   Natick
500        Northeast   New Haven
450        Northeast   Hartford
800        Northwest   SF
600        Northwest   Seattle
500        Northwest   Portland
```

400	Northwest	Dublin
500	South	Houston
400	South	Austin
300	South	Dallas
200	South	Dover

The following **SELECT** statement contains the **PERCENTILE_DISC** function:

```
SELECT region, PERCENTILE_DISC(0.1)
WITHIN GROUP ( ORDER BY sales DESC )
FROM carSales GROUP BY region;
```

The result of the **SELECT** statement lists the 10th percentile value for car sales in a region:

region	percentile_cont
Northeast	900
Northwest	800
South	500

Usage

The inverse distribution analytical functions return a k-th percentile value, which can be used to help establish a threshold acceptance value for a set of data. The function **PERCENTILE_DISC** takes a percentile value as the function argument and operates on a group of data specified in the **WITHIN GROUP** clause or operates on the entire data set. The function returns one value per group. If the **GROUP BY** column from the query is not present, the result is a single row. The data type of the results is the same as the data type of its **ORDER BY** item specified in the **WITHIN GROUP** clause. **PERCENTILE_DISC** supports all data types that can be sorted in Sybase IQ.

PERCENTILE_DISC requires a **WITHIN GROUP (ORDER BY)** clause.

The **ORDER BY** clause, which must be present, specifies the expression on which the percentile function is performed and the order in which the rows are sorted in each group. This **ORDER BY** clause is used only within the **WITHIN GROUP** clause and is *not* an **ORDER BY** for the **SELECT**.

The **WITHIN GROUP** clause distributes the query result into an ordered data set from which the function calculates a result. The **WITHIN GROUP** clause must contain a single sort item. If the **WITHIN GROUP** clause contains more or less than one sort item, an error is reported.

The ASC or DESC parameter specifies the ordering sequence ascending or descending. Ascending order is the default.

The **PERCENTILE_DISC** function is allowed in a subquery, a **HAVING** clause, a view, or a union. **PERCENTILE_DISC** can be used anywhere the simple nonanalytical aggregate functions are used. The **PERCENTILE_DISC** function ignores the NULL value in the data set.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

See also

- *NTILE Function [Analytical]* on page 244
- *PERCENTILE_CONT Function [Analytical]* on page 254

PI Function [Numeric]

Returns the numeric value PI.

Syntax

```
PI ( * )
```

Returns

DOUBLE

Example

The following statement returns the value 3.141592653....

```
SELECT PI ( * ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—The **PI()** function is supported in Adaptive Server Enterprise, but **PI(*)** is not.

POWER Function [Numeric]

Calculates one number raised to the power of another.

Syntax

```
POWER ( numeric-expression1, numeric-expression2 )
```

Parameters

Parameter	Description
numeric-expression1	The base.
numeric-expression2	The exponent.

Returns

DOUBLE

Example

The following statement returns the value 64:

```
SELECT Power( 2, 6 ) FROM iq_dummy
```

Usage

Raises *numeric-expression1* to the power *numeric-expression2*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

PROPERTY Function [System]

Returns the value of the specified server-level property as a string.

Syntax

```
PROPERTY ( { property-id | property-name } )
```

Note: CIS functional compensation performance considerations apply.

*Parameters***Table 92.**

Parameter	Description
<i>property-id</i>	An integer that is the property-number of the server-level property. This number can be determined from the PROPERTY_NUMBER function. The <i>property-id</i> is commonly used when looping through a set of properties.
<i>property-name</i>	A string giving the name of the property.

Returns

VARCHAR

Example

The following statement returns the name of the current database server:

```
SELECT PROPERTY( 'Name' ) FROM iq_dummy
```

Usage

Each property has both a number and a name, but the number is subject to change between versions, and should not be used as a reliable identifier for a given property.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *PROPERTY_NAME Function [System]* on page 261
- *PROPERTY_NUMBER Function [System]* on page 261
- *CONNECTION_PROPERTY Function [System]* on page 142
- *Properties Available for the Server* on page 120
- *Properties Available for each Database* on page 120
- *Connection Properties* on page 119

PROPERTY_DESCRIPTION Function [System]

Returns a description of a property.

Syntax

```
PROPERTY_DESCRIPTION ( { property-id | property-name } )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Parameter	Description
<i>property-id</i>	An integer that is the property number of the property. This number can be determined from the PROPERTY_NUMBER function. The <i>property-id</i> is commonly used when looping through a set of properties.
<i>property-name</i>	A string giving the name of the property.

Returns

VARCHAR

Example

The following statement returns the description “Number of index insertions:”

```
SELECT PROPERTY_DESCRIPTION( 'IndAdd' ) FROM iq_dummy
```

Usage

Each property has both a number and a name, but the number is subject to change between releases, and should not be used as a reliable identifier for a given property.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

- Sybase—Not supported by Adaptive Server Enterprise.

PROPERTY_NAME Function [System]

Returns the name of the property with the supplied property number.

Syntax

```
PROPERTY_NAME ( property-id )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Parameter	Description
property-id	The property number of the property.

Returns

VARCHAR

Example

The following statement returns the property associated with property number 126. The actual property to which this refers changes from version to version.

```
SELECT PROPERTY_NAME( 126 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *PROPERTY Function [System]* on page 259
- *PROPERTY_NUMBER Function [System]* on page 261
- *CONNECTION_PROPERTY Function [System]* on page 142
- *Properties Available for the Server* on page 120
- *Properties Available for each Database* on page 120
- *Connection Properties* on page 119

PROPERTY_NUMBER Function [System]

Returns the property number of the property with the supplied property name.

Syntax

```
PROPERTY_NUMBER ( property-name )
```

Note: CIS functional compensation performance considerations apply.

Parameters

Parameter	Description
property-name	A property name.

Returns

INT

Example

The following statement returns an integer value. The actual value changes from version to version.

```
SELECT PROPERTY_NUMBER( 'PAGESIZE' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *PROPERTY Function [System]* on page 259
- *PROPERTY_NAME Function [System]* on page 261
- *CONNECTION_PROPERTY Function [System]* on page 142
- *Properties Available for the Server* on page 120
- *Properties Available for each Database* on page 120
- *Connection Properties* on page 119

QUARTER Function [Date and Time]

Returns a number indicating the quarter of the year from the supplied date expression.

Syntax

```
QUARTER( date-expression )
```

Parameters

Parameter	Description
date-expression	A date.

Returns

INT

Example

With the **DATE_ORDER** option set to the default of *ymd*, the following statement returns the value 2:

```
SELECT QUARTER ( '1987/05/02' ) FROM iq_dummy
```

See *Reference: Statements and Options > Database Options > Alphabetical List of Options > DATE_ORDER Option*.

Usage

This table lists the dates in the quarters of the year.

Table 93. Values of quarter of the year

Quarter	Period (inclusive)
1	January 1 to March 31
2	April 1 to June 30
3	July 1 to September 30
4	October 1 to December 31

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

RADIANS Function [Numeric]

Converts a number from degrees to radians.

Syntax

```
RADIANS ( numeric-expression )
```

Parameters

Parameter	Description
numeric-expression	A number, in degrees. This angle is converted to radians

Returns

DOUBLE

Example

The following statement returns a value of approximately 0.5236:

```
SELECT RADIANS( 30 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

RAND Function [Numeric]

Returns a DOUBLE precision, random number x, where 0 <= x <1, with an optional seed.

Syntax

```
RAND ( [ integer-expression ] )
```

Parameters

Parameter	Description
integer-expression	The optional seed used to create a random number. This argument allows you to create repeatable random number sequences.

Returns

DOUBLE

Examples

The following statement returns a 5% sampling of a table:

```
SELECT AVG(table1.number_of_cars), AVG(table1.number_of_tv's)FROM
table1 WHERE RAND(ROWID(table1)) < .05 and table1.income < 50000;
```

The following statement returns a value of approximately 941392926249216914:

```
SELECT RAND( 4 ) FROM iq_dummy
```

Usage

If **RAND** is called with a **FROM** clause and an argument in a query containing only tables in IQ stores, the function returns an arbitrary but repeatable value.

When no argument is called, **RAND** is a non-deterministic function. Successive calls to **RAND** might return different values. The query optimizer does not cache the results of the **RAND** function

Note: The values returned by **RAND** vary depending on whether you use a **FROM** clause or not and whether the referenced table was created in **SYSTEM** or in an **IQ** store.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

- Sybase—Compatible with Adaptive Server Enterprise.

RANK Function [Analytical]

Ranks items in a group.

Syntax

```
RANK ( ) OVER ( [ PARTITION BY ] ORDER BY expression [ ASC | DESC ] )
```

Parameters

Parameter	Description
expression	A sort specification that can be any valid expression involving a column reference, aggregates, or expressions invoking these items.

Returns

INTEGER

Example

This statement illustrates the use of the **RANK** function:

```
SELECT Surname, Sex, Salary, RANK() OVER (PARTITION BY Sex
ORDER BY Salary DESC) AS RANK FROM Employees
WHERE State IN ('CA', 'AZ') AND DepartmentID IN (200, 300)
ORDER BY Sex, Salary DESC;
```

The results from the above query:

Surname	Sex	Salary	RANK
Savarino	F	72300.000	1
Jordan	F	51432.000	2
Clark	F	45000.000	3
Coleman	M	42300.000	1
Overbey	M	39300.000	2

Usage

RANK is a rank analytical function. The rank of row R is defined as the number of rows that precede R and are not peers of R. If two or more rows are not distinct within the groups specified in the **OVER** clause or distinct over the entire result set, then there are one or more gaps in the sequential rank numbering. The difference between **RANK** and **DENSE_RANK** is that **DENSE_RANK** leaves no gap in the ranking sequence when there is a tie. **RANK** leaves a gap when there is a tie.

RANK requires an **OVER (ORDER BY)** clause. The **ORDER BY** clause specifies the parameter on which ranking is performed and the order in which the rows are sorted in each group. This **ORDER BY** clause is used only within the **OVER** clause and is *not* an **ORDER BY** for the **SELECT**. No aggregation functions in the rank query are allowed to specify **DISTINCT**.

The **PARTITION BY** window partitioning clause in the **OVER (ORDER BY)** clause is optional.

The **ASC** or **DESC** parameter specifies the ordering sequence ascending or descending. Ascending order is the default.

The **OVER** clause indicates that the function operates on a query result set. The result set is the rows that are returned after the **FROM**, **WHERE**, **GROUP BY**, and **HAVING** clauses have all been evaluated. The **OVER** clause defines the data set of the rows to include in the computation of the rank analytical function.

RANK is allowed only in the select list of a **SELECT** or **INSERT** statement or in the **ORDER BY** clause of the **SELECT** statement. **RANK** can be in a view or a union. The **RANK** function cannot be used in a subquery, a **HAVING** clause, or in the select list of an **UPDATE** or **DELETE** statement. Only one rank analytical function is allowed per query.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise or SQL Anywhere.

See also

- *DENSE_RANK Function [Analytical]* on page 180

REGR_AVGX Function [Aggregate]

Computes the average of the independent variable of the regression line.

Syntax 1

```
REGR_AVGX (dependent-expression, independent-expression)
```

Syntax 2

```
REGR_AVGX (dependent-expression, independent-expression)
```

```
OVER (window-spec)
```

Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Returns

DOUBLE

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then **REGR_AVGX** returns NULL.

The function is applied to the set of (dependent-expression and independent-expression) pairs after eliminating all pairs for which either dependent-expression or independent-expression is NULL. The function is computed simultaneously during a single pass through the data. After eliminating NULL values, the following computation is then made, where x represents the independent-expression:

```
AVG ( x )
```

See *SQL Anywhere Server - SQL Usage > Querying and Modifying Data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions*.

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example calculates the average of the dependent variable, employee age:

```
SELECT REGR_AVGX( Salary, ( YEAR( NOW() ) - YEAR( BirthDate ) ) )FROM
Employees;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

REGR_AVGY Function [Aggregate]

Computes the average of the dependent variable of the regression line.

Syntax 1

```
REGR_AVGY( dependent-expression, independent-expression )
```

Syntax 2

REGR_AVGY (*dependent-expression*, *independent-expression*)

OVER (*window-spec*)

Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Returns

DOUBLE

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then **REGR_AVGY** returns NULL.

The function is applied to the set of (dependent-expression and independent-expression) pairs after eliminating all pairs for which either dependent-expression or independent-expression is NULL. The function is computed simultaneously during a single pass through the data. After eliminating NULL values, the following computation is then made, where y represents the dependent-expression:

$AVG(y)$

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions.*

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example calculates the average of the independent variable, employee salary. This function returns the value 49988.6232:

```
SELECT REGR_AVGY( Salary, ( YEAR( NOW( ) ) - YEAR( BirthDate ) ) )FROM
Employees;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

REGR_COUNT Function [Aggregate]

Returns an integer that represents the number of non-NULL number pairs used to fit the regression line.

Syntax 1

```
REGR_COUNT(dependent-expression, independent-expression)
```

Syntax 2

```
REGR_COUNT(dependent-expression, independent-expression)
```

```
OVER (window-spec)
```

Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Returns

INTEGER

Usage

This function returns an UNSIGNED BIGINT as the result.

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example returns a value that indicates the number of non-NULL pairs that were used to fit the regression line. This function returns the value 75:

```
SELECT REGR_COUNT( Salary, ( YEAR( NOW() ) -  
YEAR( BirthDate ) ) )FROM Employees;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

REGR_INTERCEPT Function [Aggregate]

Computes the y-intercept of the linear regression line that best fits the dependent and independent variables.

Syntax 1

```
REGR_INTERCEPT(dependent-expression, independent-expression)
```

Syntax 2

```
REGR_INTERCEPT(dependent-expression, independent-expression)
```

```
OVER (window-spec)
```

Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Returns

DOUBLE

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, **REGR_INTERCEPT** returns NULL.

The function is applied to the set of (*dependent-expression* and *independent-expression*) pairs after eliminating all pairs for which either *dependent-expression* or *independent-expression* is NULL. The function is computed simultaneously during a single pass through the data. After eliminating NULL values, the following computation is made, where y represents the *dependent-expression* and x represents the *independent-expression*:

$$AVG(y) - REGR_SLOPE(y, x) * AVG(x)$$

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions.*

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example returns the value 1874.5805688517603:

```
SELECT REGR_INTERCEPT( Salary, ( YEAR( NOW() ) -
YEAR( BirthDate ) ) )FROM Employees;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

REGR_R2 Function [Aggregate]

Computes the coefficient of determination (also referred to as R-squared or the goodness-of-fit statistic) for the regression line.

Syntax 1

```
REGR_R2(dependent-expression, independent-expression)
```

Syntax 2

```
REGR_R2(dependent-expression, independent-expression)
```

```
OVER (window-spec)
```

Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Returns

DOUBLE

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then **REGR_R2** returns NULL.

REGR_R2 is applied to the set of (dependent-expression and *independent-expression*) pairs after eliminating all pairs for which either dependent-expression or *independent-expression* is NULL. Sybase IQ then applies the following algorithm:

- **REGR_R2** calculates VAR_POP(x) and returns NULL if VAR_POP(x) = 0; otherwise, it calculates VAR_POP(y) and returns the value 1 if VAR_POP(y) = 0.
- If neither VAR_POP(x) or VAR_POP(y) is zero, the return value is POWER(CORR(y,x), 2)

where y represents the *dependent-expression* and x represents the *independent-expression*.

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions*.

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. **DISTINCT** is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example returns the value 0.19379959710325653:

```
SELECT REGR_R2( Salary, ( YEAR( NOW() ) - YEAR( BirthDate ) ) )FROM  
Employees ;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

REGR_SLOPE Function [Aggregate]

Computes the slope of the linear regression line, fitted to non-NULL pairs.

Syntax 1

```
REGR_SLOPE(dependent-expression, independent-expression)
```

Syntax 2

```
REGR_SLOPE(dependent-expression, independent-expression)
```

```
OVER (window-spec)
```

Parameters

Parameter	Description
<i>dependent-expression</i>	The variable that is affected by the independent variable.
<i>independent-expression</i>	The variable that influences the outcome.

Returns

DOUBLE

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then **REGR_SLOPE** returns NULL.

REGR_SLOPE is applied to the set of (*dependent-expression* and *independent-expression*) pairs after eliminating all pairs for which either *dependent-expression* or *independent-expression* is NULL. The function is computed simultaneously during a single pass through the data. After eliminating NULL values, the following computation is made, where *y* represents the *dependent-expression* and *x* represents the *independent-expression*:

```
COVAR_POP(x, y) / VAR_POP(y)
```

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions*.

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example returns the value 935.3429749445614:

```
SELECT REGR_SLOPE( Salary, ( YEAR( NOW() ) -  
YEAR( BirthDate ) ) )FROM Employees;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

REGR_SXX Function [Aggregate]

Computes the slope of the linear regression line, fitted to non-NULL pairs.

Syntax 1

```
REGR_SXX(dependent-expression, independent-expression)
```

Syntax 2

```
REGR_SXX(dependent-expression, independent-expression)  
OVER (window-spec)
```

Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Returns

DOUBLE

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then **REGR_SXX** returns NULL.

The function is applied to the set of (*dependent-expression* and *independent-expression*) pairs after eliminating all pairs for which either *dependent-expression* or *independent-expression* is NULL. The function is computed simultaneously during a single pass through the data. After eliminating NULL values, the following computation is made, where y represents the *dependent-expression* and x represents the *independent-expression*:

```
REGR_COUNT(y, x) * VAR_POP(x)
```

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions.*

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example returns the value 5916.4800000000105:

```
SELECT REGR_SXX( Salary, ( YEAR( NOW() ) - YEAR( BirthDate ) ) )FROM
Employees;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

REGR_SXY Function [Aggregate]

Returns the sum of products of the dependent and independent variables. Use REGR_SXY to evaluate the statistical validity of a regression model.

Syntax 1

```
REGR_SXY(dependent-expression, independent-expression)
```

Syntax 2

```
REGR_SXY(dependent-expression, independent-expression)
```

```
OVER (window-spec)
```

Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Returns

DOUBLE

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then it returns NULL.

The function is applied to the set of (dependent-expression and *independent-expression*) pairs after eliminating all pairs for which either dependent-expression or *independent-expression* is NULL. The function is computed simultaneously during a single pass through the data. After eliminating NULL values, the following computation is made, where y represents the dependent-expression and x represents the *independent-expression*:

```
REGR_COUNT(x, y) * COVAR_POP(x,
y)
```

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions*.

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example returns the value 5533938.004400015.

```
SELECT REGR_SXY( Salary, ( YEAR( NOW() ) - YEAR( BirthDate ) ) )FROM
Employees;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

REGR_SYY Function [Aggregate]

Returns values that can evaluate the statistical validity of a regression model.

Syntax 1

```
REGR_SYY(dependent-expression, independent-expression)
```

Syntax 2

```
REGR_SYY(dependent-expression, independent-expression)
```

OVER (*window-spec*)

Parameters

Parameter	Description
dependent-expression	The variable that is affected by the independent variable.
independent-expression	The variable that influences the outcome.

Returns

DOUBLE

Usage

This function converts its arguments to DOUBLE, performs the computation in double-precision floating-point, and returns a DOUBLE as the result. If applied to an empty set, then **REGR_SYY** returns NULL.

The function is applied to the set of (*dependent-expression* and independent-expression) pairs after eliminating all pairs for which either *dependent-expression* or independent-expression is NULL. The function is computed simultaneously during a single pass through the data. After eliminating NULL values, the following computation is then made, where y represents the *dependent-expression* and x represents the independent-expression:

```
REGR_COUNT(x, y) * VAR_POP(y)
```

See *SQL Anywhere Server - SQL Usage > Querying and modifying data > OLAP support > Window functions in SQL Anywhere > Row numbering functions > Mathematical formulas for the aggregate functions.*

Note: ROLLUP and CUBE are not supported in the **GROUP BY** clause with Syntax 1. DISTINCT is not supported.

Syntax 2 represents usage as a window function in a **SELECT** statement. As such, you can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

Example

The following example returns the value 26, 708, 672,843.3002:

```
SELECT REGR_SYY( Salary, ( YEAR( NOW() ) - YEAR( BirthDate ) ) )FROM
Employees;
```

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T612.
- Sybase—Compatible with SQL Anywhere.

See also

- *Windowing Aggregate Function Usage* on page 103

REMAINDER Function [Numeric]

Returns the remainder when one whole number is divided by another.

Syntax

```
REMAINDER ( dividend, divisor )
```

Parameters

Parameter	Description
dividend	The dividend, or numerator of the division.
divisor	The divisor, or denominator of the division.

Returns

INTEGER

NUMERIC

Example

The following statement returns the value 2:

```
SELECT REMAINDER( 5, 3 ) FROM iq_dummy
```

Usage

REMAINDER is the same as the **MOD** function.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported in Adaptive Server Enterprise. The % (modulo) operator and the division operator can be used to produce a remainder.

See also

- *MOD Function [Numeric]* on page 235

REPEAT Function [String]

Concatenates a string a specified number of times.

Syntax

```
REPEAT ( string-expression, integer-expression )
```


Parameters

Parameter	Description
string-expression	The string to be repeated.
integer-expression	The number of times the string is to be repeated. If <i>integer-expression</i> is negative, an empty string is returned.

Returns

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **REPEAT** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **REPEAT** to the correct data type and size.

Example

The following statement returns the value “repeatrepeatrepeat.”

```
SELECT REPEAT( 'repeat', 3 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported in Adaptive Server Enterprise, but **REPLICATE** provides the same capabilities.

See also

- *REPLACE Function [String]* on page 279
- *REPLICATE Function [String]* on page 282

REPLACE Function [String]

Replaces all occurrences of a substring with another substring.

Syntax

```
REPLACE ( original-string, search-string, replace-string )
```

Parameters

If any argument is NULL, the function returns NULL.

Parameter	Description
original-string	The string to be searched. This string can be any length.
search-string	The string to be searched for and replaced with <i>replace-string</i> . This string is limited to 255 bytes. If <i>search-string</i> is an empty string, the original string is returned unchanged.
replace-string	The replacement string, which replaces <i>search-string</i> . This can be any length. If <i>replace-string</i> is an empty string, all occurrences of <i>search-string</i> are deleted.

Returns

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **REPLACE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **REPLACE** to the correct data type and size.

Examples

The following statement returns the value “xx.def.xx.ghi:”

```
SELECT REPLACE( 'abc.def.abc.ghi', 'abc', 'xx' ) FROM iq_dummy
```

The following statement generates a result set containing **ALTER PROCEDURE** statements which, when executed, repair stored procedures that reference a table that has been renamed. (To be useful, the table name must be unique.)

```
SELECT REPLACE(
    replace(proc_defn, 'OldTableName', 'NewTableName'),
    'create procedure',
    'alter procedure')
FROM SYS.SYSPROCEDURE
WHERE proc_defn LIKE '%OldTableName%'
```

Use a separator other than the comma for the **LIST** function:

```
SELECT REPLACE( list( table_id ), ',', '--')
FROM SYS.ISYSTAB
WHERE table_id <= 5
```

Usage

The result data type of a **REPLACE** function is a LONG VARCHAR. If you use **REPLACE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license, or use **CAST** and set **REPLACE** to the correct data type and size.

There are two ways to work around this issue:

- Declare a local temporary table, then perform an **INSERT**:

```
DECLARE local temporary table #mytable
  (name_column char(10)) on commit preserve rows;
INSERT INTO #mytable SELECT REPLACE(name,'0','1') FROM
dummy_table01;
```

- Use **CAST**:

```
SELECT CAST(replace(name, '0', '1') AS Char(10)) into #mytable
from dummy_table01;
```

If you need to control the width of the resulting column when *replace-string* is wider than *search-string*, use the **CAST** function. For example,

```
CREATE TABLE aa(a CHAR(5));
INSERT INTO aa VALUES('CCCCC');
COMMIT;
SELECT a, CAST(REPLACE(a,'C','ZZ') AS CHAR(5)) FROM aa;
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *CONVERT Function [Data Type Conversion]* on page 144
- *HOURS Function [Date and Time]* on page 201
- *MINUTES Function [Date and Time]* on page 234
- *MONTHS Function [Date and Time]* on page 237
- *SECOND Function [Date and Time]* on page 290
- *WEEKS Function [Date and Time]* on page 340
- *YEAR Function [Date and Time]* on page 345
- *YEARS Function [Date and Time]* on page 346
- *LCASE Function [String]* on page 217
- *LEFT Function [String]* on page 220
- *LOWER Function [String]* on page 228
- *REVERSE Function [String]* on page 283
- *RIGHT Function [String]* on page 284
- *UCASE Function [String]* on page 331
- *UPPER Function [String]* on page 332
- *REPEAT Function [String]* on page 278
- *REPLICATE Function [String]* on page 282

REPLICATE Function [String]

Concatenates a string a specified number of times.

Syntax

```
REPLICATE ( string-expression, integer-expression )
```

Parameters

Parameter	Description
string-expression	The string to be repeated.
integer-expression	The number of times the string is to be repeated.

Returns

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **REPLICATE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **REPLICATE** to the correct data type and size.

Example

The following statement returns the value “repeatrepeatrepeat:”

```
SELECT REPLICATE( 'repeat', 3 ) FROM iq_dummy
```

Usage

REPLICATE is the same as the **REPEAT** function.

Note: The result data type of a **REPLICATE** function is a LONG VARCHAR. If you use **REPLICATE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **REPLICATE** to the correct data type and size.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *REPEAT Function [String]* on page 278
- *REPLACE Function [String]* on page 279

REVERSE Function [String]

Takes one argument as an input of type `BINARY` or `STRING` and returns the specified string with characters listed in reverse order.

Syntax

```
REVERSE ( expression | uchar_expr )
```

Parameters

Parameter	Description
expression	A character or binary-type column name, variable, or constant expression of <code>CHAR</code> , <code>VARCHAR</code> , <code>NCHAR</code> , <code>NVARCHAR</code> , <code>BINARY</code> , or <code>VARBINARY</code> type.

Returns

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **REVERSE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **REVERSE** to the correct data type and size.

Example 1

```
select reverse("abcd" )
-----
dcba
```

Example 2

```
select reverse(0x12345000)
-----
0x00503412
```

Usage

- **REVERSE**, a string function, returns the reverse of expression.
- If expression is `NULL`, reverse returns `NULL`.
- Surrogate pairs are treated as indivisible and are not reversed.

Permissions

Any user can execute **REVERSE**.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.

See also

- *String Operators* on page 25
- *LCASE Function [String]* on page 217
- *LEFT Function [String]* on page 220
- *LOWER Function [String]* on page 228
- *REPLACE Function [String]* on page 279
- *RIGHT Function [String]* on page 284
- *UCASE Function [String]* on page 331
- *UPPER Function [String]* on page 332

RIGHT Function [String]

Returns the rightmost characters of a string.

Syntax

RIGHT (*string-expression*, *numeric-expression*)

Parameters

Parameter	Description
string-expression	The string to be left-truncated.
numeric-expression	The number of characters at the end of the string to return.

Returns

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **RIGHT** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **RIGHT** to the correct data type and size.

Example

The following statement returns the value “olate:”

```
SELECT RIGHT( 'chocolate', 5 ) FROM iq_dummy
```

Usage

If the string contains multibyte characters, and the proper collation is being used, the number of bytes returned might be greater than the specified number of characters.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *LCASE Function [String]* on page 217
- *LEFT Function [String]* on page 220
- *LOWER Function [String]* on page 228
- *REPLACE Function [String]* on page 279
- *REVERSE Function [String]* on page 283
- *UCASE Function [String]* on page 331
- *UPPER Function [String]* on page 332

ROUND Function [Numeric]

Rounds the *numeric-expression* to the specified *integer-expression* number of places after the decimal point.

Syntax

```
ROUND ( numeric-expression, integer-expression )
```

Parameters

Parameter	Description
numeric-expression	The number, passed to the function, to be rounded.
integer-expression	A positive integer specifies the number of significant digits to the right of the decimal point at which to round. A negative expression specifies the number of significant digits to the left of the decimal point at which to round.

Returns

NUMERIC

Examples

The following statement returns the value 123.200:

SQL Functions

```
SELECT ROUND( 123.234, 1 ) FROM iq_dummy
```

Additional results of the **ROUND** function are shown in the following table:

Value	ROUND (Value)
123.4567	round (a.n,4)
123.4570	round (a.n,3)
123.4600	round (a.n,2)
123.5000	round (a.n,1)
123.0000	round (a.n, 0)
120.0000	round (a.n,-1)
100.0000	round (a.n,-2)
0.0000	round(a.n,-3)

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *TRUNCNUM Function [Numeric]* on page 320

ROW_NUMBER Function [Analytical]

A ranking function that returns a unique row number for each row in a window partition, restarting the row numbering at the start of every window partition.

If no window partitions exist, the function numbers the rows in the result set from 1 to the cardinality of the table.

Syntax

```
ROW_NUMBER() OVER ([PARTITION BY window partition] ORDER BY window ordering)
```

Parameters

Parameter	Description
window partition	(Optional) One or more value expressions separated by commas indicating how you want to divide the set of result rows.

Parameter	Description
window ordering	Defines the expressions for sorting rows within window partitions, if specified, or within the result set if you did not specify a window partition.

Usage

The **ROW_NUMBER** function requires an **OVER (ORDER_BY)** window specification. The window partitioning clause in the **OVER (ORDER_BY)** clause is optional. The **OVER (ORDER_BY)** clause must not contain a window frame **ROWS/RANGE** specification.

Example

The following example returns salary data from the Employees table, partitions the result set by department ID, and orders the data according to employee start date. The **ROW_NUMBER** function assigns each row a row number, and restarts the row numbering for each window partition:

```
SELECT DepartmentID dID, StartDate, Salary,
ROW_NUMBER()OVER(PARTITION BY dID ORDER BY StartDate) FROM Employees
ORDER BY 1,2;
```

The returned result set is:

dID	StartDate	Salary	Row_number()
=====	=====	=====	=====
100	1986-10-14	42,998.000	1
100	1987-07-23	39,875.500	2
100	1988-03-23	37,400.000	3
100	1989-04-20	42,500.000	4
100	1990-01-15	42,100.000	5
200	1985-02-03	38,500.000	1
200	1987-02-19	39,300.000	2
200	1988-11-22	39,800.000	3
200	1989-06-01	34,892.000	4
200	1990-05-13	33,890.000	5
200	1990-07-11	37,803.000	6

Standards and Compatibility

- SQL—ISO/ANSI SQL compliant. SQL/OLAP feature T611.

ROWID Function [Miscellaneous]

Returns the internal row ID value for each row of the table.

Syntax

```
ROWID ( table-name ) ..FROM table-name
```

Parameters

Parameter	Description
table-name	The name of the table. Specify the name of the table within the parentheses with either no quotes or with double quotes.

Returns

UNSIGNED BIGINT

Examples

The following statement returns the row ID values 1 through 10:

```
SELECT ROWID( "PRODUCTS" ) FROM PRODUCTS
```

rowid(Products)
1
2
3
.
.
.
10

The following statement returns the product ID and row ID value of all rows with a product ID value less than 400:

```
SELECT PRODUCTS.ID, ROWID ( PRODUCTS )
FROM PRODUCTS
WHERE PRODUCTS.ID < 400
```

ID	rowid(Products)
300	1
301	2
302	3

The following statement deletes all rows with row ID values greater than 50:

```
DELETE FROM PRODUCTS
WHERE ROWID ( PRODUCTS ) > 50
```

Usage

You can use the **ROWID** function with other clauses to manipulate specific rows of the table.

You must specify the **FROM** *table-name* clause.

A limitation of the **ROWID** function is that it cannot use a join index of that table, eliminating any performance benefits that would normally use that join index.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

RTRIM Function [String]

Returns a string with trailing blanks removed.

Syntax

```
RTRIM ( string-expression )
```

Parameters

Parameter	Description
string-expression	The string to be trimmed.

Returns

VARCHAR

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **RTRIM** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **RTRIM** to the correct data type and size.

Example

The following statement returns the string “Test Message” with all trailing blanks removed.

```
SELECT RTRIM( 'Test Message      ' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *LTRIM Function [String]* on page 229

SECOND Function [Date and Time]

Returns a number from 0 to 59 corresponding to the second component of the given date/time value.

Syntax

```
SECOND ( datetime-expression )
```

Parameters

Parameter	Description
<i>datetime-expression</i>	The date/time value.

Returns

SMALLINT

Example

The following statement returns the value 5:

```
SELECT SECOND( '1998-07-13 08:21:05' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *CONVERT Function [Data Type Conversion]* on page 144
- *HOURS Function [Date and Time]* on page 201
- *MINUTES Function [Date and Time]* on page 234
- *MONTHS Function [Date and Time]* on page 237
- *REPLACE Function [String]* on page 279
- *WEEKS Function [Date and Time]* on page 340
- *YEAR Function [Date and Time]* on page 345
- *YEARS Function [Date and Time]* on page 346

SECONDS Function [Date and Time]

Returns the number of seconds since an arbitrary starting date and time, the number of seconds between two times, or adds an integer amount of seconds to a time.

Syntax

```
SECONDS ( datetime-expression
| datetime-expression, datetime-expression
| datetime-expression, integer-expression )
```

Parameters

Parameter	Description
<i>datetime-expression</i>	A date and time.
<i>integer-expression</i>	The number of seconds to be added to the <i>datetime-expression</i> . If <i>integer-expression</i> is negative, the appropriate number of minutes are subtracted from the date/time value. If you supply an integer expression, the <i>datetime-expression</i> must be explicitly cast as a datetime data type.

Returns

INTEGER

TIMESTAMP

Examples

The following statement returns the value 3600:

```
SELECT ( SECONDS( '1998-07-13 06:07:12' ) -
SECONDS( '1998-07-13 05:07:12' ) ) FROM iq_dummy
```

The following statement returns the value 14400, to signify the difference between the two times:

```
SELECT SECONDS( '1999-07-13 06:07:12',
'1999-07-13 10:07:12' ) FROM iq_dummy
```

The following statement returns the datetime value 1999-05-12 21:05:12.000:

```
SELECT SECONDS( CAST( '1999-05-12 21:05:07'
AS TIMESTAMP ), 5) FROM iq_dummy
```

Usage

The second syntax returns the number of whole seconds from the first date/time to the second date/time. The number might be negative.

Standards and compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

SIGN Function [Numeric]

Returns the sign of a number.

Syntax

```
SIGN ( numeric-expression )
```

Parameters

Parameter	Description
numeric-expression	The number for which the sign is to be returned.

Returns

SMALLINT

Example

The following statement returns the value -1:

```
SELECT SIGN( -550 ) FROM iq_dummy
```

Usage

For negative numbers, **SIGN** returns -1.

For zero, **SIGN** returns 0.

For positive numbers, **SIGN** returns 1.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

SIMILAR Function [String]

Returns an integer between 0 and 100 representing the similarity between two strings.

Syntax

```
SIMILAR ( string-expression1, string-expression2 )
```

Parameters

Parameter	Description
string-expression1	The first string to be compared.
string-expression2	The second string to be compared.

Returns

SMALLINT

Example

The following statement returns the value 75:

```
SELECT SIMILAR( 'toast', 'coast' ) FROM iq_dummy
```

This signifies that the two values are 75% similar.

Usage

The function returns an integer between 0 and 100 representing the similarity between the two strings. The result can be interpreted as the percentage of characters matched between the two strings. A value of 100 indicates that the two strings are identical.

This function can be used to correct a list of names (such as customers). Some customers might have been added to the list more than once with slightly different names. Join the table to itself and produce a report of all similarities greater than 90 percent but less than 100 percent.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

SIN Function [Numeric]

Returns the sine of a number, expressed in radians.

Syntax

```
SIN ( numeric-expression )
```

Parameters

Parameter	Description
numeric-expression	The angle, in radians.

Returns

DOUBLE

Example

The following statement returns the value 0.496880:

```
SELECT SIN( 0.52 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *COS Function [Numeric]* on page 149
- *ATAN2 Function [Numeric]* on page 128
- *ATAN Function [Numeric]* on page 127
- *ASIN Function [Numeric]* on page 126
- *ACOS Function [Numeric]* on page 124
- *COT Function [Numeric]* on page 150
- *TAN Function [Numeric]* on page 318

SORTKEY Function [String]

Generates values that can be used to sort character strings based on alternate collation rules.

Syntax

```
SORTKEY ( string-expression  
[ , { collation-id  
| collation-name [(collation-tailoring-string)] } ]  
)
```


Parameters

Parameter	Description
string-expression	<p>The string expression must contain characters that are encoded in the character set of the database and must be <code>STRING</code> data type.</p> <p>If string-expression is <code>NULL</code>, the <code>SORTKEY</code> function returns a null value. An empty string has a different sort-order value than a null string from a database column.</p> <p>There is no limit on the length of the input string that the <code>SORTKEY</code> function can handle. The result of <code>SORTKEY</code> is always limited to 1024 bytes and is <code>VARBINARY</code> data type. If the actual results exceed 1024 bytes, the result contains only the first 1024 bytes.</p>
collation-name	<p>A string or character variable that specifies the name of the sort order to use. You can also specify the alias <code>char_collation</code>, or, equivalently, <code>db_collation</code>, to generate sort-keys as used by the <code>CHAR</code> collation in use by the database.</p> <p>Similarly, you can specify the alias <code>NCHAR_COLLATION</code> to generate sort-keys as used by the <code>NCHAR</code> collation in use by the database. However, Sybase IQ does not support <code>NCHAR_COLLATION</code> for Sybase IQ-specific objects. <code>NCHAR_COLLATION</code> is supported for SQL Anywhere objects on a Sybase IQ server.</p>
collation-id	<p>A variable, integer constant, or string that specifies the ID number of the sort order to use. This parameter applies only to Adaptive Server Enterprise collations, which can be referred to by their corresponding collation ID.</p>

Parameter	Description
<p>collation-tailoring-string</p>	<p>(Optional) Specify collation tailoring options (<i>collation-tailoring-string</i>) for additional control over sorting and comparison of characters. These options take the form of keyword=value pairs assembled in parentheses, following the collation name. For example,</p> <pre>'UCA(locale=es;case=Lower-First;accent=respect)'</pre> <p>The syntax for specifying these options is identical to the COLLATION clause of the CREATE DATABASE statement. .</p> <p>See <i>Reference: Statements and Options > SQL Statements > CREATE DATABASE Statement.</i></p> <hr/> <p>Note: All of the collation tailoring options are supported for SQL Anywhere databases, when specifying the Unicode Collation Algorithm (UCA) collation. For all other collations, only case sensitivity tailoring is supported.</p>

Returns
 BINARY

Example

The following statement queries the Employees table and returns the FirstName and Surname of all employees, sorted by the sort-key values for the Surname column using the dict collation (Latin-1, English, French, German dictionary):

```
SELECT Surname, GivenName FROM Employees ORDER BY SORTKEY( Surname, 'dict' );
```

Usage

The **SORTKEY** function generates values that can be used to order results based on predefined sort order behavior. This allows you to work with character sort order behaviors that may not be available from the database collation. The returned value is a binary value that contains coded sort order information for the input string that is retained from the **SORTKEY** function.

For example, you can store the values returned by the **SORTKEY** function in a column with the source character string. The following **SELECT** statement retrieves data from table T1 in the sorted order of c1 according to the Thai dictionary:

```
SELECT rid, c1 from T1 ORDER BY SORTKEY(c1)
```

You instead store the value returned by **SORTKEY** in a column with the source character string. To retrieve the character data in the required order, the **SELECT** statement needs to include only an **ORDER BY** clause on the column that contains the results of running the **SORTKEY** function:

```
UPDATE T1 SET shadowc1=SORTKEY(c1) FROM T1;
SELECT rid, c1 FROM T1 ORDER BY shadowc1
```

The **SORTKEY** function guarantees that the values it returns for a given set of sort order criteria work for the binary comparisons that are performed on **VARBINARY** data types.

Generating sort-keys for queries can be expensive. As an alternative for frequently requested sort-keys, consider creating a computed column to hold the sort-key values, and then referencing that column in the **ORDER BY** clause of the query.

If you do not specify a collation name or collation ID, the default is Default Unicode multilingual.

Valid collations are as follows:

- To see collations that are supported by Sybase IQ, listed by label, execute `iqinit -l`.
- The Adaptive Server Enterprise collations are listed in the table below.

Description	Collation name	Collation ID
Default Unicode multilingual	default	0
CP 850 Alternative: no accent	altnoacc	39
CP 850 Alternative: lowercase first	altdict	45
CP 850 Western European: no case, preference	altnocsp	46
CP 850 Scandinavian dictionary	scandict	47
CP 850 Scandinavian: no case, preference	scannocp	48
GB Pinyin	gbpinyin	n/a
Binary sort	binary	50
Latin-1 English, French, German dictionary	dict	51
Latin-1 English, French, German no case	nocase	52
Latin-1 English, French, German no case, preference	nocasep	53
Latin-1 English, French, German no accent	noaccent	54
Latin-1 Spanish dictionary	espdict	55
Latin-1 Spanish no case	espnocs	56
Latin-1 Spanish no accent	espnoac	57

Description	Collation name	Collation ID
ISO 8859-5 Russian dictionary	rusdict	58
ISO 8859-5 Russian no case	rusnoc	59
ISO 8859-5 Cyrillic dictionary	cyrdict	63
ISO 8859-5 Cyrillic no case	cyrnoc	64
ISO 8859-7 Greek dictionary	elldict	65
ISO 8859-2 Hungarian dictionary	hundict	69
ISO 8859-2 Hungarian no accents	hunnoac	70
ISO 8859-2 Hungarian no case	hunnocs	71
ISO 8859-5 Turkish dictionary	turdict	72
ISO 8859-5 Turkish no accents	turnoac	73
ISO 8859-5 Turkish no case	turnocs	74
CP 874 (TIS 620) Royal Thai dictionary	thaidict	1
ISO 14651 ordering standard	14651	22
Shift-JIS binary order	sjisbin	179
Unicode UTF-8 binary sort	utf8bin	24
EUC JIS binary order	eucjisbn	192
GB2312 binary order	gb2312bn	137
CP932 MS binary order	cp932bin	129
Big5 binary order	big5bin	194
EUC KSC binary order	euckscbn	161

With respect to collation tailoring, full sensitivity is generally the intent when creating sortkeys, so when you specify a non-UCA collation, the default tailoring applied is equivalent to case=Respect. For example, the following two statements are equivalent:

```
SELECT SORTKEY( 'abc', '1252LATIN1' );
SELECT SORTKEY( 'abc', '1252LATIN1(case=Respect)' );
```

When specifying a non-UCA collation, by default, collation tailorings are accent and case-sensitive. However, for non-UCA collations, you can override only the case sensitivity using a collation tailoring. For example:

```
SELECT SORTKEY( 'abc', '1252LATIN1(case=LowerFirst)' );
```

If the database was created without specifying tailoring options, the following two clauses may generate different sort orders, even if the database collation name is specified for the **SORTKEY** function:

```
ORDER BY string-expression
```

```
ORDER BY SORTKEY( string-expression, database-collation-name )
```

Different sort orders may be generated, because the default tailoring settings used for database creation and for the **SORTKEY** function are different. To get the same behavior from **SORTKEY** as for the database collation, either provide a tailoring syntax for *collation-tailoring-string* that matches the settings for the database collation, or specify `db_collation` for collation-name. For example:

```
SORTKEY( expression, 'db_collation' )
```

Note: Sort-key values created using a version of Sybase IQ earlier than 15.0 do not contain the same values created using version 15.0 and later. This may be a problem for your applications if your pre-15.0 database has sort-key values stored within it, especially if sort-key value comparison is required by your application. Regenerate any sort-key values in your database that were generated using a version of Sybase IQ earlier than 15.0.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

SOUNDEX Function [String]

Returns a number representing the sound of a string.

Syntax

```
SOUNDEX ( string-expression )
```

Parameters

Parameters	Description
string-expression	The string.

Returns

SMALLINT

Example

The following statement returns two numbers, representing the sound of each name. The **SOUNDEX** value for each argument is 3827.

```
SELECT SOUNDEX( 'Smith' ), SOUNDEX( 'Smythe' ) FROM iq_dummy
```

SOUNDEX ('Smith') is equal to **SOUNDEX** ('Smythe').

Usage

The **SOUNDEX** function value for a string is based on the first letter and the next three consonants other than H, Y, and W. Doubled letters are counted as one letter. For example:

```
SOUNDEX( 'apples' ) FROM iq_dummy
```

is based on the letters A, P, L, and S.

Multibyte characters are ignored by the **SOUNDEX** function.

Although it is not perfect, **SOUNDEX** normally returns the same number for words that sound similar and that start with the same letter.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise, except that Adaptive Server Enterprise returns a CHAR(4) result and Sybase IQ returns an integer.

See also

- *DIFFERENCE Function [String]* on page 181

SPACE Function [String]

Returns a specified number of spaces.

Syntax

```
SPACE ( integer-expression )
```

Parameters

Parameter	Description
integer-expression	The number of spaces to return.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **SPACE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **SPACE** to the correct data type and size.

Example

The following statement returns a string containing 10 spaces:

```
SELECT SPACE( 10 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

SQLFLAGGER Function [Miscellaneous]

Returns the conformity of a given SQL statement to a specified standard.

Syntax

```
SQLFLAGGER ( sql-standard-string, sql-statement-string )
```

Parameters

Parameter	Description
sql-standard-string	<p>The standard level against which to test compliance. Possible values are the same as for the <code>SQL_FLAGGER_ERROR_LEVEL</code> database option:</p> <ul style="list-style-type: none"> • SQL:2003/Core Test for conformance to core SQL/2003 syntax. • SQL:2003/Package Test for conformance to full SQL/2003 syntax. • SQL:1999/Core Test for conformance to core SQL/1999 syntax. • SQL:1999/Package Test for conformance to full SQL/1999 syntax. • SQL:1992/Entry Test for conformance to entry-level SQL/1992 syntax. • SQL:1992/Intermediate Test for conformance to intermediate-level SQL/1992 syntax. • SQL:1992/Full Test for conformance to full-SQL/1992 syntax.
sql-statement-string	The SQL statement to check for conformance.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **SQLFLAGGER** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **SQLFLAGGER** to the correct data type and size.

Examples

The following statement shows an example of the message that is returned when a disallowed extension is found:

```
SELECT SQLFLAGGER( 'SQL:2003/Package', 'SELECT top 1 dummy_col FROM sys.dummy ORDER BY dummy_col' );
```

This statement returns the message '0AW03 Disallowed language extension detected in syntax near 'top' on line 1'.

The following statement returns '00000' because it contains no disallowed extensions:

```
SELECT SQLFLAGGER( 'SQL:2003/Package', 'SELECT dummy_col FROM sys.dummy' );
```

Usage

You can also use the `iqsqlpp` SQL Preprocessor Utility to flag any Embedded SQL that is not part of a specified set of SQL92. See *Utility Guide > iqsqlpp SQL Preprocessor Utility*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

SQRT Function [Numeric]

Returns the square root of a number.

Syntax

```
SQRT ( numeric-expression )
```

Parameters

Parameter	Description
numeric-expression	The number for which the square root is to be calculated.

Returns

DOUBLE

Example

The following statement returns the value 3:

```
SELECT SQRT( 9 ) FROM iq_dummy
```


Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

SQUARE Function [Numeric]

Returns the square of the specified expression as a float.

Syntax

```
SQUARE ( numeric-expression )
```

Parameters

Parameter	Description
expression	Is a column, variable, or expression with a data type that is either exact numeric, approximate numeric, money, or any type that can be implicitly converted to one of these types. For other data types, the SQUARE function generates an error. The return value is of DOUBLE data type.

Usage

SQUARE function takes one argument. For example, **SQUARE** (12.01) returns 144.240100.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

STDDEV Function [Aggregate]

Returns the standard deviation of a set of numbers.

Syntax

```
STDDEV ( [ ALL ] expression )
```

Parameters

Parameter	Description
expression	Any numeric data type (FLOAT, REAL, or DOUBLE precision) expression.

Returns

DOUBLE

Examples

Given this data:

```
SELECT Salary FROM Employees WHERE DepartmentID = 300
```

Salary
51432.000
57090.000
42300.000
43700.00
36500.000
138948.000
31200.000
58930.00
75400.00

The following statement returns the value 32617.8446712838471:

```
SELECT STDDEV ( Salary ) FROM Employees  
WHERE DepartmentID = 300
```

Given this data:

```
SELECT UnitPrice FROM Products WHERE Name = 'Tee Shirt'
```

Name	UnitPrice
Tee Shirt	9.00
Tee Shirt	14.00
Tee Shirt	14.00

The following statement returns the value 2.88675134594813049:

```
SELECT STDDEV ( UnitPrice ) FROM Products  
WHERE Name = 'Tee Shirt'
```

Usage

The formula used to calculate **STDDEV** is:

$$stddev = \sqrt{variance}$$

STDDEV returns a result of data type `DOUBLE` precision floating-point. If applied to the empty set, the result is `NULL`, which returns `NULL` for a one-element input set.

STDDEV does not support the keyword `DISTINCT`. A syntax error is returned if you use `DISTINCT` with **STDDEV**.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103
- *STDDEV_SAMP Function [Aggregate]* on page 306
- *VARIANCE Function [Aggregate]* on page 338

STDDEV_POP Function [Aggregate]

Computes the standard deviation of a population consisting of a numeric-expression, as a `DOUBLE`.

Syntax

```
STDDEV_POP ( [ ALL ] expression )
```

Parameters

Parameter	Description
expression	The expression (commonly a column name) whose population-based standard deviation is calculated over a set of rows.

Returns

`DOUBLE`

Examples

The following statement lists the average and variance in the number of items per order in different time periods:

```
SELECT year( ship_date ) AS Year, quarter( ship_date )
      AS Quarter, AVG( quantity ) AS Average,
      STDDEV_POP ( quantity ) AS Variance
FROM SalesOrderItems GROUP BY Year, Quarter
ORDER BY Year, Quarter;
```

Year	Quarter	Average	Variance
2000	1	25.775148	14.2794
2000	2	27.050847	15.0270
...

Usage

Computes the population standard deviation of the provided *value expression* evaluated for each row of the group or partition (if DISTINCT was specified, then each row that remains after duplicates have been eliminated), defined as the square root of the population variance.

$$\sqrt{\frac{\sum (x_i - \bar{x})^2}{n}}$$

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103

STDDEV_SAMP Function [Aggregate]

Computes the standard deviation of a sample consisting of a numeric-expression, as a DOUBLE.

Syntax

```
STDDEV_SAMP ( [ ALL ] expression )
```

Parameters

Parameter	Description
expression	The expression (commonly a column name) whose sample-based standard deviation is calculated over a set of rows.

Returns

DOUBLE

Examples

The following statement lists the average and variance in the number of items per order in different time periods:

```
SELECT year( ship_date ) AS Year, quarter( ship_date )
      AS Quarter, AVG( quantity ) AS Average,
      STDDEV_SAMP( quantity ) AS Variance
FROM SalesOrderItems GROUP BY Year, Quarter
      ORDER BY Year, Quarter;
```

Year	Quarter	Average	Variance
2000	1	25.775148	14.3218
2000	2	27.050847	15.0696
...

Usage

Note: `STDDEV_SAMP` is an alias for `STDDEV`.

Computes the sample standard deviation of the provided *value expression* evaluated for each row of the group or partition (if `DISTINCT` was specified, then each row that remains after duplicates have been eliminated), defined as the square root of the sample variance.

`NULL` returns `NULL` for a one-element input set.

Standard deviations are computed according to the following formula, which assumes a normal distribution:

$$\sqrt{\frac{\sum (x_i - \bar{x})^2}{(n-1)}}$$

Standards and compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103
- *STDDEV Function [Aggregate]* on page 303
- *VARIANCE Function [Aggregate]* on page 338

STR Function [String]

Returns the string equivalent of a number.

Syntax

```
STR ( numeric-expression [ , length [ , decimal ] ] )
```

Parameters

Parameter	Description
numeric-expression	Any approximate numeric (FLOAT, REAL, or DOUBLE precision) expression.
length	The number of characters to be returned (including the decimal point, all digits to the right and left of the decimal point, the sign, if any, and blanks). The default is 10 and the maximum length is 255.
decimal	The number of digits to the right of the decimal point to be returned. The default is 0.

Returns

VARCHAR

Examples

The following statement returns a string of six spaces followed by 1234, for a total of ten characters:

```
SELECT STR( 1234.56 ) FROM iq_dummy
```

The following statement returns the result 1234.5:

```
SELECT STR( 1234.56, 6, 1 ) FROM iq_dummy
```

Usage

If the integer portion of the number cannot fit in the length specified, then the result is NULL. For example, the following statement returns NULL:

```
SELECT STR( 1234.56, 3 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

STR_REPLACE Function [String]

Takes three arguments as input of type BINARY or STRING and replaces any instances of the second string expression (*string_expr2*) that occur within the first string expression (*string_expr1*) with a third expression (*string_expr3*).

Syntax

```
REPLACE ( string_expr1, string_expr2, string_expr3 )
```

*Parameters***Table 94. Parameters**

Parameter	Description
string_expr1	The source string, or the string expression to be searched, expressed as CHAR, VARCHAR, UNICHAR, UNIVARCHAR, VARBINARY, or BINARY data type.
string_expr2	The pattern string, or the string expression to find within the first expression (<i>string_expr1</i>) and is expressed as CHAR, VARCHAR, UNICHAR, UNIVARCHAR, VARBINARY, or BINARY data type.
string_expr3	The replacement string expression, expressed as CHAR, VARCHAR, UNICHAR, UNIVARCHAR, VARBINARY, or BINARY data type.

Example 1

Replaces the string *def* within the string *cdefghi* with *yyy*.

```
select replace("cdefghi", "def", "yyy")
-----
cyyyghi
(1 row(s) affected)
```

Example 2

Replaces all spaces with “toyota”

```
select str_replace("chevy, ford, mercedes", " ", "toyota")
-----
chevy,toyotaford,toyotamercedes
(1 row(s) affected)
```

Example 3

Accepts NULL in the third parameter and treats it as an attempt to replace *string_expr2* with NULL, effectively turning STR_REPLACE into a “string cut” operation. Returns “abcghijklm”:

```
select str_replace("abcdefghijklm", "def", NULL)
-----
abcghijklm
(1 row affected)
```

Usage

STR_REPLACE is an alias of **REPLACE** function.

SQL Functions

- Takes any data type as input and returns **STRING** or **BINARY**.
For example, an empty string passed as an argument (“”) is replaced with one space (“ ”) before further evaluation occurs. This is true for both **BINARY** and **STRING** types.
- All arguments can have a combination of **BINARY** and **STRING** data types.
- The result length may vary, depending upon what is known about the argument values when the expression is compiled. If all arguments are columns or host variables assigned to constants, Sybase IQ calculates the result length as:

```
result_length = ((s/p)*(r-p)+s)
WHERE
  s = length of source string
  p = length of pattern string
  r = length of replacement string
IF (r-p) <= 0, result length = s
```
- If Sybase IQ cannot calculate the result length because the argument values are unknown when the expression is compiled, the result length used is 255.
- **RESULT_LEN** never exceeds 32767.

Permissions

Any user can execute **STR_REPLACE**.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.

See also

- *BIT_LENGTH Function [String]* on page 131
- *BYTE_LENGTH Function [String]* on page 132
- *CHAR_LENGTH Function [String]* on page 137
- *COL_LENGTH Function [System]* on page 141
- *DATALENGTH Function [System]* on page 155
- *LEN Function [String]* on page 222
- *LENGTH Function [String]* on page 223
- *OBJECT_NAME Function [System]* on page 249
- *OCTET_LENGTH Function [String]* on page 250

STRING Function [String]

Concatenates one or more strings into one large string.

Syntax

```
STRING ( string-expression [ , ... ] )
```


Parameters

Parameter	Description
string-expression	A string. If only one argument is supplied, it is converted into a single expression. If more than one argument is supplied, they are concatenated into a single string. A NULL is treated as an empty string (").

Returns

LONG VARCHAR

LONG NVARCHAR

LONG BINARY

Note: The result data type is a LONG VARCHAR. If you use **STRING** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **STRING** to the correct data type and size.

Example

The following statement returns the value testing123:

```
SELECT STRING( 'testing', NULL, 123 )
FROM iq_dummy
```

Usage

Numeric or date parameters are converted to strings before concatenation. You can also use the **STRING** function to convert any single expression to a string by supplying that expression as the only parameter.

If all parameters are NULL, **STRING** returns NULL.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

STRTOUUID Function [String]

Converts a string value to a unique identifier (UUID or GUID) value.

Syntax

```
STRTOUUID ( string-expression )
```

Parameters

Parameter	Description
string-expression	A string in the format <code>xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx</code>

Returns

UNIQUEIDENTIFIER

Example

```
CREATE TABLE T (
  pk uniqueidentifier primary key,
  c1 int);
INSERT INTO T (pk, c1)
VALUES (STRTOUUID
('12345678-1234-5678-9012-123456789012'), 1);
```

Usage

Converts a string in the format `xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx` where *x* is a hexadecimal digit, to a unique identifier value. If the string is not a valid UUID string, NULL is returned.

You can use **STRTOUUID** to insert UUID values into a Adaptive Server EnterpriseSybase IQ database.

Standards and Compatibility

- SQL—Transact-SQL extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Binary Data Types* on page 78
- *NEWID Function [Miscellaneous]* on page 239
- *UUIDTOSTR Function [String]* on page 334
- *Character Data Types* on page 69
- *Binary Data Types* on page 689

STUFF Function [String]

Deletes a number of characters from one string and replaces them with another string.

Syntax

```
STUFF ( string-expression1, start, length, string-expression2 )
```

Parameters

Parameter	Description
string-expression1	The string to be modified by the STUFF function.
start	The character position at which to begin deleting characters. The first character in the string is position 1.
length	The number of characters to delete.
string-expression2	The string to be inserted. To delete a portion of a string using the STUFF function, use a replacement string of NULL

Returns

LONG NVARCHAR

Example

The following statement returns the value “chocolate pie”:

```
SELECT STUFF( 'chocolate cake', 11, 4, 'pie' )
FROM iq_dummy
```

Usage

To delete a portion of a string using **STUFF**, use a replacement string of NULL. To insert a string using **STUFF**, use a length of zero.

The **STUFF** function will return a NULL result in the following situations:

- Any of the first three parameters is a NULL value.
- Either the start or length parameter is a negative value.
- The start parameter is greater than the length of string-expression1.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

SUBSTRING Function [String]

Returns a substring of a string.

Syntax

```
{ SUBSTRING | SUBSTR } ( string-expression, start [ , length ] )
```

Parameters

Parameter	Description
string-expression	The string from which a substring is to be returned.
start	The start position of the substring to return, in characters. A negative starting position specifies a number of characters from the end of the string instead of the beginning. The first character in the string is at position 1.
length	The length of the substring to return, in characters. A positive <i>length</i> specifies that the substring ends <i>length</i> characters to the right of the starting position, while a negative <i>length</i> specifies that the substring ends <i>length</i> characters to the left of the starting position.

Returns

LONG VARCHAR

LONG NVARCHAR

LONG BINARY

Note: The result data type is a LONG VARCHAR. If you use **STRING** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **STRING** to the correct data type and size.

Examples

The following statement returns “back”:

```
SELECT SUBSTRING ( 'back yard', 1 , 4 )
FROM iq_dummy
```

The following statement returns yard:

```
SELECT SUBSTR ( 'back yard', -1 , -4 )
FROM iq_dummy
```

The following statement returns 0x2233:

```
SELECT SUBSTR ( 0x112233445566, 2, 2 )
FROM iq_dummy
```

Usage

If *length* is specified, the substring is restricted to that length. If no length is specified, the remainder of the string is returned, starting at the *start* position.

Both *start* and *length* can be negative. Using appropriate combinations of negative and positive numbers, you can get a substring from either the beginning or end of the string.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—**SUBSTR** is not supported by Adaptive Server Enterprise. Use **SUBSTRING** instead

See also

- *CHARINDEX Function [String]* on page 139

SUBSTRING64 Function [String]

The **SUBSTRING64** function returns a variable-length character string of the large object column or variable parameter.

Usage

SUBSTRING64 supports searching LONG VARCHAR and LONG BINARY columns and LONG VARCHAR and LONG BINARY variables of any size of data. Currently, a SQL variable can hold up to 2GB - 1 in length.

If you are licensed to use the Unstructured Data Analytics functionality, you can use this function with large object data.

See *Unstructured Data Analytics in Sybase IQ > Function Support*.

SUM Function [Aggregate]

Returns the total of the specified expression for each group of rows.

Syntax

```
SUM ( expression | DISTINCT column-name )
```

Parameters

Parameter	Description
expression	The object to be summed. This is commonly a column name.
DISTINCT column-name	Computes the sum of the unique values in <i>column-name</i> for each group of rows. This is of limited usefulness, but is included for completeness.

Returns

INTEGER

DOUBLE

NUMERIC

Example

The following statement returns the value 3749146.740:

```
SELECT SUM( salary )
FROM Employees
```

Usage

Rows where the specified expression is NULL are not included.

Returns NULL for a group containing no rows.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *AVG Function [Aggregate]* on page 129
- *COUNT Function [Aggregate]* on page 153
- *Windowing Aggregate Function Usage* on page 103

SUSER_ID Function [System]

Returns an integer user identification number.

Syntax

```
SUSER_ID ( [ user-name ] )
```

Parameters

Parameter	Description
user-name	The user name.

Returns

INT

Examples

The following statement returns the user identification number 1:

```
SELECT SUSER_ID ( 'DBA' ) FROM iq_dummy
```

The following statement returns the user identification number 0:

```
SELECT SUSER_ID ( 'SYS' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *SUSER_NAME Function [System]* on page 317
- *USER_ID Function [System]* on page 333
- *USER_NAME Function [System]* on page 333

SUSER_NAME Function [System]

Returns the user name.

Syntax

```
SUSER_NAME ( [ user-id ] )
```

Parameters

Parameter	Description
user-id	The user identification number.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **SUSER_NAME** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license or use **CAST** and set **SUSER_NAME** to the correct data type and size.

Examples

The following statement returns the value DBA:

```
SELECT SUSER_NAME ( 1 ) FROM iq_dummy
```

The following statement returns the value SYS:

```
SELECT SUSER_NAME ( 0 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Sybase IQAdaptive Server Enterprise function implemented for Sybase IQ. In Adaptive Server Enterprise, **SUSER_NAME** returns the server user name.

See also

- *SUSER_ID Function [System]* on page 316
- *USER_ID Function [System]* on page 333
- *USER_NAME Function [System]* on page 333

TAN Function [Numeric]

Returns the tangent of a number.

Syntax

```
TAN ( numeric-expression )
```

Parameters

Parameter	Description
numeric-expression	An angle, in radians.

Returns

DOUBLE

Example

Returns the value 0.572561:

```
SELECT TAN( 0.52 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *COS Function [Numeric]* on page 149
- *ATAN2 Function [Numeric]* on page 128
- *ATAN Function [Numeric]* on page 127
- *ASIN Function [Numeric]* on page 126
- *ACOS Function [Numeric]* on page 124
- *COT Function [Numeric]* on page 150
- *SIN Function [Numeric]* on page 293

TODAY Function [Date and time]

Returns the current date. This is the historical syntax for **CURRENT DATE**.

Syntax

```
TODAY ( * )
```

Returns

DATE

Example

The following statement returns the current day according to the system clock.

```
SELECT TODAY( * ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

TRIM Function [String]

Removes leading and trailing blanks from a string.

Syntax

```
TRIM ( string-expression )
```

Parameters

Parameter	Description
string-expression	The string to be trimmed.

Returns

VARCHAR

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **TRIM** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license, or use **CAST** and set **TRIM** to the correct data type and size.

Example

The following statement returns the value “chocolate” with no leading or trailing blanks.

```
SELECT TRIM( ' chocolate ' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise. Use **LTRIM** and **RTRIM** instead

TRUNCNUM Function [Numeric]

Truncates a number at a specified number of places after the decimal point.

Syntax

```
TRUNCNUM ( numeric-expression, integer-expression )
```

Parameters

Parameter	Description
numeric-expression	The number to be truncated.
integer-expression	A positive integer specifies the number of significant digits to the right of the decimal point at which to round. A negative expression specifies the number of significant digits to the left of the decimal point at which to round.

Returns

NUMERIC

Examples

The following statement returns the value 600:

```
SELECT TRUNCNUM( 655, -2 ) FROM iq_dummy
```

The following statement: returns the value 655.340:

```
SELECT TRUNCNUM( 655.348, 2 ) FROM iq_dummy
```

Usage

This function is the same as **TRUNCATE**, but does not cause keyword conflicts.

You can use combinations of **ROUND**, **FLOOR**, and **CEILING** to provide similar functionality.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *ROUND Function [Numeric]* on page 285

TS_ARMA_AR Function [Time Series]

Calculates the least-square estimates of parameters for an autoregressive moving average (ARMA) model, and returns the requested autoregressive estimate.

Syntax

```
TS_ARMA_AR (timeseries_expression, ar_count, ar_elem, method)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_ARMA_CONST Function [Time Series]

Calculates the least-square estimates of parameters for an autoregressive moving average (ARMA) model, and returns an estimated constant.

Syntax

```
TS_ARMA_CONST (timeseries_expression, method)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_ARMA_MA Function [Time Series]

Calculates the least-square estimates of parameters for an autoregressive moving average (ARMA) model, and returns the requested moving average estimate.

Syntax

```
TS_ARMA_MA (timeseries_expression, ma_count, ma_elem, method)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTOCORRELATION function [Time Series]

Calculates the sample autocorrelation function of a stationary time series.

Syntax

```
TS_AUTOCORRELATION (timeseries_expression, lagmax, lag_elem)  
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA Function [Time Series]

Determines parameters of a multiplicative seasonal autoregressive integrated moving average (ARIMA) model, and produces forecasts that incorporate the effects of outliers whose effects persist beyond the end of the series.

Syntax

```
TS_AUTO_ARIMA(<time_value>,<timeseries_expression> [, <max_lag> [,  
<critical > [, <epsilon> [, <criteria> [, <confidence> [, <model> [,  
<n_predictions>]]]]]]])  
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_OUTLIER Function [Time Series]

Like the **TS_AUTO_ARIMA** aggregate function, **TS_AUTO_ARIMA_OUTLIER** accepts an input time series and automatically determines the parameters of a multiplicative seasonal autoregressive integrated moving average (ARIMA) model.

Where **TS_AUTO_ARIMA** uses the ARIMA model to forecast values beyond the set of inputs, **TS_AUTO_ARIMA_OUTLIER** uses the ARIMA model to identify the elements of the input time series that are statistical outliers, and returns the outlier type of each one.

Syntax

```
TS_AUTO_ARIMA_OUTLIER(<time_value>,<timeseries_expression> [,  
<max_lag> [, <critical > [, <epsilon> [, <criteria> [, <confidence> [,  
<model> [, <delta>]]]]]]])  
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_AIC Function [Time Series]

Retrieves the Akaike's Information Criterion (AIC) output parameter produced by **TS_AUTO_ARIMA**.

Syntax

```
TS_AUTO_ARIMA_RESULT_AIC(auto_arima_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_AICC Function [Time Series]

Retrieves the corrected AIC (AICC) output parameter produced by **TS_AUTO_ARIMA**.

Syntax

```
TS_AUTO_ARIMA_RESULT_AICC(auto_arima_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_BIC Function [Time Series]

Retrieves the Bayesian Information Criterion (BIC) output parameter produced by **TS_AUTO_ARIMA**.

Syntax

```
TS_AUTO_ARIMA_RESULT_BIC(auto_arima_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_FORECAST_VALUE Function [Time Series]

Retrieves the forecasted values for the requested outlier free series produced by **TS_AUTO_ARIMA**.

Syntax

```
TS_AUTO_ARIMA_RESULT_FORECAST_VALUE(auto_arima_result,  
model_element_number)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_FORECAST_ERROR Function [Time Series]

Retrieves the forecasted standard error values for the original input series produced by **TS_AUTO_ARIMA**.

Syntax

```
TS_AUTO_ARIMA_RESULT_FORECAST_ERROR(auto_arima_result,
forecast_element_number)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_MODEL_D Function [Time Series]

Retrieves the *d* value produced by **TS_AUTO_ARIMA** when computing the ARIMA model description.

Syntax

```
TS_AUTO_ARIMA_RESULT_MODEL_D(auto_arima_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_MODEL_P Function [Time Series]

Retrieves the *p* value produced by **TS_AUTO_ARIMA** when computing the ARIMA model description.

Syntax

```
TS_AUTO_ARIMA_RESULT_MODEL_P(auto_arima_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_MODEL_Q [Time Series]

Retrieves the *q* value produced by **TS_AUTO_ARIMA** when computing the ARIMA model description.

Syntax

```
TS_AUTO_ARIMA_RESULT_MODEL_Q(auto_arima_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_MODEL_S function [Time Series]

Retrieves the *s* value produced by TS_AUTO_ARIMA when computing the ARIMA model description.

Syntax

```
TS_AUTO_ARIMA_RESULT_MODEL_S(auto_arima_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_ARIMA_RESULT_RESIDUAL_SIGMA [Time Series]

Retrieves the residual standard error of the outlier-free data points.

Syntax

```
TS_AUTO_ARIMA_RESULT_RESIDUAL_SIGMA (auto_arima_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_AUTO_UNI_AR Function [Time Series]

Performs automatic selection and fitting of a univariate autoregressive time series model.

Syntax

```
TS_AUTO_UNI_AR (timeseries_expression, ar_count, ar_elem, method)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_BOX_COX_XFORM Function [Time Series]

Performs a forward or inverse Box-Cox power transformation.

Syntax

```
TS_BOX_COX_XFORM (timeseries_expression, power [, shift [,  
inverse] ]) OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_DIFFERENCE Function [Time Series]

Differences a seasonal or nonseasonal time series.

Syntax

```
TS_DIFFERENCE (timeseries_expression, period1 [, period2 [, ...period10] ]) OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the Time Series Guide for detailed information on this function.

TS_DOUBLE_ARRAY [Time Series]

A supporting function for the TS_GARCH function. Constructs a logical array consisting of three to ten constant double precision floating point values, and returns a single varbinary value.

Syntax

```
TS_DOUBLE_ARRAY(xguess1, xguess2, xguess3, [ ... [ , xguess10] ... ])
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the Time Series Guide for detailed information on this function.

TS_ESTIMATE_MISSING Function [Time Series]

Estimates the missing values in a time series and returns them as a new time series, interspersed with the original time series.

Syntax

```
TS_ESTIMATE_MISSING (timeseries_expression, method)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the Time Series Guide for detailed information on this function.

TS_GARCH Function [Time Series]

Computes estimates of the parameters of a GARCH (p, q) model.

Syntax

```
TS_GARCH (<time series expression>, <garch_count>, <arch_count>, <xguess_binary_encoding>, [, <max_sigma>])
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the Time Series Guide for detailed information on this function.

TS_GARCH_RESULT_A Function [Time Series]

A supporting function for the `TS_GARCH` function. Retrieves the Log-Likelihood output parameter, *A*, produced by the `TS_GARCH` aggregate function.

Syntax

```
TS_GARCH_RESULT_A(ts_garch_result )
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_GARCH_RESULT_AIC Function [Time Series]

A supporting function for the `TS_GARCH` function. Retrieves the Akaike Information Criterion output parameter, *AIC*, produced by the `TS_GARCH` aggregate function.

Syntax

```
TS_GARCH_RESULT_AIC(ts_garch_result)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_GARCH_RESULT_USER [Time Series]

A supporting function for the `TS_GARCH` function. Accesses each element in the logical array that describes the `GARCH(p, q)` model.

Syntax

```
TS_GARCH_RESULT_USER(ts_garch_result, model_element_number)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_INT_ARRAY [Time Series]

A supporting function for the `TS_AUTO_ARIMA` function and the `TS_AUTO_ARIMA_OUTLIER` function. Constructs a logical array of constant integer values encoded as a varbinary value.

Syntax

```
TS_INT_ARRAY(int1, int2, int3, int4, [ ... [, int10] ... ] )
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_LACK_OF_FIT Function [Time Series]

Performs the lack-of-fit test for a univariate time series or transfer function, given the appropriate correlation function.

Syntax

```
TS_LACK_OF_FIT (timeseries_expression, p_value, q_value, lagmax,  
[tolerance])
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_LACK_OF_FIT_P Function [Time Series]

Performs the lack-of-fit test for a univariate time series. This function is identical to the TS_LACK_OF_FIT function, except that it returns the p-value of q, rather than returning q.

Syntax

```
TS_LACK_OF_FIT_P (timeseries_expression, p_value, q_value, lagmax,  
[tolerance])
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_MAX_ARMA_AR Function [Time Series]

Calculates the exact maximum likelihood estimation of the parameters in a univariate ARMA (autoregressive moving average) time series model, and returns the requested autoregressive estimate.

Syntax

```
TS_MAX_ARMA_AR (timeseries_expression, ar_count, ar_elem)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_MAX_ARMA_CONST Function [Time Series]

Calculates the exact maximum likelihood estimation of the parameters in a univariate ARMA (autoregressive moving average) time series model, and returns the constant estimate.

Syntax

```
TS_MAX_ARMA_CONST (timeseries_expression)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_MAX_ARMA_LIKELIHOOD Function [Time Series]

Calculates the exact maximum likelihood estimation of the parameters in a univariate ARMA (autoregressive moving average) time series model, and returns likelihood value (ln) for the fitted model.

Syntax

```
TS_MAX_ARMA_LIKELIHOOD (timeseries_expression)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_MAX_ARMA_MA Function [Time Series]

Calculates the exact maximum likelihood estimation of the parameters in a univariate ARMA (autoregressive moving average) time series model, and returns the requested moving average estimate.

Syntax

```
TS_MAX_ARMA_MA (timeseries_expression, ma_count, ma_elem)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_OUTLIER_IDENTIFICATION Function [Time Series]

Detects and determines outliers and simultaneously estimates the model parameters in a time series where the underlying outlier-free series follows a general seasonal or non-seasonal ARMA model.

Syntax

```
TS_OUTLIER_IDENTIFICATION (timeseries_expression, p_value, q_value,  
s_value, d_value, [, delta_value[, critical_value]])
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_PARTIAL_AUTOCORRELATION Function [Time Series]

Calculates the sample partial autocorrelation function of a stationary time series.

Syntax

```
TS_PARTIAL_AUTOCORRELATION (timeseries_expression, lagmax, lag_elem)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

TS_VWAP Function [Time Series]

VWAP stands for volume-weighted average price. **TS_VWAP** calculates the ratio of the value traded to the total volume traded over a particular time horizon.

VWAP is a measure of the average price of a stock over a defined trading horizon. You can use **TS_VWAP** as both a simple and an OLAP-style aggregate function.

Unlike the other time series functions, **TS_VWAP** does not call the IMSL libraries.

Syntax 1

```
TS_VWAP (price_expression, volume_expression)
```

Syntax 2

```
TS_VWAP (price_expression, volume_expression)
```

```
OVER (window-spec)
```

Note: This function is available only with RAP – The Trading Edition Enterprise. See the *Time Series Guide* for detailed information on this function.

UCASE Function [String]

Converts all characters in a string to uppercase.

Syntax

```
UCASE ( string-expression )
```

Parameters

Parameter	Description
string-expression	The string to be converted to uppercase.

Returns

VARCHAR

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **UCASE** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license, or use **CAST** and set **UCASE** to the correct data type and size.

Example

The following statement returns the value “CHOCOLATE”:

```
SELECT UCASE( 'ChocoLate' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—**UCASE** is not supported by Adaptive Server Enterprise, but **UPPER** provides the same feature in a compatible manner

See also

- *LCASE Function [String]* on page 217
- *LEFT Function [String]* on page 220
- *LOWER Function [String]* on page 228
- *REPLACE Function [String]* on page 279
- *REVERSE Function [String]* on page 283
- *RIGHT Function [String]* on page 284
- *UPPER Function [String]* on page 332

UPPER Function [String]

Converts all characters in a string to uppercase.

Syntax

```
UPPER ( string-expression )
```

Parameters

Parameter	Description
string-expression	The string to be converted to uppercase.

Returns

VARCHAR

NVARCHAR

LONG VARCHAR

LONG NVARCHAR

Note: The result data type is a LONG VARCHAR. If you use **UPPER** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license, or use **CAST** and set **UPPER** to the correct data type and size.

Example

The following statement returns the value “CHOCOLATE”:

```
SELECT UPPER( 'ChocoLate' ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Compatible with Adaptive Server Enterprise.

See also

- *LCASE Function [String]* on page 217
- *LEFT Function [String]* on page 220
- *LOWER Function [String]* on page 228
- *REPLACE Function [String]* on page 279
- *REVERSE Function [String]* on page 283
- *RIGHT Function [String]* on page 284
- *UCASE Function [String]* on page 331

USER_ID Function [System]

Returns an integer user identification number.

Syntax

```
USER_ID ( [ user-name ] )
```

Parameters

Parameter	Description
user-name	The user name.

Returns

INT

Examples

The following statement returns the user identification number 1:

```
SELECT USER_ID ('DBA') FROM iq_dummy
```

The following statement returns the user identification number 0:

```
SELECT USER_ID ('SYS') FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ.

See also

- *SUSER_ID Function [System]* on page 316
- *SUSER_NAME Function [System]* on page 317
- *USER_NAME Function [System]* on page 333

USER_NAME Function [System]

Returns the user name.

Syntax

```
USER_NAME ( [ user-id ] )
```

Parameters

Parameter	Description
user-id	The user identification number.

Returns

LONG VARCHAR

Note: The result data type is a LONG VARCHAR. If you use **USER_NAME** in a **SELECT INTO** statement, you must have an Unstructured Data Analytics Option license, or use **CAST** and set **USER_NAME** to the correct data type and size.

Examples

The following statement returns the value “DBA”:

```
SELECT USER_NAME ( 1 ) FROM iq_dummy
```

The following statement returns the value “SYS”:

```
SELECT USER_NAME ( 0 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Adaptive Server Enterprise function implemented for Sybase IQ. In Adaptive Server Enterprise, USER_NAME returns the user name, not the server user name.

See also

- *SUSER_ID Function [System]* on page 316
- *SUSER_NAME Function [System]* on page 317
- *USER_ID Function [System]* on page 333

UUIDTOSTR Function [String]

Converts a unique identifier value (UUID, also known as GUID) to a string value.

Syntax

```
UUIDTOSTR ( uuid-expression )
```

Parameters

Table 95. Parameters

Parameter	Description
uuid-expression	A unique identifier value.

Returns

VARCHAR

Example

To convert a unique identifier value into a readable format, execute a query similar to:


```
CREATE TABLE T3 (
pk uniqueidentifier primary key,c1 int);
INSERT INTO T3 (pk, c1)
values (0x12345678123456789012123456789012, 1)
SELECT UIDTOSTR(pk) FROM T3
```

Usage

Converts a unique identifier to a string value in the format `xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx`, where x is a hexadecimal digit. If the binary value is not a valid unique identifier, NULL is returned.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Binary Data Types* on page 78
- *NEWID Function [Miscellaneous]* on page 239
- *STRTOUUID Function [String]* on page 311
- *Character Data Types* on page 69
- *Binary Data Types* on page 689

VAR_POP Function [Aggregate]

Computes the statistical variance of a population consisting of a numeric-expression, as a DOUBLE.

Syntax

```
VAR_POP ( [ ALL ] expression )
```

Parameters

Parameter	Description
expression	The expression (commonly a column name) whose population-based variance is calculated over a set of rows.

Returns

DOUBLE

Examples

The following statement lists the average and variance in the number of items per order in different time periods:

```
SELECT year( ShipDate ) AS Year, quarter( ShipDate )
AS Quarter, AVG( Quantity ) AS Average,
```

SQL Functions

```
VAR_POP( Quantity ) AS Variance
FROM SalesOrderItems GROUP BY Year, Quarter
ORDER BY Year, Quarter
```

Year	Quarter	Average	Variance
2000	1	25.775148	203.9021
2000	2	27.050847	225.8109
...

Usage

Computes the population variance of the provided *value expression* evaluated for each row of the group or partition (if `DISTINCT` was specified, then each row that remains after duplicates have been eliminated), defined as the sum of squares of the difference of *value expression*, from the mean of *value expression*, divided by the number of rows (remaining) in the group or partition.

Population-based variances are computed according to the following formula:

$$\frac{\sum (x_i - \bar{x})^2}{n}$$

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103

VAR_SAMP Function [Aggregate]

Computes the statistical variance of a sample consisting of a numeric-expression, as a `DOUBLE`.

Syntax

```
VAR_SAMP ( [ ALL ] expression )
```

Parameters

Parameter	Description
expression	The expression (commonly a column name) whose sample-based variance is calculated over a set of rows.

Returns

DOUBLE

Examples

The following statement lists the average and variance in the number of items per order in different time periods:

```
SELECT year( ShipDate ) AS Year, quarter( ShipDate )
      AS Quarter, AVG( Quantity ) AS Average,
      VAR_SAMP( Quantity ) AS Variance
FROM SalesOrderItems GROUP BY Year, Quarter
ORDER BY Year, Quarter
```

Year	Quarter	Average	Variance
2000	1	25.775148	205.1158
2000	2	27.050847	227.0939
...

Usage

Note: `VAR_SAMP` is an alias of `VARIANCE`.

Computes the sample variance of *value expression* evaluated for each row of the group or partition (if `DISTINCT` was specified, then each row that remains after duplicates have been eliminated), defined as the sum of squares of the difference of *value expression*, from the mean of *value expression*, divided by one less than the number of rows (remaining) in the group or partition.

`NULL` returns `NULL` for a one-element input set in Sybase IQ 12.7 and later. In versions earlier than 12.7, `NULL` returned zero.

Variances are computed according to the following formula, which assumes a normal distribution:

$$\frac{\sum (x_i - \bar{x})^2}{n}$$

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103

VARIANCE Function [Aggregate]

Returns the variance of a set of numbers.

Syntax

```
VARIANCE ( [ ALL ] expression )
```

Parameters

Parameter	Description
expression	Any numeric data type (FLOAT, REAL, or DOUBLE) expression. The expression (commonly a column name) whose sample-based variance is calculated over a set of rows.

Returns

DOUBLE

Examples

Given this data:

```
SELECT Salary FROM Employees WHERE DepartmentID = 300
```

salary
51432.000
57090.000
42300.000
43700.00
36500.000
138948.000
31200.000
58930.00
75400.00

The following statement returns the value 1063923790.99999994:

```
SELECT VARIANCE ( Salary ) FROM Employees  
WHERE DepartmentID = 300
```

Given this data:

```
SELECT UnitPrice FROM Products WHERE name = 'Tee Shirt'
```

UnitPrice
9.00
14.00
14.00

The following statement returns the value 8.33333333333334327:

```
SELECT VARIANCE ( UnitPrice ) FROM Products
WHERE name = 'Tee Shirt'
```

Usage

The formula used to calculate **VARIANCE** is

$$var = \frac{n \sum x^2 - (\sum x)^2}{n(n-1)}$$

VARIANCE returns a result of data type double-precision floating-point. If applied to the empty set, the result is NULL, which returns NULL for a one-element input set.

VARIANCE does not support the keyword DISTINCT. A syntax error is returned if DISTINCT is used with **VARIANCE**.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *Windowing Aggregate Function Usage* on page 103
- *STDDEV Function [Aggregate]* on page 303
- *STDDEV_SAMP Function [Aggregate]* on page 306

WEEKS Function [Date and Time]

Returns the number of weeks since an arbitrary starting date/time, returns the number of weeks between two specified date/times, or adds the specified integer-expression number of weeks to a date/time.

Syntax

```
WEEKS ( datetime-expression
| datetime-expression, datetime-expression
| datetime-expression, integer-expression )
```

Parameters

Parameter	Description
<i>datetime-expression</i>	A date and time.
<i>integer-expression</i>	The number of weeks to be added to the <i>datetime-expression</i> . If <i>integer-expression</i> is negative, the appropriate number of weeks are subtracted from the date/time value. Hours, minutes, and seconds are ignored. If you supply an integer expression, the <i>datetime-expression</i> must be explicitly cast as a DATETIME data type.

Returns

Syntax 1 returns an INTEGER.

Syntax 2 returns a TIMESTAMP.

Examples

The following statement returns the value 104278:

```
SELECT WEEKS( '1998-07-13 06:07:12' ) FROM iq_dummy
```

The following statement returns the value 9, to signify the difference between the two dates:

```
SELECT WEEKS( '1999-07-13 06:07:12',
'1999-09-13 10:07:12' ) FROM iq_dummy
```

The following statement returns the timestamp value 1999-06-16 21:05:07.000:

```
SELECT WEEKS( CAST( '1999-05-12 21:05:07'
AS TIMESTAMP ), 5) FROM iq_dummy
```

Usage

Weeks are defined as going from Sunday to Saturday, as they do in a North American calendar. The number returned by the first syntax is often useful for determining if two dates are in the same week.

```
WEEKS ( invoice_sent ) = WEEKS ( payment_received ) FROM iq_dummy
```

In the second syntax, the value of **WEEKS** is calculated from the number of Sundays between the two dates. Hours, minutes, and seconds are ignored. This function is not affected by the `DATE_FIRST_DAY_OF_WEEK` option.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *CONVERT Function [Data Type Conversion]* on page 144
- *HOURS Function [Date and Time]* on page 201
- *MINUTES Function [Date and Time]* on page 234
- *MONTHS Function [Date and Time]* on page 237
- *REPLACE Function [String]* on page 279
- *SECOND Function [Date and Time]* on page 290
- *YEAR Function [Date and Time]* on page 345
- *YEARS Function [Date and Time]* on page 346

WEIGHTED_AVG Function [Aggregate]

Calculates an arithmetically (or linearly) weighted average.

A weighted average is an average in which each quantity to be averaged is assigned a weight. Weightings determine the relative importance of each quantity that make up the average.

Syntax

```
WEIGHTED_AVG ( expression )
```

```
OVER ( window-spec )
```

window-spec: See the Usage section, below.

Parameters

Parameter	Description
expression	A numeric expression for which a weighted value is being computed.

Usage

Use the **WEIGHTED_AVG** function to create a weighted moving average. In a weighted moving average, weights decrease arithmetically over time. Weights decrease from the highest weight for the most recent data points, down to zero.

Figure 1: WEIGHTED_AVG Calculation

$$WMA_M = \frac{np_M + (n - 1)p_{M-1} + \dots + 2p_{M-n+2} + p_{M-n+1}}{n + (n - 1) + \dots + 2 + 1}$$

To exaggerate the weighting, you can average two or more weighted moving averages together, or use an **EXP_WEIGHTED_AVG** function instead.

You can specify elements of *window-spec* either in the function syntax (inline), or with a **WINDOW** clause in the **SELECT** statement.

window-spec:

- Must contain an ORDER BY specifier.
- Cannot contain FOLLOWING or RANGE specifiers.
- The second argument of the ROW specifier—if provided—must be CURRENT ROW.
- Cannot contain NULL values.
- Cannot contain the DISTINCT specifier.
- UNBOUNDED PRECEDING is supported, but may result in poor performance if used

Example

The following example returns a weighted average of salaries by department for employees in Florida, with the salary of recently hired employees contributing the most weight to the average:

```
SELECT DepartmentID, Surname, Salary,
WEIGHTED_AVG(Salary) OVER (PARTITION BY DepartmentID
ORDER BY YEAR(StartDate) DESC) as "W_AVG"
FROM Employees
WHERE State IN ('FL') ORDER BY DepartmentID
```

The returned result set is:

Table 96. WEIGHTED_AVG result set

DepartmentID	Surname	Salary	W_AVG
100	Lull	87,900.000	87,900.000000
100	Gowda	59,840.000	69,193.333333
200	Sterling	64,900.000	64,900.000000
200	Kelly	87,500.000	79,966.666667
300	Litton	58,930.000	58,930.000000
400	Evans	68,940.000	68,940.000000

DepartmentID	Surname	Salary	W_AVG
400	Charlton	28,300.000	41,846.666667
400	Francis	53,870.000	47,858.333333

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.

See also

- *EXP_WEIGHTED_AVG Function [Aggregate]* on page 188
- *Windowing Aggregate Function Usage* on page 103

WIDTH_BUCKET Function [Numerical]

For a given expression, the **WIDTH_BUCKET** function returns the bucket number that the result of this expression will be assigned after it is evaluated.

Syntax

```
WIDTH_BUCKET ( expression, min_value, max_value, num_buckets )
```

Parameters

Parameter	Description
<i>expression</i>	The expression for which the histogram is being created. This expression must evaluate to a numeric or datetime value or to a value that can be implicitly converted to a numeric or datetime value. If <i>expr</i> evaluates to null, then the expression returns null.
<i>min_value</i>	An expression that resolves to the end points of the acceptable range for <i>expr</i> . Must also evaluate to numeric or datetime values and cannot evaluate to null.
<i>max_value</i>	An expression that resolves to the end points of the acceptable range for <i>expr</i> . Must also evaluate to numeric or datetime values and cannot evaluate to null.
<i>num_buckets</i>	Is an expression that resolves to a constant indicating the number of buckets. This expression must evaluate to a positive integer.

Examples

The following example creates a ten-bucket histogram on the `credit_limit` column for customers in Massachusetts in the sample table and returns the bucket number (“Credit Group”) for each customer. Customers with credit limits greater than the maximum value are assigned to the overflow bucket, 11:

```
select EmployeeID, Surname, Salary, WIDTH_BUCKET(Salary, 29000,
60000, 4) "Wages" from Employees where State = 'FL' order by "Wages"
```

EMPLOYEEID	SURNAME	SALARY	Wages
888	Charlton	28300.000	0
1390	Litton	58930.000	4
207	Francis	53870.000	4
266	Gowda	59840.000	4
445	Lull	87900.000	5
1021	Sterling	64900.000	5
902	Kelly	87500.000	5
1576	Evans	68940.000	5

When the bounds are reversed, the buckets are open-closed intervals. For example: **WIDTH_BUCKET** (*credit_limit*, 5000, 0, 5). In this example, bucket number 1 is (4000, 5000], bucket number 2 is (3000, 4000], and bucket number 5 is (0, 1000]. The overflow bucket is numbered 0 (5000, +infinity), and the underflow bucket is numbered 6 (-infinity, 0].

Usage

You can generate equiwidth histograms with the **WIDTH_BUCKET** function. Equiwidth histograms divide data sets into buckets whose interval size (highest value to lowest value) is equal. The number of rows held by each bucket will vary. A related function, **NTILE**, creates equiheigh buckets.

Equiwidth histograms can be generated only for numeric, date or datetime data types; therefore, the first three parameters should be all numeric expressions or all date expressions. Other types of expressions are not allowed. If the first parameter is **NULL**, the result is **NULL**. If the second or the third parameter is **NULL**, an error message is returned, as a **NULL** value cannot denote any end point (or any point) for a range in a date or numeric value dimension. The last parameter (number of buckets) should be a numeric expression that evaluates to a positive integer value; 0, **NULL**, or a negative value will result in an error.

Buckets are numbered from 0 to (n+1). Bucket 0 holds the count of values less than the minimum. Bucket(n+1) holds the count of values greater than or equal to the maximum specified value.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

YEAR Function [Date and Time]

Returns a 4-digit number corresponding to the year of the given date/time.

Syntax

```
YEAR ( datetime-expression )
```

Parameters

Parameter	Description
datetime-expression	A date and time.

Returns

SMALLINT

Example

The following statement returns the value 1998:

```
SELECT YEAR( '1998-07-13 06:07:12' ) FROM iq_dummy
```

Usage

The **YEAR** function is the same as the first syntax of the **YEARS** function.

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *CONVERT Function [Data Type Conversion]* on page 144
- *HOURS Function [Date and Time]* on page 201
- *MINUTES Function [Date and Time]* on page 234
- *MONTHS Function [Date and Time]* on page 237
- *REPLACE Function [String]* on page 279
- *SECOND Function [Date and Time]* on page 290
- *WEEKS Function [Date and Time]* on page 340
- *YEARS Function [Date and Time]* on page 346
- *NTILE Function [Analytical]* on page 244

YEARS Function [Date and Time]

Returns a 4-digit number corresponding to the year of a given date/time, returns the number of years between two specified date/times, or adds the specified integer-expression number of years to a date/time.

Syntax

```
YEARS ( datetime-expression
| datetime-expression, datetime-expression
| datetime-expression, integer-expression )
```

Parameters

Parameter	Description
<i>datetime-expression</i>	A date and time.
<i>integer-expression</i>	The number of years to be added to the <i>datetime-expression</i> . If <i>integer-expression</i> is negative, the appropriate number of years are subtracted from the datetime value. If you supply an integer expression, the <i>datetime-expression</i> must be explicitly cast as a DATETIME data type.

Returns

Syntax 1 returns an INTEGER.

Syntax 2 returns a TIMESTAMP.

Examples

The following statement returns the value 1998:

```
SELECT YEARS( '1998-07-13 06:07:12' ) FROM iq_dummy
```

The following statement returns the value 2, to signify the difference between the two dates.

```
SELECT YEARS( '1997-07-13 06:07:12',
'1999-09-13 10:07:12' ) FROM iq_dummy
```

The following statement returns the YEARS(cast('1999-05-12 21:05:07' as timestamp), 5) value 2004-05-12 21:05:07.000:

```
SELECT YEARS( CAST( '1999-05-12 21:05:07'
AS TIMESTAMP ), 5) FROM iq_dummy
```

Usage

The first syntax of the **YEARS** function is the same as the **YEAR** function.

The second syntax returns the number of years from the first date to the second date, calculated from the number of first days of the year between the two dates. The number might be negative. Hours, minutes, and seconds are ignored. For example, the following statement returns 2, which is the number of first days of the year between the specified dates:

```
SELECT YEARS ( '2000-02-24', '2002-02-24' ) FROM iq_dummy
```

The next statement also returns 2, even though the difference between the specified dates is not two full calendar years. The value 2 is the number of first days of the year (in this case January 01, 2001 and January 01, 2002) between the two dates.

```
SELECT YEARS ( '2000-02-24', '2002-02-20' ) FROM iq_dummy
```

The third syntax adds an *integer-expression* number of years to the given date. If the new date is past the end of the month (such as **SELECT YEARS (CAST ('1992-02-29' AS **TIMESTAMP**), 1)**), the result is set to the last day of the month. If *integer-expression* is negative, the appropriate number of years is subtracted from the date. Hours, minutes, and seconds are ignored.

Standards and compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

See also

- *CAST Function [Data Type Conversion]* on page 133
- *CONVERT Function [Data Type Conversion]* on page 144
- *HOURS Function [Date and Time]* on page 201
- *MINUTES Function [Date and Time]* on page 234
- *MONTHS Function [Date and Time]* on page 237
- *REPLACE Function [String]* on page 279
- *SECOND Function [Date and Time]* on page 290
- *WEEKS Function [Date and Time]* on page 340
- *YEAR Function [Date and Time]* on page 345

YMD Function [Date and Time]

Returns a date value corresponding to the given year, month, and day of the month.

Syntax

```
YMD ( integer-expression1, integer-expression2, integer-expression3 )
```

Parameters

Parameter	Description
integer-expression1	The year.
integer-expression2	The number of the month. If the month is outside the range 1–12, the year is adjusted accordingly.
integer-expression3	The day number. The day is allowed to be any integer, the date is adjusted accordingly.

Returns

DATE

Examples

The following statement returns the value 1998-06-12:

```
SELECT YMD( 1998, 06, 12 ) FROM iq_dummy
```

If the values are outside their normal range, the date adjusts accordingly. For example, the following statement returns the value 1993-03-01:

```
SELECT YMD( 1992, 15, 1 ) FROM iq_dummy
```

The following statement returns the value 1993-02-28:

```
SELECT YMD ( 1992, 15, 1-1 ) FROM iq_dummy
```

The following statement returns the value 1992-02-29:

```
SELECT YMD ( 1992, 3, 1-1 ) FROM iq_dummy
```

Standards and Compatibility

- SQL—Vendor extension to ISO/ANSI SQL grammar.
- Sybase—Not supported by Adaptive Server Enterprise.

Differences from Other SQL Dialects

Sybase IQ conforms to the ANSI SQL89 standard, but has many additional features that are defined in the IBM DB2 and SAA specifications, as well as in the ANSI SQL92 standard.

Certain Sybase IQ features are not found in many other SQL implementations.

Dates

Sybase IQ has date, time, and timestamp types that include year, month, day, hour, minutes, seconds, and fraction of a second. For insertions or updates to date fields, or comparisons with date fields, a free-format date is supported.

In addition, the following operations are allowed on dates:

Table 97. Date Operations

Date Operation	Description
date + integer	Add the specified number of days to a date.
date - integer	Subtract the specified number of days from a date.
date - date	Compute the number of days between two dates.
date + time	Make a timestamp out of a date and time.

Also, many functions are provided for manipulating dates and times.

Integrity

Sybase IQ supports both entity and referential integrity.

This has been implemented via the following two extensions to the **CREATE TABLE** and **ALTER TABLE**.

```
PRIMARY KEY ( column-name, ... )
[NOT NULL] FOREIGN KEY [role-name]
    [(column-name, ...)]
    REFERENCES table-name [(column-name, ...)]
    [ CHECK ON COMMIT ]
```

The **PRIMARY KEY** clause declares the primary key for the relation. IQ will then enforce the uniqueness of the primary key, and ensure that no column in the primary key contains the **NULL** value.

Differences from Other SQL Dialects

The **FOREIGN KEY** clause defines a relationship between this table and another table. This relationship is represented by a column (or columns) in this table which must contain values in the primary key of another table. The system then ensures referential integrity for these columns; whenever these columns are modified or a row is inserted into this table, these columns are checked to ensure that either one or more is **NULL** or the values match the corresponding columns for some row in the primary key of the other table. For more information, see **CREATE TABLE** statement.

Joins

Sybase IQ allows **automatic joins** between tables.

In addition to the **NATURAL** and **OUTER** join operators supported in other implementations, Sybase IQ allows **KEY** joins between tables based on foreign-key relationships. This reduces the complexity of the **WHERE** clause when performing joins.

Updates

Sybase IQ allows more than one table to be referenced by **UPDATE**.

Views defined on more than one table can also be updated. Many SQL implementations do not allow updates on joined tables.

Altering Tables

ALTER TABLE has been extended.

In addition to changes for entity and referential integrity, the following types of alterations are allowed:

```
ADD column data-type
MODIFY column data-type
DELETE column
RENAME new-table-name
RENAME old-column TO new-column
```

You can use **MODIFY** to change the maximum length of a character column, as well as converting from one data type to another.

Subqueries Not Always Allowed

Unlike SQL Anywhere, Sybase IQ does not allow subqueries to appear wherever expressions are allowed.

Sybase IQ supports subqueries only as allowed in the SQL-1989 grammar, plus in the **SELECT** list of the top level query block or in the **SET** clause of an **UPDATE** statement. Sybase IQ does not support SQL Anywhere extensions.

Many SQL implementations allow subqueries only on the right side of a comparison operator. For example, the following command is valid in Sybase IQ but not valid in most other SQL implementations.

```
SELECT      SurName,
            BirthDate,
            ( SELECT DepartmentName
              FROM Departments
              WHERE DepartmentID = Employees.EmployeeID
              AND DepartmentID = 200 )
FROM Employees
```

Additional Functions

Sybase IQ supports several functions not in the ANSI SQL definition.

See also

- *SQL Functions* on page 99

Cursors

When using Embedded SQL, cursor positions can be moved arbitrarily on the **FETCH** statement. Cursors can be moved forward or backward relative to the current position or a given number of records from the beginning or end of the cursor.

Physical Limitations

Limitations exist on the size of objects and the number of objects in Sybase IQ databases. In most cases, computer memory and disk drive are more limiting factors.

For limitations that apply to only one platform, see the platform-specific documentation.

Table 98. Sybase IQ Database Object Size and Number Limitations

Item	Limitation
Catalog file size	Maximum is 1TB for all platforms. Windows systems with NTFS support the 1TB maximum.
Database name size	250 bytes
Database size	Maximum database size approximates the number of files times the file size on a particular platform, depending on the maximum disk configuration. Refer to your operating system documentation for kernel parameters that affect the maximum number of files.
Dbfile size	Determined by the operating system file size
DbSPACE size	Raw device: Maximum size is 4T. File system device: Maximum size is 4TB Operating system file: Maximum size supported by the operating system Sybase does not recommend creating dbspaces on NAS (Network Attached Storage) devices
Field size	255 bytes for BINARY, 32,767 bytes for VARBINARY 32,767 for CHAR, VARCHAR Up to 512 TB for 128 KB pages or 1 PB for 512 KB pages for LONG BINARY, LONG VARCHAR
IQ page size	Must be between 64KB and 512KB
Maximum key size	255 bytes for single-column index. 5300 bytes for multicolumn index
Maximum length of string literal	32KB

Physical Limitations

Item	Limitation
Maximum length of SQL statement	<p>The maximum length of SQL statements is limited to the amount of memory available for the IQ catalog, and to the size of the catalog stack</p> <p>If your SQL statements are long, increase the catalog stack size using -gss, and increase catalog memory cache size using -c or a combination of -ch and -cl</p> <p>When printing the SQL statement in error messages, the text is limited to the IQ catalog page size. To print long commands, you can start the server with an increased -gp setting, although in general Sybase recommends that you use the default of -gp 4096.</p>
Maximum length of variable-length FILLER column	512 bytes
Maximum number of dbfiles	The total number of files that can be opened depends on the number of unique file descriptors that an operating system can support
Maximum number of users (connected and concurrent)	<p>1000 on 64-bit platforms AIX, HP, Linux, and Sun Solaris</p> <p>200 on 64-bit platforms on Windows</p>
Maximum size of temp extract file	<p>Set by TEMP_EXTRACT_SIZE option. Platform limits are:</p> <p>AIX & HP-UX: 0 – 64GB</p> <p>Sun Solaris: 0 – 512GB</p> <p>Windows: 0 – 128GB</p> <p>Linux: 0 – 512GB</p>
Number of columns per table	Sybase IQ supports up to 45,000 columns in a table. You may see performance degradation if you have more than 10,000 columns in a table. The limit on the number of columns per table that allow NULLs is approximately $8 * (\text{database-page-size} - 30)$.
Number of events per database	$2^{31} - 1 = 2\ 147\ 483\ 647$
Number of files per database	Operating system limit that user can adjust; for example, using NOFILE. Typically, 2047 files per database
Number of indexes	2^{32} (~4,000,000) per table
Number of rows per table	Limited by table size, upper limit $2^{48} - 1$
Number of stored procedures per database	$2^{32} - 1 = 4\ 294\ 967\ 295$

Item	Limitation
Number of tables or views in a single FROM clause	16 – 64, depending on the query, with join optimizer turned on
Number of tables or views referenced per query	512
Number of tables per database	4,293,918,719
Number of tables per join index (number of tables that can be joined in one query block)	31
Number of tables referenced per transaction	No limit.
Number of UNION branches per query	512. If each branch has multiple tables in the FROM clause, the limit on tables per query reduces the number of UNION branches allowed.
Number of values in an IN list	250,000
Row size	Sybase recommends a limit of half the page size
Table size	Limited by database size

See also

- *String Functions* on page 114

Physical Limitations

System Procedures

Use the system-supplied stored procedures in Sybase IQ databases to retrieve system information.

Sybase IQ includes the following kinds of system procedures:

- System functions that are implemented as stored procedures.
- Catalog stored procedures, for displaying system information in tabular form.
- Multiplex stored procedures, which include both of the above types of procedures, for multiplex server operations.
See *Using Sybase IQ Multiplex > Multiplex Reference > System Procedures*.
- Transact-SQL system and catalog procedures.

System stored procedures related specifically to Large Object data, including **sp_iqsetcompression** and **sp_iqshowcompression**, are described in *Unstructured Data Analytics in Sybase IQ > Stored Procedure Support*.

Syntax Rules for Stored Procedures

Use of parentheses and quotes in stored procedure calls varies, depending on whether you enter the procedure name directly, as you can in Interactive SQL, or invoke it with a **CALL** statement.

Some variations are permitted because the product supports both Sybase IQ SQL and Transact-SQL syntax. If you need Transact-SQL compatibility, be sure to use Transact-SQL syntax.

Table 99. Stored Procedure Syntax Variations

Syntax	Syntax type	Explanation
<code>procedure_name ('param')</code>	Sybase IQ	Quotes are required if you enclose parameters in parentheses.
<code>procedure_name 'param'</code>	Sybase IQ	Parentheses are optional if you enclose parameters in quotes.

Syntax	Syntax type	Explanation
<code>procedure_name param</code>	Transact-SQL	<p>If you omit quotes around parameters, you must also omit parentheses.</p> <hr/> <p>Note: Quotes are always required around parameters when the owner is specified. For example, assuming the owner is <i>dba</i>,</p> <pre>sp_iqtablesiz 'dba.emp1'</pre> <p>requires quotes around the parameters.</p> <pre>sp_iqtablesiz emp1</pre> <p>does not.</p>
<code>procedure_name</code>	Sybase IQ or Transact-SQL	Use this syntax if you run a procedure with no parameters directly in Interactive SQL, and the procedure has no parameters.
<code>call procedure_name (param='value')</code>	Sybase IQ	Use this syntax to call a procedure that passes a parameter value.

When you use Transact-SQL stored procedures, you must use the Transact-SQL syntax.

Understanding Statistics Reported by Stored Procedures

Many stored procedures report information on the state of Sybase IQ at the time the procedure executes.

This means that you get a snapshot view. For example, a report column that lists space in use by a connection shows only the space in use at the instant the procedure executes, not the maximum space used by that connection.

To monitor Sybase IQ usage over an extended period, use the Sybase IQ monitor, which collects and reports statistics from the time you start the monitor until you stop it, at an interval you specify.

System Stored Procedures

System stored procedures carry out System Administrator tasks in the IQ main store.

System stored procedures are owned by the user ID dbo.

Note: By default, the maximum length of column values displayed by Interactive SQL Classic is 30 characters. This might be inadequate for displaying output of stored procedures such as **sp_iqstatus**. To avoid truncated output, increase the length by selecting Command > Options

from the Interactive SQL menu select and enter a higher value for Limit Display Columns, Limit Output Columns, or both.

sa_char_terms System Procedure

Breaks a CHAR string into terms and returns each term as a row along with its position.

See *Unstructured Data Analytics in Sybase IQ > Stored Procedure Support > Managing Terms in a TEXT Index > sa_char_terms System Procedure*.

sa_dependent_views system procedure

Returns the list of all dependent views for a given table or view.

Syntax

```
sa_dependent_views( 'tbl_name' [, 'owner_name' ] )
```

Arguments

- **tbl_name** – Use this CHARACTER parameter to specify the name of the table or view.
- **owner_name** – Use this optional CHARACTER parameter to specify the owner for *tbl_name*.

Result set

Column name	Data type	Description
table_id	UNSIGNED INTEGER	The object ID of the table or view.
dep_view_id	UNSIGNED INTEGER	The object ID of the dependent views.

Remarks

Use this procedure to obtain the list of IDs of dependent views. Alternatively, you can use the procedure in a statement that returns more information about the views, such as their names.

No errors are generated if no existing tables satisfy the specified criteria for table and owner names. Also:

- *tbl_name* is optional and has a default value of NULL.
- If both *owner* and *tbl_name* are null, information is returned on all tables that have dependent views.
- If *tbl_name* is null but *owner* is specified, information is returned on all tables owned by the specified owner.
- If *tbl_name* is specified but *owner* is null, information is returned on any one of the tables with the specified name.

System Procedures

By default, execution of the procedure does not require any permissions and assumes that PUBLIC has access to the catalog. DBAs can control access as needed on the view and/or the catalog.

Permissions

None

Side effects

None

Examples

In this example, the sa_dependent_views system procedure is used to obtain the list of IDs for the views that are dependent on the SalesOrders table. The procedure returns the table_id for SalesOrders, and the dep_view_id for the dependent view, ViewSalesOrders.

```
CALL sa_dependent_views( 'SalesOrders' );
```

In this example, the sa_dependent_views system procedure is used in a SELECT statement to obtain the list of names of views dependent on the SalesOrders table. The procedure returns the ViewSalesOrders view.

```
SELECT t.table_name FROM SYSTAB t,  
sa_dependent_views( 'SalesOrders' ) v  
WHERE t.table_id = v.dep_view_id;
```

sa_external_library_unload Procedure

Unloads an external library.

See *Unstructured Data Analytics in Sybase IQ > Stored Procedure Support > External Library Identification > sa_external_library_unload System Procedure*.

sa_list_external_library Procedure

Lists the external libraries currently loaded in the server.

See *Unstructured Data Analytics in Sybase IQ > Stored Procedure Support > Identifying External Libraries > sa_list_external_library Procedure*.

sa_nchar_terms system Procedure

Breaks an NCHAR string into terms and returns each term as a row along with its position.

See *Unstructured Data Analytics in Sybase IQ > Stored Procedure Support > Managing Terms in a TEXT index > sa_nchar_terms System Procedure*.

sa_text_index_vocab Procedure

Lists all terms that appear in a **TEXT** index, and the total number of indexed values that each term appears in.

See *Unstructured Data Analytics in Sybase IQ > Stored Procedure Support > Managing Terms in a TEXT Index > sa_text_index_vocab System Procedure*.

sa_get_user_status system procedure

Allows you to determine the current status of users.

Syntax

```
sa_get_user_status( )
```

Arguments

None

Result set

Column name	Data type	Description
user_id	UNSIGNED INT	A unique number identifying the user.
user_name	CHAR(128)	The name of the user.
connections	INT	The current number of connections by this user.
failed_logins	UNSIGNED INT	The number of failed login attempts made by the user.
last_login_time	TIMESTAMP	The time the user last logged in.
locked	TINYINT	Indicates if the user account is locked.
reason_locked	LONG VARCHAR	The reason the account is locked.

Remarks

This procedure returns a result set that shows the current status of users. In addition to basic user information, the procedure includes a column indicating if the user has been locked out and a column with a reason for the lockout. Users can be locked out for the following reasons: locked due to policy, password expiry, or too many failed attempts.

System Procedures

A user without DBA authority can obtain user information by creating and executing a cover procedure owned by a DBA.

Permissions

DBA authority is required to view information about all users. Users without DBA authority can view their own information. In addition, users without DBA authority can view information about other users by executing a cover procedure owned by a DBA.

Side effects

None

Example

The following example uses the `sa_get_user_status` system procedure to return the status of database users.

```
CALL sa_get_user_status;
```

sp_expireallpasswords Procedure

Causes all user passwords to expire immediately.

Syntax1

```
call sp_expireallpasswords
```

Syntax2

```
sp_expireallpasswords
```

Permissions

DBA or USER ADMIN authority required.

Examples

Causes all user passwords to expire immediately:

```
call sp_expireallpasswords
```

See also

- *sp_iqaddlogin Procedure* on page 363
- *sp_iqcopyloginpolicy Procedure* on page 394
- *sp_iqmodifyadmin Procedure* on page 457
- *sp_iqmodifylogin Procedure* on page 458
- *sp_iqpassword Procedure* on page 465

sp_iqaddlogin Procedure

Adds a new Sybase IQ user account to the specified login policy.

Syntax1

```
call sp_iqaddlogin ('username_in', 'pwd',
[ 'password_expiry_on_next_login ' ] [ , 'policy_name ' ] )
```

Syntax2

```
sp_iqaddlogin 'username_in', 'pwd', [ 'password_expiry_on_next_login ' ]
[ , 'policy_name ' ]
```

Syntax3

```
sp_iqaddlogin username_in, pwd, [ password_expiry_on_next_login ] [ ,
policy_name ]
```

Usage

Table 100. Parameters

Parameter	Description
username_in	The user's login name. Login names must conform to the rules for identifiers
pwd	The user's password. Passwords must conform to rules for passwords, that is, they must be valid identifiers.
password_expiry_on_next_login	(Optional) Specifies whether user's password expires as soon as this user's login is created. Default setting is OFF (password does not expire).
policy_name	(Optional) Creates the user under the named login policy. If unspecified, user is created under the root login policy.

A username_in/pwd created using **sp_iqaddlogin** and set to expire in one day is valid all day tomorrow and invalid on the following day. In other words, a login created today and set to expire in n days are not usable once the date changes to the $(n+1)$ th day.

Permissions

Requires DBA authority.

Description

Adds a new Sybase IQ user account, assigns a login policy to the user and adds the user to the ISYSUSER system table. If the user already has a user ID for the database but is not in

System Procedures

ISYSUSER, (for example, if the user ID was added using the **GRANT CONNECT** statement or Sybase Central), **sp_iqaddlogin** adds the user to the table.

If you do not specify a login policy name when calling the procedure, Sybase IQ assigns the user to the root login policy.

Note: If the maximum number of logins for a login policy is unlimited, then a user belonging to that login policy can have an unlimited number of connections.

The first user login forces a password change and assigns a login policy to the newly created user. Sybase recommends that you use **CREATE USER** to create new users, although, for backward compatibility, **sp_iqaddlogin** is still supported.

See *Reference: Statements and Options > SQL Statements > CREATE USER Statement*.

Examples

These calls add the user `rose` with a password `irk324` under the login policy named `expired_password`. This example assumes the `expired_password` login policy already exists.

```
call sp_iqaddlogin('rose', 'irk324', 'ON', 'expired_password')
sp_iqaddlogin 'rose', 'irk324', 'ON', 'expired_password'
```

See also

- *sp_expireallpasswords Procedure* on page 362
- *sp_iqcopyloginpolicy Procedure* on page 394
- *sp_iqmodifyadmin Procedure* on page 457
- *sp_iqmodifylogin Procedure* on page 458
- *sp_iqpassword Procedure* on page 465
- *sp_iqdroplogin Procedure* on page 414

sp_iqbackupdetails Procedure

Shows all the dbfiles included in a particular backup.

Syntax

```
sp_iqbackupdetails backup_id
```

Parameters

Table 101. Parameters

Parameter	Description
backup_id	Specifies the backup operation transaction identifier.

Note: You can obtain the backup_id value from the SYSIQBACKUPHISTORY table. Run the following query: `select * from sysiqbackuphistory`

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

sp_iqbackupdetails returns the following:

Table 102. sp_iqbackupdetails columns

Column name	Description
backup_id	Identifier for the backup transaction.
backup_time	Time of the backup.
backup_type	Type of backup: "Full," "Incremental since incremental," or "Incremental since full."
selective_type	Subtype of backup: "All inclusive", "All RW files in RW dbspaces", "Set of RO dbspace/file."
depends_on_id	Identifier for previous backup that the backup depends on.
dbspace_id	Identifier for the dbspace being backed up.
dbspace_name	Name of the dbspace from SYSIQBACKUPHISTORYDETAIL. If dbspace name matches the dbspace name in SYSDBSPACE for a given dbspace_id. Otherwise "null."
dbspace_rwstatus	"ReadWrite" or "Read Only."
dbspace_createid	Dbspace creation transaction identifier.
dbspace_alterid	Alter DBSPACE read-write mode transaction identifier.
dbspace_online	Status. Values are "Online" or "Offline."
dbspace_size	Size of dbspace, in KB, at time of backup.
dbspace_backup_size	Size of data, in KB, backed up in the dbspace.
dbfile_id	Identifier for the dbfile being backed up
dbfile_name	The logical file name, if it was not renamed after the backup operation. If renamed, "null."
dbfile_rwstatus	"ReadWrite" or "Read Only".

Column name	Description
dbfile_createid	Dbfile creation transaction identifier.
dbfile_alterid	Alter DBSPACE alter FILE read-write mode transaction identifier
dbfile_size in MB	Size of the dbfile, in KB.
dbfile_backup_size	Size of the dbfile backup, in KB.
dbfile_path	The dbfile path from SYSBACKUPDETAIL, if it matches the physical file path (“file_name”) in SYSDBFILE for a given dbspace_id and the dbfile_id. Otherwise “null.”

Example

Sample output from **sp_iqbackupdetails**:

```

backup_id      backup_time          backup_type      selective_type    d
depends_on_id
      883      2008-09-23 13:58:49.0    Full            All
inclusive                                0

dbspace_id     dbspace_name        dbspace_rwstatus  dbspace_createid
      0         system              ReadWrite         0

dbspace_alterid  dbspace_online     dbspace_size      dbspace_backup_size
dbfile_id
      0              0              2884              2884              0

dbfile_name     dbfile_rwstatus     dbfile_createid  dbfile_alterid
dbfile_size
      system          ReadWrite         0              0              2884

dbfile_backup_size  dbfile_path
      2884          C:\\Documents and Settings\\All Users\\SybaseIQ\\
\\demo\\iqdemo.db
    
```

See also

- *SYSIQBACKUPHISTORY* System View on page 608

sp_iqbackupsummary Procedure

Summarizes backup operations performed.

Syntax

```

sp_iqbackupsummary [ timestamp or backup_id ]
    
```


*Parameters***Table 103. Parameters**

Parameter	Description
timestamp or backup_id	Specifies the interval for which to report backup operations. If you specify a timestamp or a backup ID, only those records with backup_time greater than or equal to the time you enter are returned. If you specify no timestamp, the procedure returns all the backup records in ISYSIQ-BACKUPHISTORY.

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

sp_iqbackupsummary returns the following:

Table 104. sp_iqbackupsummary columns

Column name	Description
backup_id	Identifier for the backup transaction
backup_time	Time of the backup
backup_type	Type of backup: "Full," "Incremental since incremental," or "Incremental since full"
selective_type	Subtype of backup: "All Inclusive", "All RW files in RW dbspaces", "Set of RO dbspace/file"
virtual_type	Type of virtual backup: "Non-virtual," "Decoupled," or "Encapsulated"
depends_on_id	Identifier for backup that the backup depends on
creator	Creator of the backup
backup_size	Size, in KB, of the backup
user_comment	User comment
backup_command	The backup statement issued (minus the comment)

Example

Sample output of **sp_iqbackupsummary**:

```

backup_id      backup_time      backup_type      selective_type    v
virtual_type
      883  2008-09-23 13:58:49.0  Full              All inclusive    Non
virtual

depends_on_id   creator    backup_size  user_comment    backup_command
              0    DBA              10864              backup database to
              'c:\\\\temp
\\\\\\b1'
    
```

sp_iqcardinality_analysis Procedure

Helps you generate indexes by analyzing the cardinality of the columns in the table and recommending the indexes to add.

sp_iqcardinality_analysis can generate an optional SQL script containing ready-to-execute SQL statements for index creation.

sp_iqcardinality_analysis operates independently of the Index Advisor. The Index Advisor gives recommendations based on the usage of a column in actual queries, but only considers cardinality in the case where the column has an LF or HG index, and distinct count allows rebuilding 1-byte FP and 2-byte FP indexes. **sp_iqcardinality_analysis** considers cardinality in all cases, but does not include the Index Advisor recommendations in its index recommendation list.

sp_iqcardinality_analysis does not include HNG and CMP indexes in its list of recommendations. Recommended indexes include:

- LF index
- HG index
- 1-byte FP
- 2-byte FP
- 3-byte FP
- Conversion of 2-byte FP to 1-byte FP
- Primary Key Constraint (HG index)
- Unique Constraint (HG index)
- DATE/TIME/DTTM index
- WD index
- WORD index

Note: Executing **sp_iqcardinality_analysis** against a wide table may take several minutes, depending on the complexity of the column data types and indexes already existing on the table.

Syntax

```
sp_iqcardinality_analysis ( [ 'table_name' ], [ 'table_owner' ], [ 'script' ] )
```

Parameters

Parameter	Description
table_name	Name of the table.
table_owner	Name of the table owner. If this parameter is not specified, then the procedure looks for a table owned by the current user.
script	Generates a SQL script containing ready-to-execute SQL statements that create the indexes, and displays index recommendations. If this option is not specified, then the console displays: <ul style="list-style-type: none"> • table_name • table_owner • column_name • cardinality • index_type • index recommendation

Usage

If you do not specify any parameters, then IQ displays **create_index** SQL statements for all columns in all tables owned by the current user.

If you specify *script*, you can redirect the output to generate the script file:

```
OUTPUT TO 'indexfile.sql' FORMAT ASCII QUOTE '';
```

Permissions

Requires DBA authority.

Example 1

```
sp_iqcardinality_analysis 'onebytefp', 'DBA'
```

Table 105. Example 1 console output

table_name	table_owner	column_name	cardinality	index type	index recommendation
onebytefp	DBA	c1	10	One Byte FP	--Column 'c1' has no One Byte FP index and cardinality is less than 256. --One Byte FP index can be created.--Call the stored procedure: sp_iqrebuildindex 'onebytefp','column c1 255'
onebytefp	DBA	c1	10	LF	--Column 'c1' has no LF index and cardinality is less than 1000. --LF index can be created using CREATE INDEX statement: CREATE LF INEDX IQ_T400_c1_LF ON DBA.onebytefp (c1)

Example 2

```
sp_iqcardinality_analysis 'onebytefp', 'DBA', 'script'
```

Table 106. Example 2 console output

index recommendation
--Column 'c1' has no One Byte FP index and cardinality is less than 256. --One Byte FP index can be created. --Call the stored procedure: sp_iqrebuildindex 'onebytefp','column c1 255' --Column 'c1' has no LF index and cardinality is less than 1000. --LF index can be created using CREATE INDEX statement: CREATE LF INDEX IQ_T400_c1_LF ON onebytefp (c1)

sp_iqcheckdb Procedure

Checks validity of the current database. Optionally corrects allocation problems for dbspaces or databases.

sp_iqcheckdb reads all storage in the database. On successful completion, the database free list (an internal allocation map) is updated to reflect the true storage allocation for the database. **sp_iqcheckdb** then generates a report listing the actions it has performed.

If an error is found, **sp_iqcheckdb** reports the name of the object and the type of error. **sp_iqcheckdb** does not update the free list if errors are detected.

sp_iqcheckdb also allows you to check the consistency of a specified table, index, index type, or the entire database.

Note: `sp_iqcheckdb` is the user interface to the IQ database consistency checker (DBCC) and is sometimes referred to as **DBCC**.

Syntax

```
sp_iqcheckdb 'mode target [ ... ] [ resources resource-percent ]'
```

This is the general syntax of `sp_iqcheckdb`. There are three modes for checking database consistency, and one for resetting allocation maps. The syntax for each mode is listed separately below. If mode and target are not both specified in the parameter string, Sybase IQ returns the error message:

```
At least one mode and target must be specified to DBCC.
```

Parameters

mode: { **allocation** | **check** | **verify** } | **dropleaks**

target: [**indextype** *index-type* [...]] **database** | **database resetclocks** | { [**indextype** *index-type*] [...] **table** *table-name* [**partition** *partition-name*] [...] | **index** *index-name* | [...] **dbspace** *dbspace-name* }

Allocation Mode

```
sp_iqcheckdb 'allocation target [ resources resource-percent ]'
```

Check Mode

```
sp_iqcheckdb 'check target [ resources resource-percent ]'
```

Verify mode

```
sp_iqcheckdb 'verify target [ resources resource-percent ]'
```

Dropleaks Mode

```
sp_iqcheckdb 'dropleaks target [ resources resource-percent ]'
```

Usage

Table 107. Parameters

Parameter	Description
database	If the target is a database, all dbspaces must be online.

Parameter	Description
index-type	<p>One of the following index types: FP, CMP, LF, HG, HNG, WD, DATE, TIME, DTTM, TEXT.</p> <p>If the specified <i>index-type</i> does not exist in the target, an error message is returned. If multiple index types are specified and the target contains only some of these index types, the existing index types are processed by sp_iqcheckdb.</p>
index-name	<p>May contain owner and table qualifiers: [[owner .] table-name .] index-name</p> <p>If <i>owner</i> is not specified, current user and database owner (dbo) are substituted in that order. If <i>table</i> is not specified, <i>index-name</i> must be unique.</p>
table-name	<p>May contain an owner qualifier: [owner .] table-name</p> <p>If <i>owner</i> is not specified, current user and database owner (dbo) are substituted in that order. <i>table-name</i> cannot be a temporary or pre-join table.</p> <hr/> <p>Note: If either the table name or the index name contains spaces, enclose the <i>table-name</i> or <i>index-name</i> parameter in double quotation marks:</p> <pre>sp_iqcheckdb 'check index "dbo.sstab.i2" resources 75'</pre> <hr/> <p>If <i>owner</i> is not specified, current user and database owner (dbo) are substituted in that order. <i>table-name</i> cannot be a temporary or pre-join table.</p> <hr/> <p>Note: If either the table name or the index name contains spaces, enclose the <i>table-name</i> or <i>index-name</i> parameter in double quotation marks:</p> <pre>sp_iqcheckdb 'check index "dbo.sstab.i2" resources 75'</pre>

Parameter	Description
partition-name	<p>The <i>partition-name</i> parameter contains no qualifiers. If it contains spaces, enclose it in double quotation marks.</p> <p>The partition filter causes sp_iqcheckdb to examine a subset of the corresponding table's rows that belong to that partition. A partition filter on a table and table target without the partition filter are semantically equivalent when the table has only one partition.</p>
dbspace-name	<p>The <i>dbspace-name</i> parameter contains no qualifiers. If it contains spaces, enclose it in double quotation marks.</p> <p>The dbspace target examines a subset of the database's pages that belong to that dbspace. The dbspace must be online. The dbspace and database target are semantically equivalent when the table has only one dbspace.</p>
resource-percent	<p>The input parameter <i>resource-percent</i> must be an integer greater than zero. The resources percentage allows you to limit the CPU utilization of the database consistency checker by controlling the number of threads with respect to the number of CPUs. If <i>resource-percent</i> = 100 (the default value), then one thread is created per CPU. If <i>resource-percent</i> > 100, then there are more threads than CPUs, which might increase performance for some machine configurations. The minimum number of threads is one.</p>

Note: The **sp_iqcheckdb** parameter string must be enclosed in single quotes and cannot be greater than 255 bytes in length.

Allocation problems can be repaired in dropleaks mode.

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

sp_iqcheckdb checks the allocation of every block in the database and saves the information in the current session until the next **sp_iqdbstatistics** procedure is issued. **sp_iqdbstatistics** displays the latest result from the most recent execution of **sp_iqcheckdb**.

sp_iqcheckdb can perform several different functions, depending on the parameters specified.

Allocation Mode

Checks allocation with blockmap information for the entire database, a specific index, a specific index type, a specific partition, specific table, or a specific dbspace. Does not check index consistency.

Detects duplicate blocks (blocks for which two or more objects claim ownership) or extra blocks (unallocated blocks owned by an object).

Detects leaked blocks (allocated blocks unclaimed by any object in the specified target) for database or dbspace targets.

When the target is a partitioned table, **allocation mode**:

- Checks metadata of all the table's partition allocation bitmaps
- Checks metadata of the tables allocation bitmap
- Verifies that blockmap entries are consistent with the table's allocation bitmap
- Verifies that none of the table's partition allocation bitmaps overlap
- Checks that rows defined in the table's partition allocation bitmaps form a superset of the table's existence bitmap
- Checks that rows defined in the table's partition allocation bitmaps form a superset of the table's allocation bitmap

Note: **sp_iqcheckdb** cannot check all allocation problems if you specify the name of a single index, index type, or table in the input parameter string.

Run in allocation mode:

- To detect duplicate or unowned blocks (use database or specific tables or indexes as the target)
- If you encounter page header errors

The DBCC option **resetclocks** is used only with allocation mode. **resetclocks** is used with forced recovery to convert a multiplex secondary server to a coordinator. For information on multiplex capability, see *Using Sybase IQ Multiplex*. **resetclocks** corrects the values of internal database versioning clocks, in the event that these clocks are behind. Do not use the **resetclocks** option for any other purpose, unless you contact Sybase IQ Technical Support.

The **resetclocks** option must be run in single-user mode and is allowed only with the DBCC statement **allocation database**. The syntax of **resetclocks** is:


```
sp_iqcheckdb 'allocation database resetclocks'
```

Check Mode

Verifies that all database pages can be read for the entire database, specific index, specific index type, specific table, specific partition, or specific dbspace. If the table is partitioned, then check mode will check the table's partition allocation bitmaps.

Run in check mode if metadata, null count, or distinct count errors are returned when running a query.

Verify Mode

Verifies the contents of non-FP indexes with their corresponding FP indexes for the entire database, a specific index, a specific index type, specific table, specific partition, or specific dbspace. If the specified target contains all data pages for the FP and corresponding non-FP indexes, then verify mode detects the following inconsistencies:

- Missing key – a key that exists in the FP but not in the non-FP index.
- Extra key – a key that exists in the non-FP index but not in the FP index.
- Missing row – a row that exists in the FP but not in the non-FP index.
- Extra row – a row that exists in the non-FP index but not in the FP index.

If the specified target contains only a subset of the FP pages, then verify mode can detect only the following inconsistencies:

- Missing key
- Missing row

If the target is a partitioned table, then verify mode also verifies that each row in the table or table partition has been assigned to the correct partition.

Run in verify mode if metadata, null count, or distinct count errors are returned when running a query.

Note: `sp_iqcheckdb` does not check referential integrity or repair referential integrity violations.

Dropleaks Mode

When the Sybase IQ server runs in single-node mode, you can use dropleaks mode with either a database or dbspace target to reset the allocation map for the entire database or specified dbspace targets. If the target is a dbspace, then the dropleaks operation must also prevent read-write operations on the named dbspace. All dbspaces in the database or dbspace list must be online.

For information on running dropleaks mode on a multiplex, see *Using Sybase IQ Multiplex*.

The following examples illustrate the use of the `sp_iqcheckdb` procedure.

Example 1

In this example, **sp_iqcheckdb** checks the allocation for the entire database:

```
sp_iqcheckdb 'allocation database'
```

Example 2

In the second example, **sp_iqcheckdb** performs a detailed check on indexes `i1`, `i2`, and `dbo.t1.i3`. If you do not specify a new mode, **sp_iqcheckdb** applies the same mode to the remaining targets, as shown in the following command:

```
sp_iqcheckdb 'verify index i1 index i2 index dbo.t1.i3'
```

Example 3

You can combine all modes and run multiple checks on a database in a single session. In the following example, **sp_iqcheckdb** performs a quick check of partition `p1` in table `t2`, a detailed check of index `i1`, and allocation checking for the entire database using half of the CPUs:

```
sp_iqcheckdb 'check table t2 partition p1 verify index i1  
allocation database resources 50'
```

Example 4

This example checks all indexes of the type **FP** in the database:

```
sp_iqcheckdb 'check indextype FP database'
```

Example 5

The following example verifies the **FP** and **HG** indexes in the table `t1` and the **LF** indexes in the table `t2`:

```
sp_iqcheckdb 'verify indextype FP indextype HG table t1 indextype LF  
table t2'
```

Example 6

The following example illustrates one of the three “LVC cells” messages in the output of **sp_iqcheckdb**:

```
sp_iqcheckdb 'check index EFG2JKL.ASIQ_IDX_T208_C504_FP'  
-----  
Index Statistics:  
** Inconsistent Index: abcd.EFG2JKL.ASIQ_IDX_T208_C504_FP ***** FP  
Indexes Checked: 1  
** Unowned LVC Cells: 212 *****
```

The **sp_iqcheckdb** LVC cells messages include:

- Unowned LVC cells
- Duplicate LVC cell rows
- Unallocated LVC cell rows

These messages indicate inconsistencies with a `VARCHAR`, `VARBINARY`, `LONG BINARY` (BLOB), or `LONG VARCHAR` (CLOB) column. Unowned LVC cells represent a small amount

of unusable disk space and can safely be ignored. Duplicate and Unallocated LVC cells are serious errors that can be resolved only by dropping the damaged columns.

To drop a damaged column, create a new column from a copy of the old column, then drop the original column and rename the new column to the old column.

Note: LVC is a VARCHAR or VARBINARY column with a width greater than 255. LONG BINARY (BLOB) and LONG VARCHAR (CLOB) also use LVC.

DBCC performance

The execution time of DBCC varies, depending on the size of the database for an entire database check, the number of tables or indexes specified, and the size of the machine. Checking only a subset of the database (that is, only specified tables, indexes, or index types) requires less time than checking an entire database.

The processing time of **sp_iqcheckdb** dropleaks mode depends on the number of dbspace targets.

This table summarizes the actions and output of the four **sp_iqcheckdb** modes.

Table 108. Actions and output of sp_iqcheckdb modes

Mode	Errors detected	Output	Speed
Allocation	Allocation errors	Allocation statistics only	4TB per hour
Check	Allocation errors Most index errors	All available statistics	60GB per hour
Verify	Allocation errors All index errors	All available statistics	15GB per hour
Dropleaks	Allocation errors	Allocation statistics only	4TB per hour

Output

Depending on the execution mode, **sp_iqcheckdb** output includes summary results, errors, informational statistics, and repair statistics. The output may contain as many as three results sets, if you specify multiple modes in a single session. Error statistics are indicated by asterisks (*****), and appear only if errors are detected.

The output of **sp_iqcheckdb** is also copied to the Sybase IQ message file `.iqmsg`. If the **DBCC_LOG_PROGRESS** option is ON, **sp_iqcheckdb** sends progress messages to the IQ message file, allowing the user to follow the progress of the DBCC operation as it executes.

Output Example

The following is an example of the output you see when you run **sp_iqcheckdb 'allocation database'** and there is leaked space. Leaked space is a block that is allocated according to the

System Procedures

database free list (an internal allocation map), but DBCC finds that the block is not part of any database object. In this example, DBCC reports 32 leaked blocks.

Stat	Value	Flags
=====		
DBCC Allocation Mode Report		
=====		
** DBCC Status	Errors Detected	*****
=====		
Allocation Summary		
=====		
Blocks Total	8192	
Blocks in Current Version	4954	
Blocks in All Versions	4954	
Blocks in Use	4986	
% Blocks in Use	60	
** Blocks Leaked	32	*****
=====		
Allocation Statistics		
=====		
Marked Logical Blocks	8064	
Marked Physical Blocks	4954	
Marked Pages	504	
Blocks in Freelist	126553	
Imaginary Blocks	121567	
Highest PBN in Use	5432	
** 1st Unowned PBN	452	*****
Total Free Blocks	3206	
Usable Free Blocks	3125	
% Free Space Fragmented	2	
Max Blocks Per Page	16	
1 Block Page Count	97	
3 Block Page Count	153	
4 Block Page Count	14	
...		
9 Block Hole Count	2	
16 Block Hole Count	194	
Database Objects Checked	1	
B-Array Count	1	
Blockmap Identity Count	1	
=====		
====Connection Statistics		

sp_iqcheckoptions Procedure

For the connected user, **sp_iqcheckoptions** displays a list of the current value and the default value of database and server startup options that have been changed from the default.

Syntax

```
sp_iqcheckoptions
```

Permissions

None. The DBA sees all options set on a permanent basis for all groups and users and sees temporary options set for DBA. Users who are not DBAs see their own temporary options. All users see nondefault server startup options.

Usage

Requires no parameters. Returns one row for each option that has been changed from the default value. The output is sorted by option name, then by user name.

Description

For the connected user, the **sp_iqcheckoptions** stored procedure displays a list of the current value and the default value of database and server startup options that have been changed from the default. **sp_iqcheckoptions** considers all Sybase IQ and SQL Anywhere database options. Sybase IQ modifies some SQL Anywhere option defaults, and these modified values become the new default values. Unless the new Sybase IQ default value is changed again, **sp_iqcheckoptions** does not list the option.

When **sp_iqcheckoptions** is run, the DBA sees all options set on a permanent basis for all groups and users and sees temporary options set for DBA. Users who are not DBAs see their own temporary options. All users see nondefault server startup options.

Table 109. sp_iqcheckoptions columns

Column name	Description
User_name	The name of the user or group for whom the option has been set. At database creation, all options are set for the <code>public</code> group. Any option that has been set for a group or user other than <code>public</code> is displayed.
Option_name	The name of the option.
Current_value	The current value of the option.
Default_value	The default value of the option.
Option_type	“Temporary” for a TEMPORARY option, else “Permanent”.

Examples

In these examples, the temporary option APPEND_LOAD is set to ON and the group mygroup has the option MAX_WARNINGS set to 9. The user joel has a temporary value of 55 set for MAX_WARNINGS.

In the first example, **sp_iqcheckoptions** is run by the DBA.

User_name	Option_name	Current_value	Default_value	Option_type
DBA	Ansi_update_constr	CURSORS	Off	Permanent
PUBLIC	Ansi_update_constr	Cursors	Off	Permanent
DBA	Append_Load	ON	OFF	Temporary
DBA	Checkpoint_time	20	60	Temporary
DBA	Connection_authent	Company=MyComp; Application=DBTools;Signa		Temporary
DBA	Login_procedure	DBA.sp_iq_proce	sp_login_envir	Permanent
PUBLIC	Login_procedure	DBA.sp_iq_proce	sp_login_envir	Permanent
mygroup	Max_Warnings	9	281474976710655	Permanent
DBA	Thread_count	25	0	Temporary

In the second example, **sp_iqcheckoptions** is run by the user joel.

User_name	Option_name	Current_value	Default_value	Option_type
joel	Ansi_update_constr	CURSORS	Off	Permanent
PUBLIC	Ansi_update_constr	Cursors	Off	Permanent
joel	Checkpoint_time	20	60	Temporary
joel	Connection_authent	Company=MyComp; Application=DBTools;Signa		Temporary
joel	Login_procedure	DBA.sp_iq_proce	sp_login_envir	Permanent
PUBLIC	Login_procedure	DBA.sp_iq_proce	sp_login_envir	Permanent
joel	Max_Warnings	55	281474976710655	Temporary
joel	Thread_count	25	0	Temporary

sp_iqclient_lookup Procedure

Allows a client application to determine the Sybase IQ user account responsible for a particular data stream, as observed in a network analyzer originating from a specific client IP address/port.

Syntax

```
sp_iqclient_lookup [ 'IPaddress' ], [ Port ], [ UserID ]
```

*Parameters***Table 110. Parameters**

Parameter	Description
IPAddress	Specifies the IP address of the originating client application.
Port	Specifies the port number of the originating client application.
UserID	Specifies the Sybase IQ user ID.

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

The **sp_iqclient_lookup** procedure takes the client IP address and port number and returns a single row containing Number (the connection ID), IPAddress, Port, and UserID.

```
1> sp_iqclient_lookup '158.76.235.71',3360
2> go
```

Number	IPAddress	Port	UserID
15	158.76.235.71	3360	rdeniro

Optionally, you can pass a third argument to select only the UserID. If no arguments are passed, sp_iqclient_lookup returns all current logins with their IP addresses and port numbers. For example:

```
sp_iqclient_lookup
```

Number	IPAddress	Port	UserID
11	162.66.131.36	2082	mbrando
21	162.66.100.233	1863	apacino
22	162.66.100.206	8080	jcaan
23	162.66.100.119	6901	rduvall
24	162.66.100.125	7001	dkeaton
25	162.66.100.124	6347	jcazale

```
(6 rows affected)
(return status = 0)
```

If a client application is not using TCP/IP or for internal connections, the address appears as 127.0.0.1.

Note: This information is available for logged on users only. No historical login data is kept on the server for this purpose.

Side Effects

The **sp_iqclient_lookup** stored procedure may impact server performance, which varies from one installation to another. Finding the login name entails scanning through all current active connections on the server; therefore, the impact may be greater on servers with large numbers of connections. Furthermore, this information cannot be cached as it is dynamic — sometimes highly dynamic. It is, therefore, a matter for the local system administrator to manage the use of this stored procedure, as well as monitor the effects on the server, just as for any other client application that uses server facilities.

Examples

Shows IP addresses for UserID jcazale:

```
sp_iqclient_lookup null, null, jcazale
```

Number	IPaddress	Port	UserID
-----	-----	----	-----
11	162.66.131.36	2082	jcazale
15	164.66.131.36	1078	jcazale

Shows IP addresses from client IP 162.66.131.36:

```
sp_iqclient_lookup '162.66.131.36'
```

Number	IPaddress	Port	UserID
-----	-----	----	-----
11	162.66.131.36	2082	jcazale
12	162.66.131.36	1078	jcaan

Note: The result is empty when the user specifies an incorrect argument.

sp_iqcolumn Procedure

Displays information about columns in a database.

Syntax1

```
sp_iqcolumn ( [ table_name ], [ table_owner ], [ table_loc ] )
```

Syntax2

```
sp_iqcolumn [ table_name='table_name' ], [ table_owner='tableowner' ],  
[ table_loc='table_loc' ]
```


Usage

Syntax	Description
Syntax1	If you specify <i>table_owner</i> without specifying <i>table_name</i> , you must substitute NULL for <i>table_name</i> . For example, <code>sp_iqcolumn NULL, DBA</code> .
Syntax2	The parameters can be specified in any order. Enclose ' <i>table_name</i> ' and ' <i>table_owner</i> ' in single quotes.

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Displays information about columns in a database. Specifying the *table_name* parameter returns the columns only from tables with that name. Specifying the *table_owner* parameter returns only tables owned by that user. Specifying both *table_name* and *table_owner* parameters chooses the columns from a unique table, if that table exists. Specifying *table_loc* returns only tables that are defined in that segment type. Specifying no parameters returns all columns for all tables in a database. **sp_iqcolumn** does not return column information for system tables.

Table 111. sp_iqcolumn columns

Column name	Description
<i>table_name</i>	The name of the table
<i>table_owner</i>	The owner of the table
<i>column_name</i>	The name of the column
<i>domain_name</i>	The data type
<i>width</i>	The precision of numeric data types that have precision and scale or the storage width of numeric data types without scale; the width of character data types
<i>scale</i>	The scale of numeric data types
<i>nulls</i>	'Y' if the column can contain NULLS, 'N' if the column cannot contain NULLS

Column name	Description
default	'Identity/Autoincrement' if the column is an identity/autoincrement column, null if not.
cardinality	The distinct count, if known, by indexes
est_cardinality	The estimated number of distinct values, set to 255 automatically if the column was created with the <code>MINIMIZE_STORAGE</code> option ON, or a user-supplied value from the IQ UNIQUE constraint specified in CREATE TABLE
location	TEMP = IQ temporary store, MAIN = IQ main store, SYSTEM = catalog store
isPartitioned	'Y' if the column belongs to a partitioned table and has one or more partitions whose dbspace is different from the table partition's dbspace, 'N' if the column's table is not partitioned or each partition of the column resides in the same dbspace as the table partition.
remarks	User comments added with the COMMENT statement
check	the check constraint expression

See also

- *sp_iquconstraint Procedure* on page 390
- *sp_iqudatatype Procedure* on page 398
- *sp_iquevent Procedure* on page 420
- *sp_iquhelp Procedure* on page 426
- *sp_iquindex and sp_iquindex_alt Procedures* on page 433
- *sp_iqujoinindex Procedure* on page 447
- *sp_iqukeys Procedure* on page 467
- *sp_iquprocparm Procedure* on page 472
- *sp_iqu_reset_identity Procedure* on page 480
- *sp_iquitable Procedure* on page 501
- *sp_iquview Procedure* on page 515

sp_iqucolumn Procedure Example

Use the example as a reference for **sp_iqucolumn** usage.

The following variations in syntax both return all of the columns in the table Departments:

```
sp_iqucolumn Departments
```

```
call sp_iqucolumn (table_name='Departments')
```

table_name	table_owner	column_name	domain_name	width	scale
departments	GROUP0	DepartmentID	integer	4	0
N	(NULL)				
departments	GROUP0	DepartmentName	char	40	0
N	(NULL)				
departments	GROUP0	DepartmentHead	integer	4	0
Y	(NULL)				
cardinality	est_cardinality	location	isPartitioned	remarks	check
5	5	Main	N	(NULL)	(NULL)
0	5	Main	N	(NULL)	(NULL)
5	5	Main	N	(NULL)	(NULL)

The following variation in syntax returns all of the columns in all of the tables owned by table owner DBA.

```
sp_iqcolumn table_owner='DBA'
```

sp_iqcolumnuse Procedure

Reports detailed usage information for columns accessed by the workload.

Syntax

```
sp_iqcolumnuse
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Columns from tables created in SYSTEM are not reported.

Table 112. sp_iqcolumnuse columns

Column name	Description
TableName	Table name
ColumnName	Column name
Owner	User name of column owner
UID**	Column Unique Identifier
LastDT	Date/time of last access
NRef	Number of query references

**UID is a number assigned by the system that uniquely identifies the instance of the column (where instance is defined when an object is created).

System Procedures

For more information on column indexes, see *Reference: Statements and Options > Database Options, Alphabetical List of Options > INDEX_ADVISOR option.*

Example

Sample output from the **sp_iqcolumnuse** procedure:

TableName	NRef	ColumnName	Owner	UID	LastDT
orders	22:41:22.. 13	o_orderdate	DBA	151	20070917
orders	22:41:22.. 13	o_shippriority	DBA	154	20070917
lineitem	22:41:22.. 13	l_orderkey	DBA	186	20070917
lineitem	22:41:22.. 13	l_extendedp..	DBA	191	20070917
lineitem	22:41:22.. 13	l_discount	DBA	192	20070917
lineitem	22:41:22.. 13	l_shipdate	DBA	196	20070917
#tmp1	22:57:36.. 1	expression	DBA	10000000001218	20070917
#tmp1	22:41:58.. 1	expression	DBA	10000000001222	20070917
...					

Note: The long numbers in the example above are temporary IDs.

See also

- *sp_iqindexadvice Procedure* on page 437
- *sp_iqindexuse Procedure* on page 446
- *sp_iqtableuse Procedure* on page 506
- *sp_iqunusedcolumn Procedure* on page 510
- *sp_iqunusedindex Procedure* on page 511
- *sp_iqunusedtable Procedure* on page 513
- *sp_iqworkmon Procedure* on page 521

sp_iqconnection Procedure

Shows information about connections and versions, including which users are using temporary dbspace, which users are keeping versions alive, what the connections are doing inside Sybase IQ, connection status, database version status, and so on.

Syntax

```
sp_iqconnection [ connhandle ]
```

Usage

The input parameter *connhandle* is equal to the `Number` connection property and is the ID number of the connection. The **connection_property** system function returns the connection ID:

```
SELECT connection_property ( 'Number' )
```

When called with an input parameter of a valid *connhandle*, **sp_iqconnection** returns the one row for that connection only.

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

sp_iqconnection returns a row for each active connection. The columns ConnHandle, Name, Userid, LastReqTime, ReqType, CommLink, NodeAddr, and LastIdle are the connection properties Number, Name, Userid, LastReqTime, ReqType, CommLink, NodeAddr, and LastIdle respectively, and return the same values as the system function **sa_conn_info**. The additional columns return connection data from the Sybase IQ side of the Sybase IQ engine. Rows are ordered by ConnCreateTime.

The column MPXServerName stores information related to multiplex Inter-Node Communication (INC), as shown:

Table 113. MPXServerName column values

Server where run	MPXServerName column content
Simplex server	NULL (All connections are local/user connections.)
Multiplex coordinator	<ul style="list-style-type: none"> • NULL for local/user connections • Contains value of secondary node's server name (source of connection) for every INC connection (either on-demand or dedicated heartbeat connection).
Multiplex secondary	<ul style="list-style-type: none"> • NULL for local/user connections • Contains value of coordinator's server name (source of connection).

In Java applications, specify Sybase IQ-specific connection properties from TDS clients in the RemotePWD field. This example, where **myconnection** becomes the IQ connection name, shows how to specify IQ specific connection parameters:

```
p.put("RemotePWD", " , , CON=myconnection" );
```

See SQL Anywhere *Server – Programming*.

Table 114. sp_iqconnection columns

Column name	Description
ConnHandle	The ID number of the connection.
Name	The name of the server.
Userid	The user ID for the connection.
LastReqTime	The time at which the last request for the specified connection started.
ReqType	A string for the type of the last request.
IQCmdType	The current command executing on the Sybase IQ side, if any. The command type reflects commands defined at the implementation level of the engine. These commands consists of transaction commands, DDL and DML commands for data in the IQ store, internal IQ cursor commands, and special control commands such as OPEN and CLOSE DB, BACKUP, RESTORE , and others.
LastIQCmdTime	The time the last IQ command started or completed on the IQ side of the Sybase IQ engine on this connection.
IQCursors	The number of cursors open in the IQ store on this connection.
LowestIQCursorState	The IQ cursor state, if any. If multiple cursors exist on the connection, the state displayed is the lowest cursor state of all the cursors; that is, the furthest from completion. Cursor state reflects internal Sybase IQ implementation detail and is subject to change in the future. For this version, cursor states are: NONE, INITIALIZED, PARSED, DESCRIBED, COSTED, PREPARED, EXECUTED, FETCHING, END_OF_DATA, CLOSED and COMPLETED. As suggested by the names, cursor state changes at the end of the operation. A state of PREPARED, for example, indicates that the cursor is executing.
IQthreads	The number of Sybase IQ threads currently assigned to the connection. Some threads may be assigned but idle. This column can help you determine which connections are using the most resources.
TxnID	The transaction ID of the current transaction on the connection. This is the same as the transaction ID displayed in the <code>.iqmsg</code> file by the BeginTxn, CmtTxn, and PostCmtTxn messages, as well as the Txn ID Seq logged when the database is opened.
ConnCreateTime	The time the connection was created.
TempTableSpaceKB	The number of kilobytes of IQ temporary store space in use by this connection for data stored in IQ temp tables.

Column name	Description
TempWorkSpaceKB	The number of kilobytes of IQ temporary store space in use by this connection for working space such as sorts, hashes, and temporary bitmaps. Space used by bitmaps or other objects that are part of indexes on Sybase IQ temporary tables are reflected in TempTableSpaceKB.
IQConnID	The ten-digit connection ID displayed as part of all messages in the .iqmsg file. This is a monotonically increasing integer unique within a server session.
satoiq_count	An internal counter used to display the number of crossings from the SQL Anywhere side to the IQ side of the Sybase IQ engine. This might be occasionally useful in determining connection activity. Result sets are returned in buffers of rows and do not increment satoiq_count or iqtosa_count once per row.
iqtosa_count	An internal counter used to display the number of crossings from the IQ side to the SQL Anywhere side of the Sybase IQ engine. This might be occasionally useful in determining connection activity.
CommLink	The communication link for the connection. This is one of the network protocols supported by Sybase IQ, or is local for a same-machine connection.
NodeAddr	The node for the client in a client/server connection.
LastIdle	The number of ticks between requests.
MPXServerName	If an INC connection, the varchar(128) value contains the name of the multiplex server where the INC connection originates. NULL if not an INC connection.
LSName	The logical server name of the connection. NULL if logical server context is unknown or not applicable.

Example

The following is an example of **sp_iqconnection** output:

```

ConnHandle      Name          Userid          LastReqTime     ReqType
=====
1  'SQL_DBC_100525210' 'DBA'          '2011-03-28 09:29:24.466' 'OPEN'

          IQCmdType      LastIQCmdTime      IQCursors      LowestIQCursorState
          =====
'IQUTILITYOPENCURSOR' 2011-03-28 09:29:24.0      0              'NONE'

IQthreads      TxnID          ConnCreateTime  TempTableSpaceKB  TempWorkSpaceKB
=====
0  3352568      2011-03-28 09:29:20.0      0              0

IQconnID satoiq_count iqtosa_count  CommLink  NodeAdd  LastIdle  MPXServerName  LSName
=====
34      43          2  'local'    ''      244  (NULL)      Finance_LS

```

sp_iqconstraint Procedure

Lists referential integrity constraints defined using **CREATE TABLE** or **ALTER TABLE** for the specified table or column.

Syntax

```
sp_iqconstraint [ 'table-name', 'column-name', 'table-owner' ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

If table name and column name are omitted, reports all referential integrity constraints for all tables including temporary ones in the current connected database. The information includes unique or primary key constraint, referential constraint, and associated role name that are defined by the **CREATE TABLE** and/or **ALTER TABLE** statements.

Example

This is sample output that displays all primary key/foreign key pairs where either the candidate key or foreign key contains column ck1 for owner bob in all tables:

```
call sp_iqconstraint('','ck1','bob')
```

```
PTAB1 bob ASIQ_IDX_T27_HG unique ck1,ck2 selftab bob CK6FK3 Y
ASIQ_IDX_T42_HG ck1,ck2PTAB2 bob ASIQ_IDX_T27_HG unique ck1,ck2
selftab bob CK6FK4 Y
ASIQ_IDX_T206_I42_HG ck1,ck2selftab bob ASIQ_IDX_T26_HG unique
ck1,ck2 selftab bob CK3FK1 Y
ASIQ_IDX_T206_I42_HG ck1,ck2
```

The columns displayed are:

- Primary enforced table
- Table owner
- Candidate key index
- Primary key or Unique
- Primary key columns
- Foreign table
- Foreign table owner
- Foreign key role name
- Enforced status (“Y” for enforced, “N” for unenforced)
- Foreign key index
- Foreign key columns
- Location (“TEMP,” “MAIN,” or “SYSTEM”)

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqpkeys Procedure* on page 467
- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iqcontext Procedure

Tracks and displays, by connection, information about statements that are currently executing.

Syntax

```
sp_iqcontext [ connhandle ]
```

Usage

The input parameter *connhandle* is equal to the `NUMBER` connection property and is the ID number of the connection. For example, **SELECT CONNECTION_PROPERTY('NUMBER')**.

When called with an input parameter of a valid *connhandle*, **sp_iqcontext** returns the information only for that connection.

Permissions

DBA authority required. Users without DBA authority must be granted `EXECUTE` permission to run the stored procedure.

Distributed Query Processing (DQP) worker threads will not be shown if you do not have `MULTIPLY ADMIN` or DBA authority for the server that is the leader node of the thread.

Description

sp_iqcontext lets the DBA determine what statements are running on the system at any given moment, and identify the user and connection that issued the statement. With this information, you can use this utility to:

- Match the statement text with the equivalent line in **sp_iqconnection** to get resource usage and transactional information about each connection
- Match the statement text to the equivalent line in the SQL log created when the **-zr** server option is set to `ALL` or `SQL`

System Procedures

- Use connection information to match the statement text in **sp_iqcontext** to the equivalent line in the `.iqmsg` file, which includes the query plan, when Sybase IQ can collect it
- Match statement text to an IQ stack trace (`stktrc-yyyyymmdd-hhnnss_#.iq`), if one is produced
- Collate this information with an operating system stack trace that might be produced, such as `pstack` on Sun Solaris

The maximum size of statement text collected is the page size of the catalog store.

Table 115. sp_iqcontext columns

Column name	Description
ConnOrCursor	CONNECTION, CURSOR, or DQP.
ConnHandle	The ID number of the connection or 0 for DQP.
Name	The name of the server (leader name).
Userid	The user ID for the connection, cursor, or DQP worker.
numIQCursors	If column 1 is CONNECTION, the number of cursors open on this connection. If column 1 is CURSOR, a number assigned sequentially to cursors associated with this connection. If column 1 is DQP, then 0. CONNECTION can also return a value of 0.
IQthreads	The number of IQ threads currently assigned to the connection. Some threads may be assigned but idle. For DQP threads, indicates the number of threads assigned to the DQP worker.
TxnID	The transaction ID of the current transaction. In the case of a worker thread, indicates the leader's transaction ID.
ConnOrCurCreateTime	The time this connection, cursor, or DQP worker was created.
IQConnID	The connection ID displayed as part of all messages in the <code>.iqmsg</code> file. This is a monotonically increasing integer unique within a server session.

Column name	Description
IQGovernPriority	A value that indicates the order in which the queries of a user are queued for execution. 1 indicates high priority, 2 (the default) medium priority, and 3 low priority. A value of -1 indicates that IQGovernPriority does not apply to the operation. Set the IQGovernPriority value with the database option IQGOVERN_PRIORITY. For DQP connections, this column displays No command.
CmdLine	First 4096 characters of the user command being executed. For DQP connections, this column displays No command.
Attributes	Unique ID for the query being distributed.

Example

The following example shows an excerpt from output when `sp_iqcontext` is issued with no parameter, producing results for all current connections. Column names are truncated due to space considerations.

```
ConnOrCu.. ConnHandle Name UserId numIQ.. IQthr.. TxnID Conn.. IQcon..
IQGov.. Cmd.. Attributes
CONNECTION 2 sun7bar dbo 0 0 0 2010-08-04 15:15:40.0 15 No command NO
COMMAND
CONNECTION 7 sun7bar dbo 0 0 0 2010-08-04 15:16:00.0 32 No command NO
COMMAND
CONNECTION 10 sun7bar dbo 0 0 0 2010-08-04 15:16:21.0 46 No command NO
COMMAND
...
CONNECTION 229 sun7bar DBA 0 0 1250445 2010-08-05 18:28:16.0 50887 2
select server_name,
inc_state, coordinator_failover from sp_iqmpxinfo() order by server_name
...
DQP 0 dbsrv2873_node_c1DBA 0 1 10000 2010-08-05 18:28:16.0 no command no
command Query ID:
12345; Condition: c1 > 100;
DQP 0 dbsrv2873_node_c1DBA 0 1 10001 2010-08-05 18:28:16.0 no command no
command Query ID:
12346; Node #12 Join (Hash);
```

The first line of output shows connection 2 (IQ connection ID 15). This connection is on server sun7bar, user dbo. This connection was not executing a command when `sp_iqcontext` was issued.

Connection 229 shows the user command being executed (the command contains less than the maximum 4096 characters the column can display). The 2 before the user command fragment indicates that this is a medium priority query.

The connection handle (2 for the first connection in this example) identifies results in the **-zr** log. The IQ connection ID (15 for the first connection in this example) identifies results in the `.iqmsg` file. On UNIX systems, you can use **grep** to locate all instances of the connection handle or connection ID, making it easy to correlate information from all sources.

The second-last line (TxnID 10000) shows a DQP worker thread. The worker connection is running two invariant conditions.

The last line (TxnID 10001) shows connection is running a hash join.

See also

- *CONNECTION_PROPERTY* Function [System] on page 142
- *sp_iqshowpsexec* Procedure on page 484

sp_iqcopyloginpolicy Procedure

Creates a new login policy by copying an existing one.

Syntax1

```
call sp_iqcopyloginpolicy ('existing-policy-name', 'new-policy-name')
```

Syntax2

```
sp_iqcopyloginpolicy 'existing-policy-name', 'new-policy-name'
```

Syntax3

```
sp_iqcopyloginpolicy existing-policy-name, new-policy-name policy_name ]
```

Usage

Table 116. Parameters

Parameter	Description
existing policy name	The login policy to copy.
new policy name	Name of the new login policy to create (CHAR(128)).

Permissions

Requires DBA authority.

Examples

The following stored procedure creates a new login policy named *lockeduser* by copying the login policy option values from the existing login policy named *root*.

```
call sp_iqcopyloginpolicy ('root', 'lockeduser')
```

See also

- *sp_expireallpasswords Procedure* on page 362
- *sp_iqaddlogin Procedure* on page 363
- *sp_iqmodifyadmin Procedure* on page 457
- *sp_iqmodifylogin Procedure* on page 458
- *sp_iqpassword Procedure* on page 465

sp_iqcursorinfo Procedure

Displays detailed information about cursors currently open on the server.

Syntax

```
sp_iqcursorinfo [ cursor-name ] [ , conn-handle ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Parameter	Description
cursor-name	The name of the cursor. If only this parameter is specified, sp_iqcursorinfo returns information about all cursors that have the specified name in all connections. The login policy to copy.
conn-handle	An integer representing the connection ID. If only this parameter is specified, sp_iqcursorinfo returns information about all cursors in the specified connection.

The **sp_iqcursorinfo** procedure can be invoked without any parameters. If no parameters are specified, **sp_iqcursorinfo** returns information about all cursors currently open on the server. If both parameters are specified, **sp_iqcursorinfo** reports information about all of the cursors that have the specified name and are in the specified connection.

If you do not specify the first parameter, but specify the second parameter, you must substitute NULL for the omitted parameter. For example, `sp_iqcursorinfo NULL, 1`.

Table 117. sp_iqcursorinfo usage examples

Syntax	Output
sp_iqcursorinfo	Displays information about all cursors currently open on the server
sp_iqcursorinfo 'cursor1'	Displays information about the all cursors named <code>cursor1</code> in all connections
sp_iqcursorinfo NULL, 3	Displays information about all cursors in connection 3
sp_iqcursorinfo 'cursor2', 4	Displays information about all the cursors named <code>cursor2</code> in connection 4

Description

The **sp_iqcursorinfo** stored procedure displays detailed information about cursors currently open on the server. The **sp_iqcursorinfo** procedure enables database administrators to monitor cursor status using just one stored procedure and view statistics such as how many rows have been updated, deleted, and inserted.

If you specify one or more parameters, the result is filtered by the specified parameters. For example, if *cursor-name* is specified, only information about the specified cursor is displayed. If *conn-handle* is specified, **sp_iqcursorinfo** returns information only about cursors in the specified connection. If no parameters are specified, **sp_iqcursorinfo** displays information about all cursors currently open on the server.

The **sp_iqcursorinfo** procedure returns information in the following columns:

Table 118. sp_iqcursorinfo columns

Column name	Description
Name	The name of the cursor
ConnHandle	The ID number of the connection
IsUpd	Y: the cursor is updatable; N otherwise
IsHold	Y: the cursor is a hold cursor; N otherwise
IQConnID	The ten digit connection ID displayed as part of all messages in the <code>.iqmsg</code> file. This number is a monotonically increasing integer unique within a server session.
UserID	User ID (or user name) for the user who created and ran the cursor
CreateTime	The time of cursor creation
CurrentRow	The current position of the cursor in the result set

Column name	Description
NumFetch	The number of times the cursor fetches a row. The same row can be fetched more than once.
NumUpdate	The number of times the cursor updates a row, if the cursor is updatable. The same row can be updated more than once.
NumDelete	The number of times the cursor deletes a row, if the cursor is updatable.
NumInsert	The number of times the cursor inserts a row, if the cursor is updatable.
RWTabOwner	The owner of the table that is opened in RW mode by the cursor.
RWTabName	The name of the table that is opened in RW mode by the cursor.
CmdLine	The first 4096 characters of the command the user executed

See *System Administration Guide: Volume 2 > Using Procedures and Batches > Using Cursors in Procedures*. For information about cursors in transactions, see *System Administration Guide: Volume 1 > Transactions and Versioning > Cursors in Transactions*.

For CURSOR database options, see *Reference: Statements and Options > Database Options > Alphabetical List of Options*.

For CURSOR statements, see *Reference: Statements and Options > SQL Statements*.

Example

Display information about all cursors currently open on the server:

```
sp_iqcursorinfo
```

Name	ConnHandle	IsUpd	IsHold	IQConnID	UserID

crsr1	1	Y	N	118	DBA
crsr2	3	N	N	118	DBA
CreateTime		CurrentRow	NumFetch	NumUpdate	
2009-06-26 15:24:36.000		19	100000000	200000000	
2009-06-26 15:38:38.000		20000	200000000		
NumDelete	NumInsert	RWTabOwner	RWTabName	CmdLine	

20000000	3000000000	DBA	test1	call procl() call proc2()	

sp_iqdatatype Procedure

Displays information about system data types and user-defined data types.

Syntax

```
sp_iqdatatype [ type-name ], [ type-owner ], [ type-type ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 119. Parameters

Parameter	Description
type-name	The name of the data type.
type-owner	The name of the creator of the data type.
type-type	The type of data type. Allowed values are: <ul style="list-style-type: none"> • SYSTEM: displays information about system defined data types (data types owned by user SYS or dbo) only • ALL: displays information about user and system data types • Any other value: displays information about user data types

The **sp_iqdatatype** procedure can be invoked without any parameters. If no parameters are specified, only information about user-defined data types (data types not owned by dbo or SYS) is displayed by default.

If you do not specify either of the first two parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. For example, `sp_iqdatatype NULL, NULL, SYSTEM` and `sp_iqdatatype NULL, user1`.

Table 120. sp_iqdatatype usage examples

Syntax	Output
<code>sp_iqdatatype</code>	Displays information about all user-defined data types in the database

Syntax	Output
<code>sp_iqdatatype country_t</code>	Displays information about the user-defined data type named <code>country_t</code>
<code>sp_iqdatatype non_existing_type</code>	No rows returned, as the data type <code>non_existing_type</code> does not exist
<code>sp_iqdatatype NULL, DBA</code>	Displays information about all user-defined data types owned by DBA
<code>sp_iqdatatype country_t, DBA</code>	Displays information about the data type <code>country_t</code> owned by DBA
<code>sp_iqdatatype rowid</code>	<code>rowid</code> is a system-defined data type. If there is no user-defined data type also named <code>rowid</code> , no rows are returned. (By default, only user-defined data types are returned.)
<code>sp_iqdatatype rowid, SYS</code>	No rows returned, as the data type <code>rowid</code> is not a user-defined data type (by default, only user-defined data types are returned)
<code>sp_iqdatatype NULL, NULL, SYSTEM</code>	Displays information about all system defined data types (owned by <code>dbo</code> or <code>SYS</code>)
<code>sp_iqdatatype rowid, NULL, SYSTEM</code>	Displays information about the system data type <code>rowid</code>
<code>sp_iqdatatype NULL, NULL, 'ALL'</code>	Displays information about the user-defined and system data types

Description

The **sp_iqdatatype** stored procedure displays information about system and user-defined data types in a database. User-defined data types are also referred to as domains. Predefined domain names are not included in the **sp_iqdatatype** output.

If you specify one or more parameters, the **sp_iqdatatype** result is filtered by the specified parameters. For example, if *type-name* is specified, only information about the specified data type is displayed. If *type-owner* is specified, **sp_iqdatatype** only returns information about data types owned by the specified owner. If no parameters are specified, **sp_iqdatatype** displays information about all the user-defined data types in the database.

The **sp_iqdatatype** procedure returns information in the following columns:

Table 121. sp_iqdatatype columns

Column name	Description
type_name	The name of the data type
creator	The owner of the data type
nulls	Y indicates the user-defined data type allows nulls; N indicates the data type does not allow nulls and U indicates the null value for the data type is unspecified.
width	Displays the length of string columns, the precision of numeric columns, and the number of bytes of storage for all other data types
scale	Displays the number of digits after the decimal point for numeric data type columns and zero for all other data types
“default”	The default value for the data type
“check”	The CHECK condition for the data type

Example

Display information about the user-defined data type `country_t`:

```
sp_iqdatatype country_t
```

```
type_name  creator  nulls  width  scale  "default"  "check"
country_t  DBA     U      15     0      (NULL)     (NULL)
```

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqpkeys Procedure* on page 467
- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iqdbsize Procedure

Displays the size of the current database.

Syntax

```
sp_iqdbsize( [ main ] )
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Returns the total size of the database. Also returns the number of pages required to hold the database in memory and the number of IQ pages when the database is compressed (on disk).

Table 122. sp_iqdbsize columns

Column name	Description
Database	The path name of the database file.
Physical Blocks	Total database size in blocks. An IQ database consists of one or more dbspaces. Each dbspace has a fixed size, which is originally specified in units of megabytes. This megabyte quantity is converted to blocks using the IQ page size and the corresponding block size for that IQ page size. The Physical Blocks column reflects the cumulative total of each Sybase IQ dbspace size, represented in blocks.
KBytes	The total size of the database in kilobytes. This value is the total size of the database in blocks (Physical Blocks in the previous sp_iqdb-size column) multiplied by the block size. The block size depends on the IQ page size.
Pages	The total number of IQ pages necessary to represent in memory all of the data stored in tables and join indexes, as well as the metadata for these objects. This value is always greater than or equal to the value of Compressed Pages (the next sp_iqdbsize column).
Compressed Pages	The total number of IQ pages necessary to store on disk the data in tables and join indexes as well as the metadata for these objects. This value is always less than or equal to the value of Pages (the previous sp_iqdbsize column), because Sybase IQ compresses pages when the IQ page is written from memory to disk. The sp_iqdbsize Compressed Pages column represents the number of compressed pages.

Column name	Description
NBlocks	The total size in blocks used to store the data in tables and join indexes. This value is always less than or equal to the sp_iqdbsize Physical Blocks value.
Catalog Blocks	The total size in blocks used to store the metadata for tables and join indexes.

You can use this procedure in a multiplex environment. See *Using Sybase IQ Multiplex > Multiplex Reference > System Procedures > sp_iqdbsize Procedure*.

Example

Displays size information for the database iqdemo:

```
sp_iqdbsize

Database

PhysicalBlocks KBytes Pages CompressedPages NBlocks CatalogBlocks
=====
/system1/sybase/IQ-15_4/demo/iqdemo.db
                1280    522    688                257    1119                18
```

sp_iqdbspace Procedure

Displays detailed information about each IQ dbspace.

Syntax

```
sp_iqdbspace [ dbspace-name ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

You can use the information from **sp_iqdbspace** to determine whether data must be moved, and for data that has been moved, whether the old versions have been deallocated.

sp_iqdbspace displays this information:

Table 123. sp_iqdbspace columns

Column name	Description
DBSpaceName	Name of the dbspace as specified in the CREATE DBSPACE statement. Dbspace names are always case-insensitive, regardless of the CREATE DATABASE...CASE IGNORE or CASE RESPECT specification.

Column name	Description
DBSpaceType	Type of the dbspace (MAIN, SHARED_TEMP, or TEMPORARY).
Writable	T (writable) or F (not writable).
Online	T (online) or F (offline).
Usage	Percent of dbspace currently in use by all files in the dbspace.
TotalSize	Total size of all files in the dbspace in the units B (bytes), K (kilobytes), M (megabytes), G (gigabytes), T (terabytes), or P (petabytes).
Reserve	Total reserved space that can be added to all files in the dbspace.
NumFiles	Number of files in the dbspace.
NumRWFiles	Number of read/write files in the dbspace.
Stripingon	T (On) or F (Off).
StripeSize	Amount of data written to the dbspace before moving to the next dbspace, if disk striping is on.
BlkTypes	Space used by both user data and internal system structures.
OkToDrop	'Y' indicates the dbspace can be dropped; otherwise 'N'.

Values of the BlkTypes block type identifiers:

Table 124. sp_iqdbspace block types

Identifier	Block type
A	Active Version
B	Backup Structures
C	Checkpoint Log
D	Database Identity
F	Free list
G	Global Free list Manager
H	Header Blocks of the free list
I	Index advice storage
M	Multiplex CM*
O	Old Version

Identifier	Block type
T	Table use
U	Index use
N	Column use
X	Drop at checkpoint

*The multiplex commit identity block (actually 128 blocks) exists in all IQ databases, even though it is not used by simplex databases.

For information on multiplex capability, see *Using Sybase IQ Multiplex > Multiplex Reference > System Procedures > sp_iqdbspace Procedure*.

Example

The following output displays information about dbspaces.

```
sp_iqdbspace;
```

Note: The following example shows objects in the iqdemo database to better illustrate output. Note that iqdemo includes a sample user dbspace named iq_main that may not be present in your own databases.

DBSpaceName	DBSpaceType	Writable	Online	Usage	Total Size	Reserve	Num Files	Num RWF Files	Striping	Stripe Size	Blk Types	Ok To Drop
IQ_MAIN	MAIN	T	T	55	75M	200M	1	1	T	1K	1H, 5169A, 190	N
IQ__SYS-TEM_MAIN	MAIN	T	T	21	300M	50M	1	1	F	8K	1H, 7648F, 32D, 128M	N
IQ_SYS-TEM_TEMP	TEMPORARY	T	T	1	100M	50M	1	1	F	8K	1H, 64F, 32A	N

See also

- *sp_iqindexinfo Procedure* on page 440
- *sp_iqdbspaceinfo Procedure* on page 405
- *sp_iqspaceinfo Procedure* on page 486

sp_iqdbspaceinfo Procedure

Displays the size of each object and subobject used in the specified table or join index.

Syntax

```
sp_iqdbspaceinfo [ dbspace-name ] [ , owner_name ] [ ,  
object_name ] [ , object-type ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Parameter	Description
<i>dbspace_name</i>	If specified, sp_iqdbspaceinfo displays one line for each table that has any component in the specified dbspace. Otherwise, the procedure shows information for all dbspaces in the database.
<i>owner_name</i>	Owner of the object. If specified, sp_iqdbspaceinfo displays output only for tables and join indexes with the specified owner. If not specified, sp_iqdbspaceinfo displays information on tables and join indexes for all users in the database.
<i>object_name</i>	Name of the table or join index . If not specified, sp_iqdbspaceinfo displays information on all tables and join indexes in the database.
<i>object_type</i>	Valid object types are table (for table information) or joinindex (for join index information). If unspecified, object type defaults to table .

All parameters are optional and any parameter may be supplied independent of another parameter's value.

The **sp_iqdbspaceinfo** stored procedure supports wildcard characters for interpreting *dbspace_name*, *object_name*, and *owner_name*. It shows information for all dbspaces that match the given pattern in the same way the **LIKE** clause matches patterns inside queries.

Description

sp_iqdbspaceinfo shows the DBA the amount of space used by objects that reside on each dbspace. The DBA can use this information to determine which objects must be relocated before a dbspace can be dropped. The sub-object columns display sizes reported in integer quantities followed by the suffix B, K, M, G, T, or P, representing bytes, kilobytes, megabytes, gigabytes, terabytes, and petabytes, respectively.

For tables, **sp_iqdbspaceinfo** displays sub-object sizing information for all sub-objects (using integer quantities with the suffix B, K, M, G, T, or P). For join indexes, the procedure displays sizing information for the join index and all of its associated sub-objects. The output is sorted by *dbspace_name*, *object_name*, and *owner_name*.

Table 125. sp_iqdbspaceinfo columns

Column name	Description
dbspace_name	Name of the dbspace.
object_type	Type of the object (table or joinindex only).
owner	Name of the owner of the object.
object_name	Name of the object (of type tables and join indexes only) located on the dbspace.
object_id	Global object ID of the object.
id	Table id or join-index ID of the object.
columns	Size of column storage space on the given dbspace.
indexes	Size of index storage space on the given dbspace. Does not use system-generated indexes (for example, HG indexes in unique constraints or FP indexes).
metadata	Size of storage space for metadata objects on the given dbspace.
primary_key	Size of storage space for primary key related objects on the given dbspace.
unique_constraint	Size of storage space for unique constraint-related objects on the given dbspace.
foreign_key	Size of storage space for foreign-key-related objects on the given dbspace.
dbspace_online	Indicates if the dbspace is online (Y) or offline (N).

Note: If you run **sp_iqdbspaceinfo** against a server you have started with the **-r** switch (read-only), you see the error **Msg 13768, Level 14, State 0: SQL Anywhere Error -757: Modifications not permitted for read-only database**. This behavior is expected. The error does not occur on other stored procedures

such as `sp_iqdbspace`, `sp_iqfile`, `sp_iqdbspaceobjectinfo` or `sp_iqobjectinfo`.

You can use this procedure in a multiplex environment. See *Using Sybase IQ Multiplex > Multiplex Reference > System Procedures > sp_iqdbspaceinfo Procedure*.

Examples

Displays the size of all objects and sub-objects in all tables in all dbspaces in the database:

Note: The following three examples show objects in the `iqdemo` database to better illustrate output. Note that `iqdemo` includes a sample user dbspace named `iq_main` that may not be present in your own databases.

`sp_iqdbspaceinfo`

dbspace_name	object_type	owner	object_name	object_id	id	columns
iq_main	table	DBA	empl	3689	741	96K
iq_main	table	DBA	iq_dummy	3686	740	24K
iq_main	table	DBA	sale	3698	742	96K
iq_main	table	GROUP0	Contacts	3538		732
iq_main	table	GROUP0	Customers	3515		731
iq_main	table	GROUP0	Departments	3632	738	72K
iq_main	table	GROUP0	Employees	3641		739
iq_main	table	GROUP0	FinancialCodes	3612		736
iq_main	table	GROUP0	FinancialData	3621	737	96K
iq_main	table	GROUP0	Products	3593	735	272K
iq_main	table	GROUP0	SalesOrderItems	3580		734
iq_main	table	GROUP0	SalesOrders	3565		733
indexes	metadata	primary_key	unique_constraint	foreign_key		dbsp
ace_online						
0B	1.37M	0B	0B	0B		Y
0B	464K	0B	0B	0B		Y
0B	1.22M	0B	0B	0B		Y
0B	5.45M	24K	0B	48K		Y
48K	4.63M	24K	0B	0B		Y
0B	1.78M	24K	0B	48K		Y
0B	8.03M	24K	0B	48K		Y
0B	1.53M	24K	0B	0B		Y
0B	2.19M	24K	0B	48K		Y
192K	4.67M	24K	0B	0B		Y
0B	2.7M	24K	0B	104K		Y
0B	3.35M	24K	0B	144K		Y

System Procedures

Displays the size of all objects and sub-objects owned by a specified user in a specified dbspace in the database:

```
sp_iqdbspaceinfo iq_main,GROUPO
```

dbspace_name	object_type	owner	object_name	object_id	id
columns					
iq_main	table	GROUPO	Contacts	3538	732
288K					
iq_main	table	GROUPO	Customers	3515	731
240K					
iq_main	table	GROUPO	Departments	3632	738
72K					72K
iq_main	table	GROUPO	Employees	3641	739
408K					
iq_main	table	GROUPO	FinancialCodes	3612	736
72K					
iq_main	table	GROUPO	FinancialData	3621	737
96K					96K
iq_main	table	GROUPO	Products	3593	735
272K					
iq_main	table	GROUPO	SalesOrderItems	3580	734
120K					
iq_main	table	GROUPO	SalesOrders	3565	733
144K					
indexes	metadata	primary_key	unique_constraint	foreign_key	dbspace_online
0B	5.45M	24K	0B	48K	Y
48K	4.63M	24K	0B	0B	Y
0B	1.78M	24K	0B	48K	Y
0B	8.03M	24K	0B	48K	Y
0B	1.53M	24K	0B	0B	Y
0B	2.19M	24K	0B	48K	Y
192K	4.67M	24K	0B	0B	Y
0B	2.7M	24K	0B	104K	Y
0B	3.35M	24K	0B	144K	Y

Displays the size of a specified object and its sub-objects owned by a specified user in a specified dbspace in the database:

```
sp_iqdbspaceinfo iq_main,GROUPO,Departments
```

dbspace_name	object_type	owner	object_name	object_id	id
columns					
iq_main	table	GROUPO	Departments	3632	738
72K					72K
indexes	metadata	primary_key	unique_constraint	foreign_key	dbspace_online
0B	1.78M	24K	0B	48K	Y

See also

- *sp_iqindexinfo Procedure* on page 440
- *sp_iqdbspace Procedure* on page 402
- *sp_iqspaceinfo Procedure* on page 486

sp_iqdbspaceobjectinfo Procedure

Lists objects of type table and join index and their sub-objects (columns, indexes, metadata, primary keys, unique constraints, foreign keys, and partitions) for a given dbspace.

Syntax

```
sp_iqdbspaceobjectinfo [ dbspace-name ] [ , owner_name ] [ ,  
object_name ] [ , object-type ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

All parameters are optional and any parameter may be supplied independent of the value of other parameters.

Table 126. Parameters

Parameter	Description
dbspace-name	If specified, sp_iqdbspaceobjectinfo displays output only for the specified dbspace. Otherwise, it shows information for all dbspaces in the database.
owner-name	Owner of the object. If specified, sp_iqdbspaceobjectinfo displays output only for tables and join indexes with the specified owner. If not specified, sp_iqdbspaceobjectinfo displays information for tables and join indexes for all users in the database.
object-name	Name of the table or join index. If not specified, sp_iqdbspaceobjectinfo displays information for all tables and join indexes in the database.
object-type	Valid object types are table (for table information) or joinindex (for join index information). If unspecified, object type defaults to table .

The **sp_iqdbspaceobjectinfo** stored procedure supports wildcard characters for interpreting *dbspace_name*, *object_name*, and *owner_name*. It displays information for all dbspaces that match the given pattern in the same way as the **LIKE** clause matches patterns inside queries.

Description

For tables, **sp_iqdbspaceobjectinfo** displays summary information for all associated sub-objects. For join indexes, it displays sizing information for the join index and for all of its

associated sub-objects. The output of the stored procedure is sorted by `dbspace_name`, `owner` and `object_name`.

`sp_iqdbspaceobjectinfo` displays the following information, based on the input parameter values:

Table 127. `sp_iqdbspaceobjectinfo` columns

Column name	Description
<code>dbspace_name</code>	Name of the dbspace.
<code>dbspace_id</code>	Identifier of the dbspace.
<code>object_type</code>	Table or join index.
<code>owner</code>	Name of the owner of the object.
<code>object_name</code>	Name of the object (of type tables and join indexes only) located on the dbspace.
<code>object_id</code>	Global object id of the object.
<code>id</code>	Table id or join-index id of the object.
<code>columns</code>	Number of table columns which are located on the given dbspace. If a column or one of the column-partitions is located on a dbspace, it is counted to be present on that dbspace. The result is displayed in the form <code>n/N</code> (n out of total N columns of the table are on the given dbspace).
<code>indexes</code>	Number of user defined indexes on the table which are located on the given dbspace. Displayed in the form <code>n/N</code> (n out of total N indexes on the table are on the given dbspace). This does not contain indexes which are system generated, such as FP indexes and HG indexes in the case of unique constraints.
<code>metadata</code>	Boolean field (Y/N) to denote if the metadata information of the sub-object is also located on this dbspace.
<code>primary_key</code>	Boolean field (1/0) to denote if the primary key of the table, if any, is located on this dbspace.
<code>unique_constraint</code>	Number of unique constraints on the table which are located on the given dbspace. Displayed in the form <code>n/N</code> (n out of total N unique constraints on the table are in the given dbspace).
<code>foreign_key</code>	Number of foreign_keys on the table which are located on the given dbspace. Displayed in the form <code>n/N</code> (n out of total N foreign keys on the table are in the given dbspace).

Column name	Description
partitions	Number of partitions of the table which are located on the given dbspace. Displayed in the form n/N (n out of total N partitions of the table are in the given dbspace).

Examples

Note: The following two examples show objects in the `iqdemo` database to better illustrate output. Note that `iqdemo` includes a sample user dbspace named `iq_main` that may not be present in your own databases.

Displays information about a specific dbspace in the database:

```
sp_iqdbspaceobjectinfo iq_main
```

dbspace_name	dbspace_id	object_type	owner	object_name	object_id
iq_main	16387	table	DBA	emp1	3689
741	4/4				
iq_main	16387	table	DBA	iq_dummy	3686
740	1/1				
iq_main	16387	table	DBA	sale	3698
742	4/4				
iq_main	16387	table	GROUPO	Contacts	3538
732	12/12				
iq_main	16387	table	GROUPO	Customers	3515
731	10/10				
iq_main	16387	table	GROUPO	Departments	3632
738	3/3				
iq_main	16387	table	GROUPO	Employees	3641
739	21/21				
iq_main	16387	table	GROUPO	FinancialCodes	3612
736	3/3				
iq_main	16387	table	GROUPO	FinancialData	3621
737	4/4				
iq_main	16387	table	GROUPO	Products	3593
735	8/8				
iq_main	16387	table	GROUPO	SalesOrderItems	3580
734	5/5				
iq_main	16387	table	GROUPO	SalesOrders	3565
733	6/6				
indexes	metadata	primary_key	unique_constraint	foreign_key	partitions
0/0	Y	0	0/0	0/0	0/0
0/0	Y	0	0/0	0/0	0/0
0/0	Y	0	0/0	0/0	0/0
0/0	Y	1	0/0	1/1	0/0
1/1	Y	1	0/0	0/0	0/0
0/0	Y	1	0/0	1/1	0/0
0/0	Y	1	0/0	1/1	0/0
0/0	Y	1	0/0	0/0	0/0
0/0	Y	1	0/0	1/1	0/0

System Procedures

4/4	Y	1	0/0	0/0	0/0
0/0	Y	1	0/0	2/2	0/0
0/0	Y	1	0/0	3/3	0/0

Displays information about the objects owned by a specific user in a specific dbspace in the database:

```
sp_iqdbspaceobjectinfo iq_main,GROUPO
```

dbspace_name	dbspace_id	object_type	owner	object_name	object_id
iq_main	16387	table	GROUPO	Contacts	3538
	732	2/12			
iq_main	16387	table	GROUPO	Customers	3515
	731	10/10			
iq_main	16387	table	GROUPO	Departments	3632
	738	3/3			
iq_main	16387	table	GROUPO	Employees	3641
	739	21/21			
iq_main	16387	table	GROUPO	FinancialCodes	3612
	736	3/3			
iq_main	16387	table	GROUPO	FinancialData	3621
	737	4/4			
iq_main	16387	table	GROUPO	Products	3593
	735	8/8			
iq_main	16387	table	GROUPO	SalesOrderItems	3580
	734	5/5			
iq_main	16387	table	GROUPO	SalesOrders	3565
	733	6/6			

indexes	metadata	primary_key	unique_constraint	foreign_key	partitions
0/0	Y	1	0/0	1/1	0/0
1/1	Y	1	0/0	0/0	0/0
0/0	Y	1	0/0	1/1	0/0
0/0	Y	1	0/0	1/1	0/0
0/0	Y	1	0/0	0/0	0/0
0/0	Y	1	0/0	1/1	0/0
4/4	Y	1	0/0	0/0	0/0
0/0	Y	1	0/0	2/2	0/0
0/0	Y	1	0/0	3/3	0/0

Uses **sp_iqdbspaceobjectinfo** to construct commands that can be used to move objects. In this example, the commands move all tables on `dbspace_x` to `dbspace_y`.

```
SELECT 'ALTER TABLE ' || owner || '.' ||
object_name || ' MOVE TO dbspace_y;'
FROM sp_iqdbspaceobjectinfo()
WHERE object_type = 'table' AND
dbspace_name = 'dbspace_x';
```

The following **ALTER TABLE** commands are the result:

```
ALTER TABLE DBA.dt1 MOVE TO dbspace_y;
ALTER TABLE DBA.dt2 MOVE TO dbspace_y;
ALTER TABLE DBA.dt3 MOVE TO dbspace_y;
```

sp_iqdbstatistics Procedure

Reports results of the most recent **sp_iqcheckdb**.

Syntax

```
sp_iqdbstatistics
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Displays the database statistics collected by the most recent execution of **sp_iqcheckdb**.

For more information on the use of **sp_iqcheckdb** and the interpretation of the **sp_iqcheckdb** output, see *System Administration Guide: Volume 1 > System Recovery and Database Repair*.

Example

The following example shows the output from **sp_iqdbstatistics**. For this example, the most recent execution of **sp_iqcheckdb** was the command **sp_iqcheckdb 'allocation database'**.

DB Statistics	Value	Flags
===== =====		
=====		
DBCC Allocation Mode Report		
===== =====		
=====		
** DBCC Status	Errors Detected	*****
DBCC Work units Dispatched	163	
DBCC Work units Completed	163	
===== =====		
=====		
Allocation Summary		
===== =====		
=====		
Blocks Total	8192	
Blocks in Current Version	4954	
Blocks in All Versions	4954	
Blocks in Use	4986	
% Blocks in Use	60	
** Blocks Leaked	32	*****
===== =====		
=====		
Allocation Statistics		
===== =====		
=====		
Blocks Created in Current TXN	382	
Blocks To Drop in Current TXN	382	

System Procedures

```
Marked Logical Blocks      | 8064
Marked Physical Blocks    | 4954
Marked Pages               | 504
Blocks in Freelist        | 126553
Imaginary Blocks          | 121567
Highest PBN in Use        | 5432
** 1st Unowned PBN        | 452          | *****
Total Free Blocks         | 3206
Usable Free Blocks        | 3125
% Free Space Fragmented   | 2
Max Blocks Per Page       | 16
1 Block Page Count        | 97
3 Block Page Count        | 153
4 Block Page Count        | 14
...
9 Block Hole Count        | 2
16 Block Hole Count       | 194

Database Objects Checked  | 1
B-Array Count             | 1
Blockmap Identity Count   | 1
=====
====
Connection Statistics     |
=====
====
```

sp_iqdroplogin Procedure

Drops a Sybase IQ user account.

Syntax1

```
call sp_iqdroplogin ('userid')
```

Syntax2

```
sp_iqdroplogin 'userid'
```

Syntax3

```
sp_iqdroplogin userid
```

Syntax4

```
sp_iqdroplogin ('userid')
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

*Usage***Table 128. Parameters**

Parameter	Description
userid	User ID of the user to drop.

Description

sp_iqdroplogin drops the specified user.

Examples

The following stored procedure calls remove the user `rose`.

```
sp_iqdroplogin 'rose'
```

```
sp_iqdroplogin rose
```

```
call sp_iqdroplogin ('rose')
```

See also

- *sp_iqaddlogin Procedure* on page 363

sp_iqemptyfile Procedure

Empties a dbfile and moves the objects in the dbfile to another available read-write dbfile in the same dbspace.

Syntax

```
sp_iqemptyfile ( logical-file--name )
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

sp_iqemptyfile empties a dbfile. The dbspace must be read-only before you can execute the **sp_iqemptyfile** procedure. The procedure moves the objects in the file to another available read-write dbfile in the same dbspace. If there is no other read-write dbfile available, then Sybase IQ displays an error message.

Note: In a multiplex environment, you can run **sp_iqemptyfile** only on the coordinator. There must be one read-write dbspace available for the procedure to succeed.

Example

The following empties dbfile **dbfile1**:

```
sp_iqemptyfile 'dbfile1'
```

sp_iqestjoin Procedure

Estimates the space needed to create join indexes for the tables you specify.

Syntax

```
sp_iqestjoin ( table1_name, table1_row_#, table2_name,
table2_row_#, relation, iq_page_size )
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Returns the amount of space a join index uses based on the tables being joined. This procedure assumes that the database was created with the default block size for the specified IQ page size (or else the estimate is incorrect).

If you specify unqualified table names, then ensure that you are the owner of the tables being joined. If you are not the table owner, then provide a qualified table name for each table, such as 'owner.tablename'.

Table 129. sp_iqestjoin parameters

Name	Datatype	Description
<i>table1_name</i>	char (256)	Name of the first table in the join.
<i>table1_row_#</i>	int	Number of rows in the first table that participates in the join.
<i>table2_name</i>	char (256)	Name of the second table in the join.
<i>table2_row_#</i>	int	Number of rows in the second table that participates in the join.
<i>relation</i>	char (9)	Type of join, which can be “one>>many” or “one>>one” (do not leave any spaces between the words and the operator). The default is “one>>many”.
<i>iq_page_size</i>	smallint	The page size defined for the IQ segment of the database (must be a power of 2 between 1024 and 524288; the default is 131072).

Example

```
call sp_iqestjoin ( 'Customers', 1500000, 'SalesOrders', 1500000,
'one>>many', 65536 )
```

Cases	Indexsize	Create time	Msg
Table1:Customers			
Rows: 1500000			
Columns:			
8			
Width:			
223			
Table2: SalesOrders			
Rows: 15000000			
Columns:			
9			
Width:			
134			
IQpagesize:			
65536			
Min Case	48001024	3h0m/CPU	
Max Case	95449088	9h6m/CPU	
Avg Case	70496256	5h53m/CPU	

sp_iqestdbspaces Procedure

Estimates the number and size of dbspaces needed for a given total index size.

Syntax

```
sp_iqestdbspaces ( db_size_in_bytes, iq_page_size,  
min_#_of_bytes, max_#_of_bytes )
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Displays information about the number and size of dbspace segments based on the size of the database, the IQ page size, and the range of bytes per dbspace segment. This procedure

assumes that the database was created with the default block size for the specified IQ page size (or else the estimate is incorrect).

Table 130. sp_iquestdbspaces parameters

Name	Datatype	Description
<i>db_size_in_bytes</i>	decimal(16)	Size of the database in bytes.
<i>iq_page_size</i>	smallint	The page size defined for the IQ segment of the database (must be a power of 2 between 65536 and 524288; the default is 131072).
<i>min_#_of_bytes</i>	int	The minimum number of bytes per dbspace segment. The default is 20,000,000 (20MB).
<i>max_#_of_bytes</i>	int	The maximum number of bytes per dbspace segment. The default is 2,146,304,000 (2.146GB).

Usage

sp_iquestdbspaces displays four types of recommendations, depending on how much of the data is unique:

Table 131. Recommendations

Recommendation	Description
min	If there is little variation in data, you can choose to create only the dbspace segments of the sizes recommended as min . These recommendations reflect the best possible compression on data with the least possible variation.
avg	If your data has an average amount of variation, create the dbspace segments recommended as min , plus additional segments of the sizes recommended as avg .
max	If your data has a high degree of variation (many unique values), create the dbspace segments recommended as min , avg , and max .

Recommendation	Description
spare	If you are uncertain about the number of unique values in your data, create the dbspace segments recommended as min , avg , max , and spare . You can always delete unused segments after loading your data, but creating too few can cost you some time.

sp_iquestdbspaces Procedure Example

Use the example as a reference for **sp_iquestdbspaces** usage.

```
sp_iquestdbspaces 12000000000, 65536, 500000000, 2146304000
```

dbspace files	Type	Size	Msg
1	min	2146304000	
2	min	2146304000	
3	min	507392000	
4	avg	2146304000	
5	max	2053697536	
6	spare	1200001024	

This example estimates the size and number of dbspace segments needed for a 12GB database. Sybase IQ recommends that you create a minimum of 3 segments (listed as **min**) for the best compression, if you expect little uniqueness in the data. If the data has an average amount of variation, 1 more segment (listed as **avg**) should be created. Data with a lot of variation (many unique values, requiring extensive indexing), may require 1 more segment (listed as **max**). You can ensure that your initial load succeeds by creating a spare segment of 1200001024 bytes. Once you have loaded the database, you can delete any unused dbspace segments.

Using sp_iquestdbspaces With Other System Stored Procedures

You need to run two stored procedures to provide the *db_size_in_bytes* parameter needed by **sp_iquestdbspaces**.

Results of **sp_iquestdbspaces** are only estimates, based on the average size of an index. The actual size depends on the data stored in the tables, particularly on how much variation there is in the data.

Sybase strongly recommends that you create the spare dbspace segments, because you can delete them later if they are unused.

1. Run **sp_iquestjoin** for all the table pairs you expect to join frequently.
2. Select one of the suggested index sizes for each pair of tables.

System Procedures

3. Total the index sizes you selected for all tables.
4. Run **sp_iqestspace** for all tables.
5. Total all of the RAW DATA index sizes returned by **sp_iqestspace**.
6. Add the total from step 3 to the total from step 5 to determine total index size.
7. Use the total index size calculated in step 6 as the *db_size_in_bytes* parameter in **sp_iqestdbspaces**.

sp_iqestspace Procedure

Estimates the amount of space needed to create an index based on the number of rows in the table.

Syntax

```
sp_iqestspace ( table_name, #_of_rows, iq_page_size )
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Displays the amount of space that a database requires based on the number of rows in the underlying database tables and on the database IQ page size. This procedure assumes that the database was created with the default block size for the specified IQ page size (or else the estimate is incorrect). This table lists the **sp_iqestspace** parameters.

Table 132. sp_iqestspace parameters

Name	Datatype	Description
<i>table_name</i>	char (256)	Name of the table
<i>#_of_rows</i>	int	Number of rows in the table
<i>iq_page_size</i>	smallint	The page size defined for the IQ segment of the database (must be a power of 2 between 65536 and 524288; the default is 131072)

sp_iqevent Procedure

Displays information about system and user-defined events.

Syntax

```
sp_iqevent [ event-name ], [ event-owner ], [ event-type ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

*Usage***Table 133. Parameter**

Parameter	Description
event-name	The name of the event.
event-owner	The owner of the event.
event-type	The type of event. Allowed values are: <ul style="list-style-type: none"> • SYSTEM: displays information about system events (events owned by user SYS or dbo) only • ALL: displays information about user and system events • Any other value: displays information about user events

The `sp_iqevent` procedure can be invoked without any parameters. If no parameters are specified, only information about user events (events not owned by dbo or SYS) is displayed by default.

If you do not specify either of the first two parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. For example, `sp_iqevent NULL, NULL, SYSTEM` and `sp_iqevent NULL, user1`.

Table 134. sp_iqevent usage examples

Syntax	Output
<code>sp_iqevent</code>	Displays information about all user events in the database
<code>sp_iqevent e1</code>	Displays information about the event e1
<code>sp_iqevent non_existing_event</code>	No rows returned, as the event <code>non_existing_event</code> does not exist
<code>sp_iqevent NULL, DBA</code>	Displays information about all events owned by DBA
<code>sp_iqevent e1, DBA</code>	Displays information about the event e1 owned by DBA

Syntax	Output
sp_iqevent ev_iqbegintxn	ev_iqbegintxn is a system-defined event. If there is no user-defined event also named ev_iqbegintxn, no rows are returned. (By default only user-defined events are returned.)
sp_iqevent ev_iqbegintxn, dbo	No rows returned, as the event ev_iqbegintxn is not a user event (by default only user events returned)
sp_iqevent NULL, NULL, SYSTEM	Displays information about all system events (owned by dbo or SYS)
sp_iqevent ev_iqbegintxn, NULL, SYSTEM	Displays information about the system event ev_iqbegintxn
sp_iqevent ev_iqbegintxn, dbo, ALL	Displays information about the system event ev_iqbegintxn owned by dbo

Description

The **sp_iqevent** stored event displays information about events in a database. If you specify one or more parameters, the result is filtered by the specified parameters. For example, if *event-name* is specified, only information about the specified event is displayed. If *event-owner* is specified, **sp_iqevent** only returns information about events owned by the specified owner. If no parameters are specified, **sp_iqevent** displays information about all the user events in the database.

The **sp_iqevent** procedure returns information in the following columns:

Table 135. sp_iqevent columns

Column name	Description
event_name	The name of the event
event_owner	The owner of the event
event_type	For system events, the event type as listed in the SYSEVENTTYPE system table
enabled	Indicates whether or not the event is allowed to fire (Y/N)
action	The event handler definition
condition	The WHERE condition used to control firing of the event handler

Column name	Description
location	The location where the event is allowed to fire: <ul style="list-style-type: none"> • C = consolidated • R = remote • A = all
remarks	A comment string

Examples

Display information about the user-defined event e1:

```
sp_iquevent e1
```

event_name	event_owner	event_type	enabled	action
e1	DBA	(NULL)	Y	(NULL)
condition	location	remarks		
(NULL)	A	(NULL)		

Display information about all system events:

```
sp_iquevent NULL, NULL, SYSTEM
```

event_name	event_owner	event_type	enabled	action
ev_iqbegintxn	dbo	IQTLVAvailable	Y	begin call
ev_iqmpxcompact	dbo	(NULL)	N	dbo.sp_iqlog... begin Declare _Catalog...
condition	location	remarks		
(NULL)	A	(NULL)		
(NULL)	A	(NULL)		

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqpkeys Procedure* on page 467
- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iqfile Procedure

Displays detailed information about each dbfile in a dbspace.

Syntax

```
sp_iqfile [ dbspace-name ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

sp_iqfile displays the usage, properties, and types of data in each dbfile in a dbspace. You can use this information to determine whether data must be moved, and for data that has been moved, whether the old versions have been deallocated.

sp_iqfile displays the following information:

Table 136. sp_iqfile columns

Column name	Description
DBSpaceName	Name of the dbspace as specified in the CREATE DBSPACE statement. Dbspace names are always case-insensitive, regardless of the CREATE DATABASE...CASE IGNORE or CASE RESPECT specification.
DBFileName	Logical file name.
Path	Location of the physical file or raw partition.
SegmentType	Type of dbspace (MAIN or TEMPORARY).
RWMode	Mode of the dbspace: read-write (RW) or read-only (RO).
Online	T (online) or F (offline).
Usage	Percent of dbspace currently in use by this file in the dbspace.
DBFileSize	Current size of the file or raw partition. For a raw partition, this size value can be less than the physical size.
Reserve	Reserved space that can be added to this file in the dbspace.
StripeSize	Amount of data written to the file before moving to the next file, if disk striping is on.
BlkTypes	Space used by both user data and internal system structures.
FirstBlk	First IQ block number assigned to the file.

Column name	Description
LastBlk	Last IQ block number assigned to the file.
OkToDrop	'Y' indicates the file can be dropped; otherwise 'N'.

The values of the block type identifiers:

Table 137. sp_iqfile block types

Identifier	Block Type
A	Active Version
B	Backup Structures
C	Checkpoint Log
D	Database Identity
F	Free list
G	Global Free list Manager
H	Header Blocks of the Free List
I	Index Advice Storage
M	Multiplex CM*
O	Old Version
T	Table Use
U	Index Use
N	Column Use
X	Drop at Checkpoint

*The multiplex commit identity block (actually 128 blocks) exists in all IQ databases, even though it is not used by simplex databases.

Example

Displays information about the files in the dbspaces:

```
sp_iqfile;
```

```
sp_iqfile;
DBSpaceName, DBFileName, Path, SegmentType, RWMode, Online,
Usage, DBFileSize, Reserve, StripeSize, BlkTypes, FirstBlk,
LastBlk, OkToDrop
```

System Procedures

```
'IQ_SYSTEM_MAIN','IQ_SYSTEM_MAIN','/sun1-c1/users/smith/mpx/m/
mpx_db.iq','MAIN','RW','T','21','
2.92G','0B','1K','1H',76768F,32D,19A,1850,128M,34B,32C'
,1,384000,'N'

'mpx_main1','mpx_main1','/sun1-c1/users/smith/mpx/m/
mpx_main1.iq','MAIN','RW','T','1'
,'100M','0B','1K','1H',1045440,1058239,'N'

'IQ_SHARED_TEMP','sharedfile1_bcp','/sun1-c1/users/smith/mpx/m/
f1','SHARED_TEMP','RO','T','0',
'50M','0B','1K','1H',1,6400,'N'

'IQ_SHARED_TEMP','sharedfile2_bcp','/sun1-c1/users/smith/mpx/m/
f2','SHARED_TEMP','RO','T','0',
'50M','0B','1K','1H',1045440,1051839,'N'

'IQ_SYSTEM_TEMP','IQ_SYSTEM_TEMP','/sun1-c1/users/smith/mpx/m/
mpx_db.iqtmp','TEMPORARY','RW',
'T','1','2.92G','0B','1K','1H',64F,33A',1,384000,'N'
```

sp_iqhelp Procedure

Displays information about system and user-defined objects and data types.

Syntax

```
sp_iqhelp [ obj-name ], [ obj-owner ], [ obj-category ], [ obj-type ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 138. Parameters

Parameter	Description
obj-name	The name of the object.
obj-owner	The owner of the object.

Parameter	Description
obj-category	<p>An optional parameter that specifies the category of the object.</p> <p>Columns, constraints, and indexes are associated with tables and cannot be queried directly. When a table is queried, the information about columns, indexes, and constraints associated with that table is displayed.</p> <p>If the specified object category is not one of the allowed values, an "Invalid object category" error is returned.</p>
obj-type	<p>The type of object. Allowed values are:</p> <ul style="list-style-type: none"> • SYSTEM: displays information about system objects (objects owned by user SYS or dbo) only • ALL: displays information about all objects <p>By default, only information about non-system objects is displayed. If the specified object type is not SYSTEM or ALL, an "Invalid object type" error is returned.</p>

Table 139. sp_iqhelp obj-category parameter values

<i>object-type</i> parameter	Specifies
"table"	The object is a base table
"view"	The object is an view
"procedure"	The object is a stored procedure or function
"event"	The object is an event
"datatype"	The object is a system or user-defined data type

The **sp_iqhelp** procedure can be invoked without any parameters. If no parameters are specified, **sp_iqhelp** displays information about all independent objects in the database, that is, base tables, views, stored procedures, functions, events, and data types.

If you do not specify any of the first three parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. For example, `sp_iqhelp NULL, NULL, NULL, SYSTEM` and `sp_iqhelp NULL, user1, "table"`.

Enclose the *obj-category* parameter in single or double quotes., except when NULL.

If **sp_iqhelp** does not find an object in the database that satisfies the specified description, the error "No object found for the given description" is returned.

Table 140. sp_iqhelp usage examples

Syntax	Output
sp_iqhelp	Displays summary information about all user-defined tables, views, procedures, events, and data types in the database
sp_iqhelp t1, u1, "table"	Displays information about table t1 owned by user u1 and the columns, indexes, and constraints associated with t1
sp_iqhelp NULL, u1, "view"	Displays information about view v1 owned by user u1 and the columns associated with v1
sp_iqhelp sp2	Displays information about the procedure sp2 and the parameters of sp2
sp_iqhelp e1	Displays information about the event e1
sp_iqhelp dt1	Displays information about the data type dt1
sp_iqhelp NULL, NULL, NULL, SYSTEM	Displays summary information about all system objects (owned by dbo or SYS)
sp_iqhelp non_existing_obj	Error "Object 'non_existing_obj' not found" returned, as the object non_existing_obj does not exist
sp_iqhelp NULL, non_existing_user	Error "User 'non_existing_user' not found" returned, as the user non_existing_user does not exist
sp_iqhelp t1, NULL, "apple"	Error "Invalid object category 'apple'" returned, as " apple " is not an allowed value for <i>obj-category</i>
sp_iqhelp t1, NULL, NULL, "USER"	Error "Invalid object type 'USER'" returned, as " USER " is not an allowed value for <i>obj-type</i>

Description

The **sp_iqhelp** stored procedure displays information about system and user-defined objects and data types in an IQ database. Objects supported by **sp_iqhelp** are tables, views, columns, indexes, join indexes, constraints, stored procedures, functions, events, and data types.

If you specify one or more parameters, the result is filtered by the specified parameters. For example, if *obj-name* is specified, only information about the specified object is displayed. If

obj-owner is specified, **sp_iqhelp** returns information only about objects owned by the specified owner. If no parameters are specified, **sp_iqhelp** displays summary information about all user-defined tables, views, procedures, events, and data types in the database.

The **sp_iqhelp** procedure returns either summary or detailed information, depending on whether the specified parameters match multiple objects or a single object. The output columns of **sp_iqhelp** are similar to the columns displayed by the stored procedures **sp_iqtable**, **sp_iqindex**, **sp_iqview**, and **sp_iqconstraint**.

When multiple objects match the specified **sp_iqhelp** parameters, **sp_iqhelp** displays summary information about those objects.

Table 141. sp_iqhelp summary information

Object type	Columns displayed
base table	table_name, table_owner, server_type, location, table_constraints, remarks
view	view_name, view_creator, view_def, server_type, location, remarks
stored procedure	proc_name, proc_creator, proc_defn, replicate, srvid, remarks
function	proc_name, proc_creator, proc_defn, replicate, remarks
event	event_name, event_creator, enabled, location, event_type, action, external_action, condition, remarks
system and user-defined data types	type_name, creator, nulls, width, scale, default, check

When a single object matches the specified **sp_iqhelp** parameters, **sp_iqhelp** displays detailed information about the object.

Table 142. sp_iqhelp detailed information

Object type	Description	Columns
table	Displays information about the specified base table, its columns, indexes, constraints, and join indexes (if the table participates in any join indexes)	<ul style="list-style-type: none"> • Table columns: table_name, table_owner, server_type, location, table_constraints, remarks • Column columns: column_name, domain_name, width, scale, nulls, default, check, pkey, user_type, cardinality, est_cardinality, remarks • Index columns: index_name, column_name, index_type, unique_index, location, remarks • Constraint columns: constraint_name (role), column_name, index_name, constraint_type, foreigntable_name, foreigntable_owner, foreigncolumn_name, foreignindex_name, location • Join index columns: joinindex_name, creator, left_table_name, left_table_owner, left_column_name, join_type, right_table_name, right_table_owner, right_column_name, key_type, valid, remarks
view	Displays information about the specified view and its columns	<ul style="list-style-type: none"> • View columns: view_name, view_creator, view_def, server_type, location, remarks • Column columns: column_name, domain_name, width, scale, nulls, default, check, pkey, user_type, cardinality, est_cardinality, remarks
stored procedure	Displays information about the specified procedure and its parameters	<ul style="list-style-type: none"> • Procedure columns: proc_name, proc_creator, proc_defn, replicate, srvid, remarks • Parameter columns: parameter_name, type, width, scale, default, mode
function	Displays information about the specified function and its parameters	<ul style="list-style-type: none"> • Function columns: proc_name, proc_creator, proc_defn, replicate, srvid, remarks • Parameter columns: parameter_name, type, width, scale, default, mode

Object type	Description	Columns
event	Displays information about the specified event	<ul style="list-style-type: none"> Event columns: event_name, event_creator, enabled, location, event_type, action, external_action, condition, remarks
data type	Displays information about the specified data type	<ul style="list-style-type: none"> Data type columns: type_name, creator, nulls, width, scale, default, check

Note: Procedure definitions (proc-defn) of system procedures are encrypted and hidden from view.

For descriptions of the individual output columns, refer to the related stored procedure. For example, for a description of the table column, see the **sp_iqtable** procedure.

Examples

Display detailed information about the table sale:

```
sp_iqhelp sale
```

```
Table_name Table_owner Server_type Location dbspace_id isPartitioned
table_constraints
=====
sale      DBA          IQ           Main      16387      N

Remarks  table_constraints
===== (NULL) (NULL)

column_name domain_name width scale nulls default cardinality
=====
prod_id     integer    4      0      Y      (NULL)  0
month_num   integer    4      0      Y      (NULL)  0
rep_id      integer    4      0      Y      (NULL)  0
sales       integer    4      0      Y      (NULL)  0

  est_cardinality  isPartitioned  remarks  check
  =====
  0                N              (NULL)  (NULL)
  0                N              (NULL)  (NULL)
  0                N              (NULL)  (NULL)
  0                N              (NULL)  (NULL)

index_name          column_name  index_type  unique_index  location
=====
ASIQ_IDX_T463_C2_FP month_num   FP          N             Main
ASIQ_IDX_T463_C1_FP prod_id     FP          N             Main
ASIQ_IDX_T463_C3_FP rep_id      FP          N             Main
ASIQ_IDX_T463_C4_FP sales       FP          N             Main

remarks
```

```

=====
(NULL)
(NULL)
(NULL)
(NULL)

```

Display detailed information about the procedure `sp_customer_list`:

```

sp_iqhelp sp_customer_list
proc_name      proc_owner      proc_defn
=====      =====      =====
sp_customer_list      DBA      create procedure DBA.sp_customer_list()
                                     result(id integer company_name char(35))
                                     begin
                                     select id company_name from Customers
                                     end

replicate      srvid      remarks
=====      =====      =====
N              (NULL)      (NULL)

parm_name      parm_type      parm_mode      domain_name      width      scale
=====      =====      =====      =====      =====      =====
id              result          out              integer          4          0
company_name    result          out              char              35         0

default
=====
(NULL)

```

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqpkeys Procedure* on page 467
- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iqhelp Compatibility with Adaptive Server Enterprise

The Sybase IQ `sp_iqhelp` stored procedure is similar to the Adaptive Server Enterprise `sp_help` procedure, which displays information about any database object listed in the `SYSOBJECTS` system table and about system and user-defined data types.

Sybase IQ has some architectural differences from Adaptive Server Enterprise in terms of types of objects supported and the namespace of objects. In Adaptive Server Enterprise, all

objects (tables, views, stored procedures, logs, rules, defaults, triggers, check constraints, referential constraints, and temporary objects) are stored in the `SYSOBJECTS` system table and are in the same namespace. The objects supported by Sybase IQ (tables, views, stored procedures, events, primary keys, and unique, check, and referential constraints) are stored in different system tables and are in different namespaces. For example, in Sybase IQ a table can have the same name as an event or a stored procedure.

Because of the architectural differences between Sybase IQ and Adaptive Server Enterprise, the types of objects supported by and the syntax of Sybase IQ `sp_iqhelp` are different from the supported objects and syntax of Adaptive Server Enterprise `sp_help`; however, the type of information about database objects that is displayed by both stored procedures is similar.

sp_iqindex and sp_iqindex_alt Procedures

Lists information about indexes.

Syntax 1

```
sp_iqindex ( [ table_name ], [column_name ], [table_owner ] )
```

Syntax 2

```
sp_iqindex [ table_name='tablename' ],  
[column_name='columnname' ], [table_owner='tableowner' ]
```

Syntax 3

```
sp_iqindex_alt ( [ table_name ], [column_name ], [table_owner ] )
```

Syntax 4

```
sp_iqindex_alt [ table_name='tablename' ],  
[column_name='columnname' ], [table_owner='tableowner' ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 143. Parameters

Parameter	Description
Syntax1	If you do not specify either of the first two parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. For example, <code>sp_iqindex NULL, NULL, DBA</code> and <code>sp_iqindex Departments, NULL, DBA</code> .

Parameter	Description
Syntax2	You can specify the parameters in any order. Enclose them in single quotes.
Syntax 3 and 4	Produces slightly different output when a multicolumn index is present. Allows the same options as Syntax 1 and 2.

Description

Displays information about indexes in the database. Specifying one of the parameters returns the indexes from only that table, column, or tables owned by the specified user. Specifying more than one parameter filters the results by all of the parameters specified. Specifying no parameters returns all indexes for all tables in the database.

Table 144. sp_iqindex and sp_iqindex_alt columns

Column name	Description
table_name	The name of the table
table_owner	The owner of the table
column_name	The name of the column; multiple names can appear in a multicolumn index
index_type	The abbreviated index type (for example, HG , LF)
index_name	The name of the index
unique_index	'U' indicates the index is a unique index; otherwise, 'N'
location	TEMP = IQ temporary store, MAIN = IQ store, SYSTEM = catalog store
remarks	User comments added with the COMMENT statement

sp_iqindex always produces one line per index. **sp_iqindex_alt** produces one line per index per column if there is a multicolumn index.

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqpkeys Procedure* on page 467

- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iqindex and sp_iqindex_alt Procedure Examples

Use the example as a reference for **sp_iqindex** and **sp_iqindex_alt** usage.

The following variations in syntax both return all indexes on columns with the name DepartmentID:

```
call sp_iqindex (NULL,'DepartmentID')
```

```
sp_iqindex column_name='DepartmentID'
```

table_name	table_owner	column_name	index_type	index_name	unique_index	location	dbspace_id	remarks
Departments	GROUPO	DepartmentID	FP	ASIQ_IDX_T201_C1_FP	N	Main	16387	(NULL)
Departments	GROUPO	DepartmentID	HG	ASIQ_IDX_T201_C1_HG	U	Main	16387	(NULL)
Employees	GROUPO	DepartmentID	FP	ASIQ_IDX_T202_C5_FP	N	Main	16387	(NULL)

The following variations in syntax both return all indexes in the table Departments that is owned by table owner GROUPO:

```
sp_iqindex Departments,NULL,GROUPO
```

```
sp_iqindex table_name='Departments',table_owner='DBA'
```

table_name	table_owner	column_name	index_type	index_name	unique_index	location	dbspace_id	remarks
Departments	GROUPO	DepartmentHeadID	FP	ASIQ_IDX_T201_C3_FP	N	Main	16387	(NULL)
Departments	GROUPO	DepartmentID	FP	ASIQ_IDX_T201_C1_FP	N	Main	16387	(NULL)

table_name	table_owner	column_name	index_type	index_name	unique_index	location	dbspace_id	remarks
Departments	GROUPO	DepartmentID	HG	ASIQ_IDX_T201_C1_HG	U	Main	16387	(NULL)
Departments	GROUPO	Department-Name	FP	ASIQ_IDX_T201_C2_FP	N	Main	16387	(NULL)

The following variations in syntax for **sp_iqindex_alt** both return indexes on the table Employees that contain the column City. The index emp_loc is a multicolumn index on the columns City and State. **sp_iqindex_alt** displays one row per column for a multicolumn index.

```
sp_iqindex_alt Employees, City
```

```
sp_iqindex_alt table_name='Employees',
column_name='City'
```

table_name	table_owner	column_name	index_type	index_name	unique_index	dbspace_id	remarks
Employees	GROUPO	City	FP	ASIQ_IDX_T452_C7_FP	N	16387	(NULL)
Employees	GROUPO	City	HG	emp_loc	N	16387	(NULL)
Employees	GROUPO	State	HG	emp_loc	N	16387	(NULL)

The output from **sp_iqindex** for the same table and column is slightly different:

```
sp_iqindex Employees, City
```

```
sp_iqindex table_name='Employee', column_name='City'
```

table_name	table_owner	column_name	index_type	index_name	unique_index	dbspace_id	location	remarks
Employees	GROUPO	City	FP	ASIQ_IDX_T452_C7_FP	N	16387	Main	(NULL)

table_name	table_owner	column_name	index_type	index_name	unique_index	dbspace_id	location	remarks
Employees	GROUPO	City,State	HG	emp_loc	N	16387	Main	(NULL)

sp_iqindexadvice Procedure

Displays stored index advice messages. Optionally clears advice storage.

Syntax

```
sp_iqindexadvice ( [ resetflag ] )
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 145. Parameters

Parameter	Description
resetflag	Lets the caller clear the index advice storage. If <i>resetflag</i> is nonzero, all advice is removed after the last row has been retrieved.

Description

Allows users to query aggregated index advisor messages using SQL. Information can be used to help decide which indexes or schema changes will affect the most queries.

INDEX_ADVISOR columns:

Table 146. sp_iqindexadvice columns

Column name	Description
Advice	Unique advice message
NInst	Number of instances of message
LastDT	Last date/time advice was generated

Examples

Sample output from the **sp_iqindexadvice** procedure:

Table 147. Sample sp_iqindexadvice output

Advice	NInst	LastDT
Add a CMP index on DBA.tb (c2, c3) Predicate: (tb.c2 = tb.c3)	2073	2009-04-07 16:37:31.000
Convert HG index on DBA.tb.c4 to a unique HG	812	2009-04-06 10:01:15.000
Join Key Columns DBA.ta.c1 and DBA.tb.c1 have mismatched data types	911	2009-02-25 20:59:01.000

See also

- *sp_iqcolumnuse Procedure* on page 385
- *sp_iqindexuse Procedure* on page 446
- *sp_iqtableuse Procedure* on page 506
- *sp_iqunusedcolumn Procedure* on page 510
- *sp_iqunusedindex Procedure* on page 511
- *sp_iqunusedtable Procedure* on page 513
- *sp_iqworkmon Procedure* on page 521

sp_iqindexfragmentation Procedure

Reports information about the percentage of page space taken up within the B-trees, garrays, and bitmap structures in Sybase IQ indexes.

For garrays, the fill percentage calculation does not take into account the reserved space within the garray groups, which is controlled by the **GARRAY_FILL_FACTOR_PERCENT** option.

Syntax

```
dbo.sp_iqindexfragmentation ( 'target' )
```

```
target: table table-name | index index-name [...]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 148. Parameter

Parameter	Description
table-name	Target table table-name reports on all nondefault indexes in the named table.

Parameter	Description
index-name	Target index <i>index-name</i> reports on the named index. Each index-name is a qualified index name. You can specify multiple indexes within the table, but you must repeat the index keyword with each index specified.

Example

Reports the internal index fragmentation for nonunique **HG** index `cidhg` in table `Customers`:

```
dbo.sp_iqindexfragmentation ( 'index customers.cidhg ' )
```

Index	Index type	btree node pages	GAR- RAY_FILL_FAC- TOR_PERCENT
dba.customers.cidhg	HG	3	75
SQLCODE	0		
Fill Percent	btree pages	garray pages	bitmap pages
0 - 10%	0	0	0
11 - 20%	0	0	0
21 - 30%	0	0	0
31-40%	0	0	22
41 - 50%	0	0	0
51 - 60%	0	0	10
61 - 70%	2	0	120
71 - 80%	138	3	64
81 - 90%	24	122	14
91 - 100%	18	1	0

According to this output, of the 182 B-tree pages in nonunique **HG** index `cidhg`, 2 are between 61% and 70% full, 138 are 71% to 80% full, 24 are 81% - 90% full, and 18 are 91% - 100% full. Usage for garray and bitmap pages is reported in the same manner. All percentages are truncated to the nearest percentage point. **HG** indexes also display the value of option `GARRAY_FILL_FACTOR_PERCENT`. Those index types that use a B-tree also display the number of node (nonleaf) pages. These are **HG**, **LF**, **WD**, **DATE**, and **DTTM**.

System Procedures

If an error occurred during execution of the stored procedure for this index, the SQLCODE would be nonzero.

See also

- *sp_iqindexmetadata Procedure* on page 442
- *sp_iqindexinfo Procedure* on page 440
- *sp_iqindexsize Procedure* on page 443
- *sp_iqrebuildindex Procedure* on page 475
- *sp_iqrowdensity Procedure* on page 482

sp_iqindexinfo Procedure

Displays the number of blocks used per index per main dbspace for a given object. If the object resides on several dbspaces, **sp_iqindexinfo** returns the space used in all dbspaces, as shown in the example.

Syntax

```
sp_iqindexinfo '{ database  
| [ table table-name | index index-name ] [...] }  
[ resource-percent ]'
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

You can request index information for the entire database, or you can specify any number of table or index parameters. If a table name is specified, **sp_iqindexinfo** returns information on all indexes in the table. If an index name is specified, only the information on that index is returned.

You cannot specify a join index by name. Use the **database** keyword to display join indexes.

If the specified *table-name* or *index-name* is ambiguous or the object cannot be found, an error is returned.

By default in a multiplex database, **sp_iqindexinfo** displays information about the shared IQ store on a secondary node. If individual tables or indexes are specified, the store to display is automatically selected.

resource-percent must be an integer greater than 0. The resources percentage allows you to limit the CPU utilization of the **sp_iqindexinfo** procedure by specifying the percent of total CPUs to use.

You can use this procedure in a multiplex environment. See *Using Sybase IQ Multiplex*.

Description

sp_iqindexinfo shows the DBA on which dbspaces a given object resides. The DBA can use this information to determine which dbspaces must be given **relocate** mode to relocate the object.

The results of **sp_iqindexinfo** are from the point of view of the version seen by the transaction running the command. Blocks used by other versions are not shown.

Table 149. sp_iqindexinfo columns

Column name	Description
Object	Table, index, or join index name
Dbpace_name	Name of the dbspace
ObjSize	Size of data for this object on this dbspace
DBSpPct	Percent of dbspace used by this object
MinBlk	First block used by this object on this dbspace
MaxBlk	Last block used by this object on this dbspace; useful for determining which objects must be relocated before the dbspace is resized to a smaller size

Examples

Displays index information about the table t2:

```
sp_iqindexinfo 'table t2';
```

Object	dbspace_name	Obj-Size	DBSpPct	MinBlk	MaxBlk
t2	IQ_SYS-TEM_MAIN	32K	1	84	107
t2	dbspacedb2	160K	2	1045495	1045556
t2	dbspacedb3	8K	1	2090930	2090930
t2.DBA.ASIQ_IDX_T430_C1_FP	IQ_SYS-TEM_MAIN	136K	2	126	321
t2.DBA.ASIQ_IDX_T430_C1_FP	dbspacedb3	152K	2	2091032	2091053
t2.DBA.t2c1hng	dbspacedb2	136K	2	1045537	1045553

See also

- *sp_iqdbspace Procedure* on page 402

- *sp_iqdbspaceinfo Procedure* on page 405
- *sp_iqspaceinfo Procedure* on page 486
- *sp_iqindexmetadata Procedure* on page 442
- *sp_iqindexfragmentation Procedure* on page 438
- *sp_iqindexsize Procedure* on page 443

sp_iqindexmetadata Procedure

Displays index metadata for a given index.

You can optionally restrict the output to only those indexes on a specified table, and to only those indexes belonging to a specified owner.

Syntax

```
dbo.sp_iqindexmetadata { 'index-name'
[ , 'table-name' [ , 'owner-name' ] ] }
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Specifying a table name limits output to those indexes belonging to that table. Specifying an owner name limits output to indexes owned by that owner. Omitted parameters default to NULL. You can specify only one index per procedure.

Description

The first row of output is the owner name, table name, and index name for the index.

Subsequent rows of output are specific to the type of index specified.

Table 150. sp_iqindexmetadata output rows

Index type	Metadata returned
CMP, DATE, DTTM, TIME	Type, Version
FP	Type, Version, LookupPages, Style, LookupEntries, 1stLookupPage, LargeLOBs, SmallLOBs, IQ Unique, LOB Compression (only if column datatype is LONG VARCHAR or LONG BINARY)
HG	Type, Version, Distinct Keys
HNG	Type, Version, BitsPerBlockmap, NumberOfBits
LD	Type, Version<ld>, Version, Distinct Keys

Index type	Metadata returned
LF	Type, Version, IndexStatus, NumberOfBlockmaps, BitsPerBlockmap, Distinct Keys
WD	Type, Version, KeySize, Delimiters, DelimiterCount, MaxKeyWordLength, PermitEmptyWord

Examples

The following command displays index information about the **HG** index `hg_index_col54`:

```
sp_iqindexmetadata 'hg_index_col54' , 'metal' , 'DBA' ;
```

'DBA',	'metal'	'hg_index_col54'
'Type',	'HG',	"
'Version',	'2',	"
'Distinct Keys',	'0',	"

See also

- *sp_iqindexfragmentation Procedure* on page 438
- *sp_iqindexinfo Procedure* on page 440
- *sp_iqindexsize Procedure* on page 443

sp_iqindexsize Procedure

Gives the size of the specified index.

Syntax

```
sp_iqindexsize [ [ owner.] table.] index_name
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Table 151. sp_iqindexsize columns

Column name	Description
Username	Index owner.
Indexname	Index for which results are returned, including the table name.

Column name	Description
Type	Index type.
Info	Component of the IQ index for which the KBytes, Pages, and Compressed Pages are being reported. The components vary by index type. For example, the default (FP) index includes BARRAY (barray) and Bitmap (bm) components. The Low_Fast (LF) index includes B-tree (bt) and Bitmap (bm) components.
KBytes	Physical object size in KB.
Pages	Number of IQ pages needed to hold the object in memory.
Compressed Pages	Number of IQ pages when the object is compressed (on disk).

Returns the total size of the index in bytes and kilobytes, and an Info column that describes the component of the IQ index for which the KBytes, Pages, and Compressed Pages are reported. The components described vary by index type. For example, the default (FP) index includes BARRAY (barray) and Bitmap (bm) components. The Low_Fast (LF) index includes B-tree (bt) and Bitmap (bm) components.

Also returns the number of pages required to hold the object in memory and the number of IQ pages when the index is compressed (on disk).

You must specify the *index_name* parameter with this procedure. To restrict results to this index name in a single table, include *owner.table*. when specifying the index.

Example

```
sp_iqindexsize ASIQ_IDX_T452_C19_FP
```

User-name	Indexname	Type	Info	KBytes	Pages	Compressed Pages
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	Total	288	4	2
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	vdo	0	0	0
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	bt	0	0	0
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	garray	0	0	0

User-name	Indexname	Type	Info	KBytes	Pages	Compressed Pages
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	bm	136	2	1
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	barray	152	2	1
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	dpstore	0	0	0
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	large-lob	0	0	0
DBA	Employees.ASIQ_IDX_T452_C19_FP	FP	txtPst	0	0	0

```
CREATE TEXT INDEX ti ON Employees( Street ) IMMEDIATE
REFRESH;sp_iqindexsize 'ti';
```

User-name	Indexname	Type	Info	KBytes	Pages	Compressed Pages
GROUPO	GROUPO.Employees.ti	TEXT	Total	896	12	6
GROUPO	GROUPO.Employees.ti	TEXT	vdo	0	0	0
GROUPO	GROUPO.Employees.ti	TEXT	bt	304	4	2
GROUPO	GROUPO.Employees.ti	TEXT	garray	152	2	1
GROUPO	GROUPO.Employees.ti	TEXT	bm	136	2	1
GROUPO	GROUPO.Employees.ti	TEXT	barray	152	2	1
GROUPO	GROUPO.Employees.ti	TEXT	dpstore	0	0	0
GROUPO	GROUPO.Employees.ti	TEXT	large-lob	0	0	0

User-name	Indexname	Type	Info	KBytes	Pages	Compressed Pages
GRUPO	GRUPO.Employees.ti	TEXT	txtPst	304	4	2

See also

- *sp_iqindexmetadata Procedure* on page 442
- *sp_iqindexfragmentation Procedure* on page 438
- *sp_iqindexinfo Procedure* on page 440

sp_iqindexuse Procedure

Reports detailed usage information for secondary (non-FP) indexes accessed by the workload.

Syntax

sp_iqindexuse

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Each secondary index accessed by the workload displays a row. Indexes that have not been accessed do not appear. Index usage is broken down by optimizer, constraint, and query usage.

Indexes from tables created in SYSTEM are not reported.

Table 152. sp_iqindexuse columns

Column name	Description
IndexName	Index name
TableName	Table name
Owner	User name of index owner
UID**	Index unique identifier
Type	Index type
LastDT	Date/time of last access
NOpt	Number of metadata/uniqueness accesses
NQry	Number of query accesses

Column name	Description
NConstraint	Number of accesses for unique or referential integrity checks

**UID is a number assigned by the system that uniquely identifies the instance of the index (where instance is defined when an object is created).

Example

Sample output from the **sp_iqindexuse** procedure.

IndexName	TableName	Owner	UID	Type	LastDT	NQry	NConstraint	NOpt
n_nationkey_hg	nation	DBA	29	HG	20070917	12	0	12
n_regionkey_hg	nation	DBA	31	HG	20070917	12	0	0
r_regionkey_hg	region	DBA	47	HG	20070917	12	0	12
s_suppkey_hg	supplier	DBA	64	HG	20070917	12	0	12
p_partkey_hg	part	DBA	87	HG	20070917	6	0	6
s_suppkey_hg	supplier	DBA	64	HG	20070917	12	0	12
...								

See also

- *sp_iqcolumnuse Procedure* on page 385
- *sp_iqindexadvice Procedure* on page 437
- *sp_iqtableuse Procedure* on page 506
- *sp_iqunusedcolumn Procedure* on page 510
- *sp_iqunusedindex Procedure* on page 511
- *sp_iqunusedtable Procedure* on page 513
- *sp_iqworkmon Procedure* on page 521

sp_iqjoinindex Procedure

Displays information about join indexes.

Syntax

```
sp_iqjoinindex [ left-table-name ], [ left-column-name ], [ left-table-owner ], [ right-table-name ], [ right-column-name ], [ right-table-owner ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

The **sp_iqjoinindex** procedure can be invoked without any parameters. If no parameters are specified, **sp_iqjoinindex** displays information about all join indexes on IQ base tables. Join index tables are always IQ base tables. Join index tables cannot be temporary tables, remote tables, or proxy tables.

Table 153. Parameters

Parameter	Description
left-table-name	The name of the table that forms the left side of the join operation.
left-column-name	The name of the column that is part of the left side of the join.
left-table-owner	The owner of the table that forms the left side of the join operation.
right-table-name	The name of the table that forms the right side of the join operation.
right-column-name	The name of the column that is part of the right side of the join.
right-table-owner	The owner of the table that forms the right side of the join operation.

If you do not specify any of the first five parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. For example, `sp_iqjoinindex NULL, NULL, NULL, t2, n2, DB'` and `sp_iqjoinindex t1, NULL, NULL, t2`.

Table 154. sp_iqjoinindex usage examples

Syntax	Output
<code>sp_iqjoinindex</code>	Displays information about all the join indexes
<code>sp_iqjoinindex t1, NULL, DBA</code>	Displays information about all join indexes in which t1 owned by DBA forms the left side of the operation
<code>sp_iqjoinindex t2, n1, DBA</code>	Displays join index information with column n1 of table t2 owned by DBA as left side of the join
<code>sp_iqjoinindex NULL, NULL, DBA, NULL, NULL, DBA</code>	Displays information about all join indexes in which the left and right side tables are owned by DBA

Syntax	Output
sp_iqjoinindex NULL, NULL, NULL, t2, NULL, NULL	Displays information about all join indexes in which the table t2 is on the right side of the join operation
sp_iqjoinindex t1, n1, DBA, t2, n1, DBA	Displays information about join indexes in which the left side is column n1 of table t1 owned by DBA and the right side is column n1 of table t2 owned by DBA
sp_iqjoinindex non_existing_table	No rows returned, as the table non_existing_table does not exist
sp_iqjoinindex NULL, NULL, non_existing_user	No rows returned, as the user non_existing_user does not exist

Description

The **sp_iqjoinindex** stored procedure displays information about join indexes in a database. If you specify one or more parameters, the result is filtered by the specified parameters. For example, if *left-table-name* is specified, **sp_iqjoinindex** displays all the join indexes in which that table forms the left side of the join. If *left-table-owner* is specified, **sp_iqjoinindex** only returns join indexes in which the left table is owned by the specified owner. If no parameters are specified, **sp_iqjoinindex** displays information about all join indexes in the database.

The **sp_iqjoinindex** procedure returns information in the following columns:

Table 155. sp_iqjoinindex columns

Column name	Description
joinindex_name	The name of the join index.
creator	The owner of the join index.
left_table_name	The name of the table that forms the left side of the join operation.
left_table_owner	The name of the owner of the table that forms the left side of the join operation.
left_column_name	The name of the column that is part of the left side of the join.
join_type	The only currently supported value is “=”.
right_table_name	The name of the table that forms the right side of the join operation.
right_table_owner	The name of the owner of the table that forms the right side of the join operation.
right_column_name	The name of the column that is part of the right side of the join.

Column name	Description
key_type	Defines the type of join on the keys: <ul style="list-style-type: none"> NATURAL: a natural join KEY: a key join ON: a left outer/right outer/full join
valid	Indicates whether this join index needs to be synchronized. 'Y' means that it does not require synchronization; 'N' means that it does require synchronization.
remarks	A comment string.
dbspace_id	Name of the dbspace in which specified join indexes reside.

Examples

Displays information about the join index in which table t1 forms the left side of the join operation:

```
sp_iqjoinindex t1

joindex_name creator left_table_name left_table_owner left_column_name
join_type right_table_name right_table_owner right_column_name key_type
valid dbspace_id remarks
t1_t2_t3_join DBA t1 DBA n1
= t2 DBA n1 NATURAL
Y 16387 (NULL)
```

Displays information about the join index in which table t2 forms the left side of the join operation:

```
sp_iqjoinindex t2

joindex_name creator left_table_name left_table_owner left_column_name
join_type right_table_name right_table_owner right_column_name key_type
valid dbspace_id remarks
t1_t2_t3_join DBA t2 DBA n1
= t3 DBA n1 NATURAL
Y (NULL)
t1_t2_t3_join DBA t2 DBA name
= t3 DBA name NATURAL
Y 16387 ((NULL))
```

Displays information about join indexes in which the left side is column name of table t2 owned by DBA and the right side is column name of table t3 owned by DBA:

```
sp_iqjoinindex t2, name, DBA, t3, name, DBA
```

```

joinindex_name creator left_table_name left_table_owner left_column_
name
join_type right_table_name right_table_owner right_column_name key_t
ype
valid dbspace_id remarks
t1_t2_t3_join DBA t2 DBA DBA name name
= t3 DBA name NATURAL
Y 16387 ((NULL))

```

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqpkkeys Procedure* on page 467
- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iqjoinindexsize Procedure

Gives the size of the specified join index.

Syntax

```
sp_iqjoinindexsize ( join_index_name )
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Returns the total size of the index in bytes, KBytes, and NBlocks (IQ blocks). Also returns the number of pages required to hold the join index in memory and the number of IQ pages when the join index is compressed (on disk). You must specify the *join_index_name* parameter with this procedure.

Table 156. sp_iqjoinindexsize columns

Column name	Description
Username	Owner of the join index

Column name	Description
JoinIndexName	Join index for which results are returned
Number of Tables	Number of tables in the join index
KBytes	Physical object size in KB
Pages	Number of IQ pages needed to hold the object in memory
Compressed Pages	Number of IQ pages when the object is compressed (on disk)
NBlocks	Number of IQ blocks

Example

```
sp_iqjoinindexsize ( 't1t2' )
```

User-name	JoinIndex-Name	Number of Tables	KBytes	Pages	Compressed Pages	NBlocks
DBA	t1t2	2	13	15	4	26

sp_iqlmconfig Procedure

Controls license management configurations, displays and sets license type and status.

Syntax 1

```
sp_iqlmconfig 'edition', { 'SE' | 'SA' | 'EE' }
```

Table 157. Summary information for "edition" parameter

Topic	Value
Default	'EE' (Enterprise Edition)
Range of values	'SE' (Small Business) 'SA' (Single Application) 'EE' (Enterprise Edition)
Status	Static

Syntax 2

```
sp_iqlmconfig 'license type', { 'CP' | 'DT' | 'SF' | 'AC' | 'BC' | 'CH' | 'DH' | 'SH' | 'AH' | 'BH' }
```

Table 158. Summary information for "license type" parameter

Topic	Value
Default	'DT' (Development and Testing)
Range of values	'AC' (OEM CPU License) 'AH' (OEM CPU License Chip) 'BC' (OEM Standby License) 'BH' (OEM Standby License Chip) 'CP' (CPU License) 'CH' (CPU License Chip) 'DH' (Development and Testing License Chip) 'DT' (Development and Testing) 'EV' (Evaluation) 'SF' (Standby CPU License) 'SH' (Standby CPU License Chip)
Status	Static

Syntax 3

```
sp_iqlmconfig 'email severity', { 'ERROR' | 'WARNING'
| 'INFORMATIONAL' | 'NONE' }
```

NONE designates that e-mail notification is disabled.

Syntax 4

```
sp_iqlmconfig 'smtp host', '<host name>' | '
```

Table 159. Parameters

Parameter	Description
host name	Specifies SMTP host used for e-mail notification.

Syntax 5

```
sp_iqlmconfig 'email sender', '<email address>' |
```

Table 160. Parameters

Parameter	Description
<email address>	Specifies the sender's e-mail address used for e-mail notification.

Syntax 6

```
sp_iqlmconfig 'email recipients', '<email recipients>' |
```

Table 161. Parameters

Parameter	Description
<email recipients>	Specifies a comma-separated list of e-mail addresses to whom e-mail notifications will be sent.

Syntax 7

```
sp_iqlmconfig |
```

Permissions

DBA permissions necessary.

Usage

At startup, **sp_iqlmconfig** checks the edition type and license type of the license specified.

- If the specified license is not found, the server falls to grace mode.
- The specified license type becomes valid only when a non-null edition value is specified.
- If **sp_iqlmconfig** is called with no parameters (Syntax 3), it displays all the information above, as well as other information, such as:
 - Product Edition and License Type
 - Which of the optional licenses are in use
 - License count
 - E-mail information
 - General information about the license

See also

- *Properties Available for the Server* on page 120

sp_iqlocks Procedure

Shows information about locks in the database, for both the IQ store and the catalog store.

Syntax

```
sp_iqlocks ([connection,] [[owner.]table_name] max_locks,]
[sort_order])
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Optional **sp_iqlocks** parameters you can specify to restrict results:

Table 162. Optional sp_iqlocks parameters

Parameter	Data Type	Description
<i>connection</i>	integer	Connection ID. With this option, the procedure returns information about locks for the specified connection only. Default is zero, which returns information about all connections.
<i>owner.table_name</i>	char(128)	Table name. With this option, the procedure returns information about locks for the specified table only. Default is NULL, which returns information about all tables in the database. If you do not specify owner, it is assumed that the caller of the procedure owns the table.
<i>max_locks</i>	integer	Maximum number of locks for which to return information. Default is 0, which returns all lock information.
<i>sort_order</i>	char(1)	Order in which to return information: <ul style="list-style-type: none"> • C sorts by connection (default) • T sorts by table_name

Description

Displays information about current locks in the database. Depending on the options you specify, you can restrict results to show locks for a single connection, a single table, or a specified number of locks.

sp_iqlocks displays the following information, sorted as specified in the sort_order parameter:

Table 163. sp_iqlocks columns

Column	Data type	Description
conn_name	VARCHAR(128)	The name of the current connection.
conn_id	INTEGER	Connection ID that has the lock.
user_id	CHAR(128)	User associated with this connection ID.
table_type	CHAR(6)	The type of table. This type is either BASE for a table, GLBTMP for global temporary table, or MVIEW for a materialized view. Materialized views are only supported for SQL Anywhere tables in the IQ catalog store.

Column	Data type	Description
creator	VARCHAR(128)	The owner of the table.
table_name	VARCHAR(128)	Table on which the lock is held.
index_id	INTEGER	The index ID or NULL
lock_class	CHAR(8)	<p>String of characters indicating the type of lock:</p> <ul style="list-style-type: none"> • S – share. • SW – share and write. • EW – exclusive and write. • E – exclusive. • P – phantom. • A – antiphantom. • W – write. <p>All locks listed have one of S, E, EW, or SW, and may also have P, A, or both. Phantom and anti-phantom locks also have a qualifier of T or *:</p> <ul style="list-style-type: none"> • T – the lock is with respect to a sequential scan. • * – the lock is with respect to all scans. • <i>nnn</i> – Index number; the lock is with respect to a particular index. <p>Sybase IQ obtains a share lock before a write lock. If a connection has exclusive lock, share lock does not appear. For write locks, if a connection has all-exclusive, share, and write locks, it is EW.</p>
lock_duration	CHAR(11)	The duration of the lock. One of Transaction, Position, or Connection.
lock_type	CHAR(9)	Value identifying the lock (dependent on the lock class)
row_identifier	UNSIGNED BIGINT	The identifier for the row or NULL.

If **sp_iqlocks** cannot find the connection ID or user name of the user who has a lock on a table, it displays a 0 (zero) for the connection ID and `User unavailable` for the user name.

Note: Exclusive, phantom, or antiphantom locks can be placed on SQL Anywhere tables, but not on Sybase IQ tables. Unless you have explicitly taken out locks on a table in the catalog store, you never see these types of locks (or their qualifiers T, *, and *nnn*) in a Sybase IQ database. For information on how locking works in SQL Anywhere tables, see *SQL Anywhere Server – SQL Usage*.

Examples

The example shows the **sp_iqlocks** procedure call and its output in a Sybase IQ database. The procedure is called with all default options, so that the output shows all locks, sorted by connection.

```
call sp_iqlocks()

conn_name  conn_id  user_id  table_type  creator  table_name
=====  =====  =====  =====  =====  =====
con1       70187172  'mary'   BASE        DBA      t1

index_id   lock_class  lock_duration  lock_type  row_iden
tifier
=====  =====  =====  =====  =====
ASIQ_IDX_T452_C19_FP  Table        Position      Table      1
```

sp_iqmodifyadmin Procedure

Sets an option on a named login policy to a certain value. If no login policy is specified, the option is set on the root policy. In a multiplex, **sp_iqmodifyadmin** takes an optional parameter that is the multiplex server name.

Syntax1

```
call sp_iqmodifyadmin ('policy_option_name', 'value_in' ,
['login_policy_name'] )
```

Syntax2

```
sp_iqmodifyadmin 'policy_option_name', 'value_in' , 'login_policy_name '
```

Syntax3

```
sp_iqmodifyadmin policy_option_name, value_in, ,login_policy_name
```

Syntax 4

```
sp_iqmodifyadmin 'policy_option_name',
'value_in' , 'login_policy_name ' , 'server_name '
```

Usage

Table 164. Parameters

Parameter	Description
policy_option_name	The login policy option to be changed.
value_in	New value for the login policy option.

Parameter	Description
login_policy_name	Name of the login policy whose login policy option needs to be changed.

Permissions

Requires DBA authority.

Examples

Sets the login option locked to ON for the policy named *lockeduser*:

```
call sp_iqmodifyadmin ('locked','on','lockeduser')
```

Sets the login option locked to ON for the policy named *lockeduser* on the multiplex server named *Writer1*:

```
call sp_iqmodifyadmin ('locked','on','lockeduser','Writer1')
```

See also

- *sp_expireallpasswords Procedure* on page 362
- *sp_iqaddlogin Procedure* on page 363
- *sp_iqcopyloginpolicy Procedure* on page 394
- *sp_iqmodifylogin Procedure* on page 458
- *sp_iqpassword Procedure* on page 465

sp_iqmodifylogin Procedure

Assigns a user to a login policy.

Syntax1

```
call sp_iqmodifylogin 'userid', ['login_policy_name']
```

Syntax2

```
sp_iqmodifylogin 'userid', ['login_policy_name']
```

Permissions

DBA authority required.

*Usage***Table 165. Parameters**

Parameter	Description
userid	Variable that holds the name of the account to modify.
login_policy_name	(Optional) Specifies the name of the login policy to which the user will be assigned. If no login policy name is specified, the user is assigned to the root login policy.

Examples

Assigns user joe to a login policy named expired_password:

```
sp_iqmodifylogin 'joe', 'expired_password'
```

Assigns user joe to the root login policy:

```
call sp_iqmodifylogin ('joe')
```

See also

- *sp_expireallpasswords Procedure* on page 362
- *sp_iqaddlogin Procedure* on page 363
- *sp_iqcopyloginpolicy Procedure* on page 394
- *sp_iqmodifyadmin Procedure* on page 457
- *sp_iqpassword Procedure* on page 465

sp_iqmpxcheckdqpconfig Procedure

sp_iqmpxcheckdqpconfig is a diagnostic tool that checks the DQP configuration for the current connection. If DQP fails, run **sp_iqmpxcheckdqpconfig** to determine if DQP configuration issues are causing the query distribution failure.

Syntax

```
sp_iqmpxcheckdqpconfig
```

Permissions

No special privileges are required to execute the procedure.

Description

Table 166. Column Descriptions

Column Name	Description
DiagMsgID	Uniquely identifies a diagnostic message
Description	Diagnostic message describing the issue found with DQP configuration

Table 167. Diagnostic Messages

DiagMsgID	Description
0	No issues found with DQP configuration
1	Database is a simplex
2	Multiplex is running in single-node configuration mode
3	Login policy option <code>dqp_enabled</code> is set to OFF
4	Temporary <code>dqp_enabled</code> connection option is set to OFF
5	Logical server context has only one member node
6	Coordinator does not participate in DQP since its named membership in the logical server is currently ineffective
7	Coordinator does not participate in DQP since its logical membership in the logical server is currently ineffective because <code>ALLOW_COORDINATOR_AS_MEMBER</code> option in Root Logical server policy set to OFF

DiagMsgID	Description
8	There is no dbfile in IQ_SHARED_TEMP dbspace
9	All dbfiles in IQ_SHARED_TEMP dbspace are READ ONLY
10	IQ_SHARED_TEMP dbspace is dynamically offline

Example

Sample output from the **sp_iqmpxcheckdqpconfig** procedure:

```
diagmsgid      description
3              Login policy option dqp_enabled is set to OFF
5              Logical server context has only one member node
6              Coordinator does not participate in DQP since its
              named membership in the logical server is
              currently ineffective
7              Coordinator does not participate in DQP since
              its logical membership in the logical server
              is currently ineffective because
              ALLOW_COORDINATOR_AS_MEMBER option in Root
              Logical server policy set to OFF
8              There is no dbfile in IQ_SHARED_TEMP dbspace
```

sp_iqmpxfilestatus procedure

When executed on the coordinator node, displays file status for every shared dbfile on every included secondary node.

Shared dbfiles include all files in IQ_SYSTEM_MAIN, IQ_SHARED_TEMP, and user dbspaces. When executed on a secondary node, displays file status for only the current node. Users must be licensed for the Multiplex Grid Option to run secondary nodes. For complete information, including syntax, see *Using Sybase IQ Multiplex*.

sp_iqmpxinconnpoolinfo Procedure

If run on the coordinator node, displays INC connection pool status for every node. If executed on a secondary node, displays INC connection pool status for just the current node.

Users must be licensed for the Multiplex Grid Option to run secondary nodes. For **sp_iqmpxinconnpoolinfo** syntax and complete description, see *Using Sybase IQ Multiplex*.

sp_iqmpxinheartbeatinfo Procedure

If run on the coordinator node, displays INC heartbeat status for every node. If executed on a secondary node, displays INC heartbeat status for just the current node.

Users must be licensed for the Multiplex Grid Option to run secondary nodes. For **sp_iqmpxinheartbeatinfo** syntax and complete description, see *Using Sybase IQ Multiplex*.

sp_iqmpxinfo Procedure

If run on the coordinator node, displays complete multiplex configuration info for all nodes. If run on a secondary node, displays complete multiplex configuration info for only that node.

Users must be licensed for the Multiplex Grid Option to run secondary nodes. For **sp_iqmpxinfo** syntax and complete description, see *Using Sybase IQ Multiplex*.

sp_iqmpxvalidate Procedure

Checks multiplex configuration for inconsistencies.

For **sp_iqmpxvalidate** syntax and complete description, see *Using Sybase IQ Multiplex*.

sp_iqmpxversioninfo Procedure

Shows the current version information for this server.

For **sp_iqmpxversioninfo** syntax and complete description, see *Using Sybase IQ Multiplex*.

sp_iqobjectinfo Procedure

Returns partitions and dbspace assignments of database objects and sub-objects.

Syntax

```
sp_iqobjectinfo [ owner_name ] [ , object_name ] [ , object-type ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 168. Parameter

Parameter	Description
owner_name	Owner of the object. If specified, sp_iqobjectinfo displays output only for tables and join indexes with the specified owner. If not specified, sp_iqobjectinfo displays information on tables and join indexes for all users in the database.
object_name	Name of the table or join index. If not specified, sp_iqobjectinfo displays information on all tables and join indexes in the database.
object-type	Valid object types are table (the default) or join-index . If the object-type is a table, it must be enclosed in quotation marks.

All parameters are optional, and any parameter may be supplied independent of the value of another parameter.

Sybase recommends that you use input parameters with **sp_iqobjectinfo**; you can query the results of the **sp_iqobjectinfo** and it performs better if you use input parameters rather than using predicates in the **WHERE** clause of the query. For example, Query A is written as:

```
SELECT COUNT(*) FROM sp_iqobjectinfo()
WHERE owner = 'DBA'
AND object_name = 'tab_case510'
AND object_type = 'table'
AND sub_object_name is NULL
AND dbspace_name = 'iqmain7'
AND partition_name = 'P1'
```

Query B is Query A rewritten to use **sp_iqobjectinfo** input parameters:

```
SELECT COUNT(*) FROM sp_iqobjectinfo('DBA','tab_case510','table')
WHERE sub_object_name is NULL
AND dbspace_name = 'iqmain7'
AND PARTITION_NAME = 'P1'
```

Query B returns results faster than Query A. When the input parameters are passed to **sp_iqobjectinfo**, the procedure compares and joins fewer records in the system tables, thus doing less work compared to Query A. In Query B, the predicates are applied in the procedure itself and the procedure returns a smaller result set, so a smaller number of predicates is applied in the query.

System Procedures

The **sp_iqobjectinfo** stored procedure supports wildcard characters for interpreting *owner_name*, *object_name*, and *object_type*. It shows information for all dbspace that match the given pattern in the same way the **LIKE** clause matches patterns inside queries.

Description

Returns all the partitions and the dbspace assignments of a particular or all database objects (of type table and join index only) and its sub-objects. The sub-objects are columns, indexes, primary key, unique constraints, and foreign keys.

Table 169. sp_iqobjectinfo columns

Column name	Description
owner	Name of the owner of the object.
object_name	Name of the object (of type table and join index only) located on the dbspace.
sub_object_name	Name of the object located on the dbspace.
object_type	Type of the object (column, index, primary key, unique constraint, foreign key, partition, join index or table).
object_id	Global object id of the object.
id	Table id or join-index id of the object.
dbspace_name	Name of the dbspace on which the object resides. The string “[multiple]” is displayed for a special meta row for partitioned objects. The [multiple] row indicates that multiple rows follow in the output to describe the table or column.
partition_name	Name of the partition for the given object.

Examples

Note: The following two examples show objects in the `iqdemo` database to better illustrate output. Note that `iqdemo` includes a sample user dbspace named `iq_main` that may not be present in your own databases.

Displays information about partitions and dbspace assignments of a specific database object and sub-objects owned by a specific user:

```
sp_iqobjectinfo GROUPO,Departments
```

owner	object_name	sub_object_name	object_type	obj
ect_id	id			
GROUPO	Departments	(NULL)	table	3
632	738			
GROUPO	Departments	DepartmentID	column	3
633	738			
GROUPO	Departments	DepartmentName	column	3
634	738			
GROUPO	Departments	DepartmentHeadID	column	3

```

635      738
GROUPO Departments DepartmentsKey          primary
key    83      738
GROUPO Departments FK_DepartmentHeadID_EmployeeID foreign
key    92      738

dbspace_name      partition_name
iq_main           (NULL)
iq_main           (NULL)
iq_main           (NULL)
iq_main           (NULL)
iq_main           (NULL)
iq_main           (NULL)

```

Displays information about partitions and dbspace assignments of a specific database object and sub-objects owned by a specific user for *object-type table*:

```
sp_iqobjectinfo DBA,sale,'table'
```

```

owner  object_name  sub_object_name  object_type  object_id  id
DBA    sale        (NULL)          table        3698      742
DBA    sale        prod_id         column       3699      742
DBA    sale        month_num      column       3700      742
DBA    sale        rep_id         column       3701      742
DBA    sale        sales          column       3702      742

dbspace_name      partition_name
iq_main           (NULL)
iq_main           (NULL)
iq_main           (NULL)
iq_main           (NULL)
iq_main           (NULL)

```

sp_iqpassword Procedure

Changes a user's password.

Syntax1

```
call sp_iqpassword ('caller_password', 'new_password' [, 'user_name'])
```

Syntax2

```
sp_iqpassword 'caller_password', 'new_password' [, 'user_name']
```

Permissions

None to set your own password; DBA or PERMS ADMIN authority required to set other users' passwords.

*Usage***Table 170. Parameters**

Parameter	Description
caller_password	Your password. When you are changing your own password, this is your old password. When the DBA or PERMS ADMIN is changing another user's password, caller_password is the DBA or PERMS ADMIN password.
new_password	New password for the user, or for <i>loginname</i> .
user_name	Login name of the user whose password is being changed by the DBA or PERMS ADMIN. Do not specify user_name when changing your own password.

Description

A user password is an identifier. Any user can change his or her own password using **sp_iqpassword**. The DBA or PERMS ADMIN authority can change the password of any existing user.

The preferred way to create a user is by using the **CREATE USER** statement.

Identifiers have a maximum length of 128 bytes. They must be enclosed in double quotes or square brackets if any of these conditions are true:

- The identifier contains spaces.
- The first character of the identifier is not an alphabetic character (as defined below).
- The identifier contains a reserved word.
- The identifier contains characters other than alphabetic characters and digits.

Alphabetic characters include the alphabet, as well as the underscore character (_), at sign (@), number sign (#), and dollar sign (\$). The database collation sequence dictates which characters are considered alphabetic or digit characters.

Examples

Changes the password of the logged in user from irk103 to exP984:

```
sp_iqpassword 'irk103', 'exP984'
```

Changes the password of user joe from epr45 to pdi032 only if the logged in user has DBA or PERMS ADMIN privileges or the user is joe himself:

```
call sp_iqpassword ('epr45', 'pdi932', 'joe')
```

See also

- *sp_expireallpasswords Procedure* on page 362
- *sp_iqaddlogin Procedure* on page 363
- *sp_iqcopyloginpolicy Procedure* on page 394
- *sp_iqmodifyadmin Procedure* on page 457
- *sp_iqmodifylogin Procedure* on page 458

sp_iqpkkeys Procedure

Displays information about primary keys and primary key constraints by table, column, table owner, or for all Sybase IQ tables in the database.

Syntax

```
sp_iqpkkeys { [ table-name ], [ column-name ], [ table-owner ] }
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

*Usage***Table 171. Parameter**

Parameter	Description
table-name	The name of a base or global temporary table. If specified, the procedure returns information about primary keys defined on the specified table only.
column-name	The name of a column. If specified, the procedure returns information about primary keys on the specified column only.
table-owner	The owner of a table or table. If specified, the procedure returns information about primary keys on tables owned by the specified owner only.

One or more of the parameters can be specified. If you do not specify either of the first two parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. If none of the parameters are specified, a description of all primary keys on all tables in the database is displayed. If any of the specified parameters is invalid, no rows are displayed in the output.

Table 172. sp_iqpkeys usage examples

Syntax	Output
sp_iqpkeys sales	Displays information about primary keys defined on table sales
sp_iqpkeys sales, NULL, DBA	Displays information about primary keys defined on table sales owned by DBA
sp_iqpkeys sales, store_id, DBA	Displays information about primary key defined on column store_id of table sales owned by DBA
sp_iqpkeys NULL, NULL, DBA	Displays information about primary keys defined on all tables owned by DBA

Description

The **sp_iqpkeys** stored procedure displays the following information about primary keys on base and global temporary tables in a database:

Table 173. sp_iqpkeys columns

Column name	Description
table_name	The name of the table
table_owner	The owner of the table
column_name	The name of the column(s) on which the primary key is defined
column_id	The column ID
constraint_name	The name of the primary key constraint
constraint_id	The primary key constraint ID

Note: The **sp_iqpkeys** stored procedure exists only in databases created with Sybase IQ version 12.6 or later.

Examples

Display the primary keys defined on columns of table sales1:

```
sp_iqpkeys sales1

table_name table_owner column_name column_id constraint_name constraint_id
sales1      DBA          store_id      1          MA114          114
```

Display the primary keys defined on columns of table sales2:

```
sp_iqpkeys sales2
```

```

table_name table_owner column_name column_id constraint_name constra
int_id
sales2          DBA          store_id,      1,2          MA115          115
                order_num

```

Display the primary keys defined on the column `store_id` of table `sales2`:

```

sp_iqpkkeys sales2, store_id

table_name table_owner column_name column_id constraint_name constra
int_id
sales2          DBA          store_id      1          MA115          115

```

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iqprocedure Procedure

Displays information about system and user-defined procedures.

Syntax

```

sp_iqprocedure [ proc-name ], [ proc-owner ], [ proc-type ]

```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 174. Parameters

Parameter	Description
proc-name	The name of the procedure.
proc-owner	The owner of the procedure.

Parameter	Description
proc-type	<p>The type of procedure. Allowed values are:</p> <ul style="list-style-type: none"> • SYSTEM: displays information about system procedures (procedures owned by user <i>SYS</i> or <i>dbo</i>) only • ALL: displays information about user and system procedures • Any other value: displays information about user procedures

The **sp_iqprocedure** procedure can be invoked without any parameters. If no parameters are specified, only information about user-defined procedures (procedures not owned by *dbo* or *SYS*) is displayed by default.

If you do not specify either of the first two parameters, but specify the next parameter in the sequence, you must substitute *NULL* for the omitted parameters. For example, `sp_iqprocedure NULL, NULL, SYSTEM` and `sp_iqprocedure NULL, user1`.

Table 175. sp_iqprocedure usage examples

Syntax	Output
<code>sp_iqprocedure</code>	Displays information about all procedures in the database not owned by <i>dbo</i> or <i>SYS</i>
<code>sp_iqprocedure sp_test</code>	Displays information about the procedure sp_test
<code>sp_iqprocedure non_existing_proc</code>	No rows returned, as the procedure non_existing_proc does not exist
<code>sp_iqprocedure NULL, DBA</code>	Displays information about all procedures owned by <i>DBA</i>
<code>sp_iqprocedure sp_test, DBA</code>	Displays information about the procedure sp_test owned by <i>DBA</i>
<code>sp_iqprocedure sp_iqtable</code>	The procedure sp_iqtable is not a system procedure. If there is no user-defined procedure also named sp_iqtable , no rows are returned. (By default only user-defined procedures are returned.)
<code>sp_iqprocedure sp_iqtable, dbo</code>	No rows returned, as the procedure sp_iqtable is not a user procedure (by default only user procedures returned)

Syntax	Output
<code>sp_iqprocedure NULL, NULL, SYSTEM</code>	Displays information about all system procedures (owned by <code>dbo</code> or <code>SYS</code>)
<code>sp_iqprocedure sp_iqtable, NULL, 'SYSTEM</code>	Displays information about the system procedure sp_iqtable
<code>sp_iqprocedure sp_iqtable, dbo, ALL</code>	Displays information about the system procedure sp_iqtable owned by <code>dbo</code>

Description

The **sp_iqprocedure** stored procedure displays information about procedures in a database. If you specify one or more parameters, the result is filtered by the specified parameters. For example, if *proc-name* is specified, only information about the specified procedure is displayed. If *proc-owner* is specified, **sp_iqprocedure** returns only information about procedures owned by the specified owner. If no parameters are specified, **sp_iqprocedure** displays information about all the user-defined procedures in the database.

The **sp_iqprocedure** procedure returns information in the following columns:

Table 176. sp_iqprocedure columns

Column name	Description
<code>proc_name</code>	The name of the procedure
<code>proc_owner</code>	The owner of the procedure
<code>proc_defn</code>	The command used to create the procedure. For hidden procedures, the keyword 'HIDDEN' is displayed.
<code>replicate</code>	Displays Y if the procedure is a primary data source in a Replication Server installation; N if not.
<code>srvid</code>	Indicates the remote server, if the procedure is on a remote database server
<code>remarks</code>	A comment string

Examples

Displays information about the user-defined procedure `sp_test`:

```
sp_iqprocedure sp_test

proc_name      proc_owner      proc_defn      replicate      srvid      r
remarks
sp_test        DBA              create procedure N              (NULL)      (NULL)
                DBA.sp_test(in n1
                integer)
                begin message 'sp_test' end
```

System Procedures

Displays information about all procedures owned by user DBA:

```
sp_iqprocedure NULL, DBA
```

proc_name	proc_owner	proc_defn	replicate	srvid	remarks
sp_test	DBA	create procedure DBA.sp_test(in n1 integer) begin message 'sp_test' end	N	(NULL)	(NULL)
sp_dept	DBA	create procedure DBA.sp_dept() begin end	N	(NULL)	(NULL)

sp_iqprocparm Procedure

Displays information about stored procedure parameters, including result set variables and SQLSTATE/SQLCODE error values.

Syntax

```
sp_iqprocparm [ proc-name ], [ proc-owner ], [ proc-type ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 177. Parameters

Parameter	Description
proc-name	The name of the procedure.
proc-owner	The owner of the procedure.
proc-type	The type of procedure. Allowed values are: <ul style="list-style-type: none">• SYSTEM: displays information about system procedures (procedures owned by user SYS or dbo) only• ALL: displays information about user and system procedures• Any other value: displays information about user procedures

You can invoke **sp_iqprocparm** without parameters. If you do not specify any parameters, input/output and result parameters of user-defined procedures (procedures not owned by dbo or SYS) appear.

If you do not specify either of the first two parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. For example, `sp_iqprocparm NULL, NULL, SYSTEM` and `sp_iqprocparm NULL, user1`.

Table 178. sp_iqprocparm usage examples

Syntax	Output
<code>sp_iqprocparm</code>	Displays parameters for all procedures in the database not owned by <code>dbo</code> or <code>SYS</code>
<code>sp_iqprocparm sp_test</code>	Displays information about the procedure sp_test
<code>sp_iqprocparm non_existing_proc</code>	No rows returned, as the procedure non_existing_proc does not exist
<code>sp_iqprocparm NULL, DBA</code>	Displays parameters for all procedures owned by <code>DBA</code>
<code>sp_iqprocparm sp_test, DBA</code>	Displays parameters for the procedure sp_test owned by <code>DBA</code>
<code>sp_iqprocparm sp_iqtable</code>	sp_iqtable is a system procedure. If there is no user-defined procedure also named sp_iqtable , no rows are returned. (By default, only user-defined procedures are returned.)
<code>sp_iqprocparm sp_iqtable, dbo</code>	No rows returned, as the procedure sp_iqtable is not a user procedure. (By default, only user procedures are returned.)
<code>sp_iqprocparm NULL, NULL, SYSTEM</code>	Displays parameters for all system procedures (owned by <code>dbo</code> or <code>SYS</code>)
<code>sp_iqprocparm sp_iqtable, NULL, SYSTEM</code>	Displays parameters of the system procedure sp_iqtable
<code>sp_iqprocparm sp_iqtable, dbo, ALL</code>	Displays parameters of the system procedure sp_iqtable owned by <code>dbo</code>

Description

The **sp_iqprocparm** stored procedure displays information about stored procedure parameters, including result set variables and `SQLSTATE/SQLCODE` error values. If you specify one or more parameters, the result is filtered by the specified parameters. For example, if *proc-name* is specified, only information about parameters to the specified procedure displays. If *proc-owner* is specified, **sp_iqprocparm** only returns information about parameters to procedures owned by the specified owner. If no parameters are specified, **sp_iqprocparm** displays information about parameters to all the user-defined procedures in the database.

The **sp_iqprocparm** procedure returns information in the following columns:

Table 179. sp_iqprocparm columns

Column name	Description
proc_name	The name of the procedure
proc_owner	The owner of the procedure
parm_name	The name of the parameter
parm_type	The type of parameter is one of the following values: <ul style="list-style-type: none"> • normal parameter (variable) • result variable: used with procedures that return result sets • SQLSTATE error value • SQLCODE error value
parm_mode	The mode of the parameter: whether a parameter supplies a value to the procedure, returns a value, does both, or does neither. Parameter mode is one of the following: <ul style="list-style-type: none"> • in: parameter supplies a value to the procedure • out: parameter returns a value • inout: parameter supplies as well as returns a value • NULL: parameter neither supplies nor returns a value
domain_name	The name of the data type of the parameter as listed in the SYSDOMAIN system table
width	The length of string parameters, the precision of numeric parameters, and the number of bytes of storage for all other data types
scale	The number of digits after the decimal point for numeric data type parameters and zero for all other data types
default	The default value of the parameter, held as a string

Examples

Display information about the parameters of the user-defined procedure sp_test:

```
sp_iqprocparm sp_test

proc_name proc_owner parm_name parm_type parm_mode domain_name width
h scale default
sp_test DBA ID
normal in integer 4 0 (NULL)
```

Display information about the parameters of the system procedure **sp_iqshowcompression**:

```
sp_iqprocparm sp_iqshowcompression, dbo, system
```

proc_name	width	proc_owner	scale	parm_name	parm_type	parm_mode
domain_name		default				
sp_iqshowcompression		dbo		@owner_name	normal	in
char	128		0 (NULL)			
sp_iqshowcompression		dbo		@table_name	normal	in
char	128		0 (NULL)			
sp_iqshowcompression		dbo		@column_name	normal	in
char	128		0 (NULL)			
sp_iqshowcompression		dbo		Column	result	out
char	128		0 (NULL)			
sp_iqshowcompression		dbo		Compression	result	out
char	3		0 (NULL)			

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqpkeys Procedure* on page 467
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iqrebuildindex Procedure

Rebuilds one or more indexes on a table with the original **IQ UNIQUE** value specified in the **CREATE TABLE** statement, or a new **IQ UNIQUE** value to change storage required and/or query performance.

To rebuild an index other than the default index, specify the index name.

Syntax

```
sp_iqrebuildindex (table_name, index_clause)
```

Permissions

You must have EXECUTE permission to run this procedure. You must have INSERT permission on a table to rebuild an index on that table.

Usage

Table 180. Parameters

Parameter	Description
table_name	Partial or fully qualified table name on which the index rebuild process takes place. If the user both owns the table and executes the procedure, a partially qualified name may be used; otherwise, the table name must be fully qualified.
index_clause	<p>One or more of the following strings, separated by spaces:</p> <p>column <i>column_name</i> [<i>count</i>]</p> <p>index <i>index_name</i></p> <p>Each <i>column_name</i> or <i>index_name</i> must refer to a column or index on the specified table. If you specify a <i>column_name</i> or <i>index_name</i> multiple times, the procedure returns an error and no index is rebuilt.</p> <p>The <i>count</i> is a nonnegative number that represents the IQ UNIQUE value. In a CREATE TABLE statement, IQ UNIQUE (count) approximates how many distinct values can be in a given column. The number of distinct values affects query speed and storage requirements.</p>

You must specify the keywords **column** and **index**. These keywords are not case-sensitive.

Note: This procedure does not support **TEXT** indexes. To rebuild a **TEXT** index you must drop and re-create the index.

Description

If you specify a column name, the procedure rebuilds the default index for that column, and no index name is needed. Specifying the name of the default index assigned by Sybase IQ in addition to the column name in this situation returns an error. If you omit *count* after the *column_name*, value 0 (zero) is used as the default.

If the default index is a one-byte index, **sp_iqrebuildindex** always rebuilds it as a one-byte index no matter what **IQ UNIQUE** value the user specified.

For one-byte default indexes, if the specified value in *column_name (count)* is 0 or greater than 256, the column's cardinality value is used to update the `approx_unique_count` column in `SYS.SYSIQCOLUMN`.

sp_iqrebuildindex rebuilds a **WD** index on a column of data type `LONG VARCHAR (CLOB)`.

If the default index is a two-byte index, and the specified count is 0 or greater than 65536, the column's cardinality value determines whether to rebuild the default into a one-byte or two-byte index, and that value is used to update the `approx_unique_count` column in **SYS.SYSIQCOLUMN**.

If you specify a nonzero **IQ UNIQUE** value, the default index is rebuilt as a one-byte, two-byte, or flat default index, with exceptions described above.

If you specify an **IQ UNIQUE** value of zero or no **IQ UNIQUE** value, the `MINIMIZE_STORAGE` option controls how the index is rebuilt:

- If `MINIMIZE_STORAGE` option is set `ON`, the index is rebuilt as a one-byte default index first, and converted to two-byte or flat if necessary.
- If `MINIMIZE_STORAGE` is set `OFF`, the index is rebuilt using the default for the data type.

Examples

Rebuilds the default index on column *Surname*:

```
sp_iqrebuildindex 'empl', 'column dept_id'
```

or:

```
call sp_iqrebuildindex ('empl', 'column dept_id')
```

Creates a flat default index on column `c1`:

```
CREATE TABLE mytable (c1 int IQ UNIQUE 1000000000)
```

Converts the default one-byte index to a two-byte index:

```
sp_iqrebuildindex 'mytable', 'column c1 1024'
```

or:

```
call sp_iqrebuildindex ('mytable', 'column c1 1024')
```

Note: Users can expect to see a temporary performance drop when **sp_iqrebuildindex** runs on a large HG index.

See also

- *sp_iqindexfragmentation Procedure* on page 438
- *sp_iqrowdensity Procedure* on page 482

sp_iqrename Procedure

Renames user-created tables, columns, indexes, constraints (unique, primary key, foreign key, and check), stored procedures, and functions.

Syntax

```
sp_iqrename object-name, new-name [, object-type ]
```

Permissions

Must be the owner of the table or have DBA authority or alter permission on the object.
Requires exclusive access to the object.

Usage

Table 181. Parameters

Parameter	Description
object-name	<p>The original name of the user-created object.</p> <p>Optionally, <i>owner-name</i> can be specified as part of <i>object-name</i> as <i>owner-name.object-name</i>, where <i>owner-name</i> is the name of the owner of the object being renamed. If <i>owner-name</i> is not specified, the user calling sp_iqrename is assumed to be the owner of the object. The object is successfully renamed only if the user calling sp_iqrename has the required permissions to rename the object.</p> <p>If the object to be renamed is a column, index, or constraint, you <i>must</i> specify the name of the table with which the object is associated. For a column, index, or constraint, <i>object-name</i> can be of the form <i>table-name.object-name</i> or <i>owner-name.table-name.object-name</i>.</p>
new-name	The new name of the object. The name must conform to the rules for identifiers and must be unique for the type of object being renamed.
object-type	An optional parameter that specifies the type of the user-created object being renamed, that is, the type of the object <i>object-name</i> . The <i>object-type</i> parameter can be specified in either upper or lowercase.

Values for the object-type parameter:

Table 182. sp_iqrename object-type parameter values

<i>object-type</i> parameter	Specifies
column	The object being renamed is a column

<i>object-type</i> parameter	Specifies
index	The object being renamed is an index
constraint	The object being renamed is a unique, primary key, check, or referential (foreign key) constraint
procedure	The object being renamed is a function
<i>object-type</i> not specified	The object being renamed is a table

Warning! You must change appropriately the definition of any dependent object (procedures, functions, and views) on an object being renamed by **sp_iqrename**. The **sp_iqrename** procedure does not automatically update the definitions of dependent objects. You must change these definitions manually.

Description

The **sp_iqrename** stored procedure renames user-created tables, columns, indexes, constraints (unique, primary key, foreign key, and check), and functions.

If you attempt to rename an object with a name that is not unique for that type of object, **sp_iqrename** returns the message "Item already exists."

sp_iqrename does not support renaming a view, a procedure, an event or a data type. The message "Feature not supported." is returned by **sp_iqrename**, if you specify **event** or **datatype** as the *object-type* parameter.

You can also rename using the **RENAME** clause of the **ALTER TABLE** statement and **ALTER INDEX** statement. See *Reference: Statements and Options*.

Examples

Renames the table `titles` owned by user `shweta` to `books`:

```
sp_iqrename shweta.titles, books
```

Renames the column `id` of the table `books` to `isbn`:

```
sp_iqrename shweta.books.id, isbn, column
```

Renames the index `idindex` on the table `books` to `isbnindex`:

```
sp_iqrename books.idindex, isbnindex, index
```

Renames the primary key constraint `prim_id` on the table `books` to `prim_isbn`:

```
sp_iqrename books.prim_id, prim_isbn, constraint
```

sp_iq_reset_identity Procedure

Sets the seed of the Identity/Autoincrement column associated with the specified table to the specified value.

Syntax

```
sp_iq_reset_identity (table_name, table_owner, value)
```

Usage

You must specify *table_name*, *table owner*, and *value*.

Permissions

None required.

Description

The Identity/Autoincrement column stores a number that is automatically generated. The values generated are unique identifiers for incoming data. The values are sequential, are generated automatically, and are never reused, even when rows are deleted from the table. The seed value specified replaces the default seed value and persists across database shutdowns and failures.

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqpkeys Procedure* on page 467
- *sp_iqprocparm Procedure* on page 472
- *sp_iqtable Procedure* on page 501
- *sp_iqview Procedure* on page 515

sp_iq_reset_identity Procedure Example

Use the example as a reference for **sp_iq_reset_identity** usage.

The following example creates an Identity column with a starting seed of 50:

```
CREATE TABLE mytable(c1 INT identity)
```

```
call sp_iq_reset_identity('mytable', 'dba', 50)
```

sp_iqrestoreaction Procedure

Shows what restore actions are needed to bring database to a consistent state with a given past date.

Syntax

```
sp_iqrestoreaction [ timestamp ]
```

Usage

Table 183. Parameters

Parameter	Description
timestamp	Specifies the past date target.

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

`sp_iqrestoreaction` returns an error if the database cannot be brought to a consistent state for the timestamp. Otherwise, suggests restore actions that will return the database to a consistent state.

The common point to which the database can be restored coincides with the last backup time that backed up read-write files just before the specified timestamp. The backup may be all-inclusive or read-write files only.

Output may not be in exact ascending order based on backup time. If a backup archive consists of multiple read-only dbfiles, it may contain multiple rows (with the same backup time and backup id).

If you back up a read-only dbfile or dbspace multiple times, the restore uses the last backup. The corresponding backup time could be after the specified timestamp, as long as the dbspace/dbfile alter ID matches the dbspace/dbfile alter ID recorded in the last read-write backup that is restored.

`sp_iqrestoreaction` returns the following:

Table 184. sp_iqrestoreaction columns

Column name	Description
sequence_number	Orders the steps to be taken
backup_id	Identifier for the backup transaction

Column name	Description
backup_archive_list	List of archive files in the backup
backup_time	Time of the backup taken
virtual_type	Type of virtual backup: “Non-virtual,” “Decoupled,” or “Encapsulated”
restore_dbpace	Can be empty. Indicates that all dbspaces are to be restored from the backup archive
restore_dbfile	Could be empty. Indicates that all dbfiles in the given dbspace are to be restored from the backup archive
backup_comment	User comment

Example

Sample output of **sp_iqrestoreaction**:

```

sequence_number  backup_id  backup_archive_list  backup_time
14:47:40.0      1          1192  c:\\\\temp\\\\b1      2008-09-23
14:48:05.01     2          1201  c:\\\\temp\\\\b2.inc  2008-09-23
14:48:13.0      3          1208  c:\\\\temp\\\\b3.inc  2008-09-23

virtual_type     restore_dbpace  restore_dbfile  backup_comment
Nonvirtual
Nonvirtual
Nonvirtual
    
```

sp_iqrowdensity Procedure

Reports information about the internal row fragmentation for a table at the FP index level.

Syntax

```

dbo.sp_iqrowdensity ( 'target ' )
target: (table table-name | (column column-name (...))
    
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Parameter	Description
table-name	Target table <i>table-name</i> reports on all columns in the named table.
column-name	Target column <i>column-name</i> reports on the named column in the target table. You may specify multiple target columns, but must repeat the keyword each time.

You must specify the keywords **table** and **column**. These keywords are not case-sensitive.

Description

`sp_iqrowdensity` measures row fragmentation at the default index level. Density is the ratio of the minimum number of pages required by an index for existing table rows to the number of pages actually used by the index. This procedure returns density as a number such that $0 < density < 1$. For example, if an index that requires 8 pages minimum storage occupies 10 pages, its density is .8.

The density reported does not indicate the number of disk pages that may be reclaimed by re-creating or reorganizing the default index.

This procedure displays information about the row density of a column, but does not recommend further action. You must determine whether or not to re-create, reorganize, or rebuild an index.

Example

Reports the row density on column *ID* in table *SalesOrders*:

```
sp_iqrowdensity('column groupo.SalesOrders.ID')
```

Tablename	ColumnName	IndexType	Density
'GROUPO.SalesOrders'	'ID'	'Flat style FP'	'1.0'

See also

- *sp_iqindexfragmentation Procedure* on page 438
- *sp_iqrebuildindex Procedure* on page 475

sp_iqsharedtempdistrib procedure

Returns diagnostic information about how shared temp space is distributed among nodes of the multiplex and whether that space is usable or quarantined.

Users must be licensed for the Multiplex Grid Option to run secondary nodes. For **sp_iqsharedtempdistrib** syntax and complete description, see *Using Sybase IQ Multiplex*.

sp_iqshowpsex Procedure

Displays information about the settings of database options that control the priority of tasks and resource usage for connections.

Syntax

```
sp_iqshowpsex [ connection-id ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

Table 185. Parameters

Parameter	Description
connection-id	<p>An integer representing the connection ID.</p> <p>If <i>connection-id</i> is specified, sp_iqshowpsex returns information only about the specified connection. If <i>connection-id</i> is not specified, sp_iqshowpsex returns information about all connections.</p> <p>If the specified <i>connection-id</i> does not exist, sp_iqshowpsex returns no rows.</p>

Description

The **sp_iqshowpsex** stored procedure displays information about the settings of database options that control the priority of tasks and resource usage for connections, which is useful to database administrators for performance tuning.

Table 186. sp_iqshowpsex columns

Column name	Description
connectionid	The connection ID
application	Information about the client application that opened the connection. Includes the AppInfo connection property information: HOST: the host name of the client machine EXE: the name of the client executable (Windows only) APPINFO: the APPINFO in the client connection string, if specified
userid	Login name of the user that opened the connection

Column name	Description
iqgovern_priority	Value of the database option IQGOVERN_PRIORITY that assigns a priority to each query waiting in the -iqgovern queue. By default, this option has a value of 2 (MEDIUM). The values 1, 2, and 3 are shown as HIGH, MEDIUM, and LOW, respectively.
max_query_time	Value of the database option MAX_QUERY_TIME that sets a limit, so that the optimizer can disallow very long queries. By default, this option is disabled and has a value of 0.
query_row_limit	Value of the database option QUERY_ROWS_REJECTED_LIMIT that sets the row threshold for rejecting queries based on the estimated size of the result set. The default is 0, which means there is no limit.
query_temp_space_limit	Value of the database option QUERY_TEMP_SPACE_LIMIT (in MB) that constrains the use of temporary IQ dbspace by user queries. The default value is 2000MB.
max_cursors	Value of the database option MAX_CURSOR_COUNT that specifies a resource governor to limit the maximum number of cursors a connection can use at once. The default value is 50. A value of 0 implies no limit.
max_statements	Value of the database option MAX_STATEMENT_COUNT that specifies a resource governor to limit the maximum number of prepared statements that a connection can use at once. The default value is 100. A value of 0 implies no limit.

For information about the database options referenced in the columns, see *Reference: Statements and Options*.

Note: The **AppInfo** property may not be available from Open Client or jConnect applications such as Interactive SQL or Sybase Central. If the **AppInfo** property is not available, the application column is blank.

Example

Display information about the settings of database options that control the priority of tasks and resource usage for connection ID 2:

```
sp_iqshowpsexec 2

connectionid      application
                2      HOST=GOODGUY-XP;EXE=C:\\Program Files\\Sybase\\
                    IQ-15_4\\bin32\\dbisqlg.exe;

userid   iqgovern_priority   max_query_time   query_row_limit
DBA      MEDIUM              0                 0
```

query_temp_space_limit	max_statements	max_cursors
2000	50	100

See also

- *CONNECTION_PROPERTY* Function [System] on page 142
- *sp_iqcontext* Procedure on page 391

sp_iqspaceinfo Procedure

Displays the number of blocks used by each object in the current database and the name of the dbspace in which the object is located.

Syntax

```
sp_iqspaceinfo [ 'main
| [table table-name | index index-name] [...] ' ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

For the current database, displays the object name, number of blocks used by each object, and the name of the dbspace. **sp_iqspaceinfo** requires no parameters.

The information returned by **sp_iqspaceinfo** is helpful in managing dbspaces.

You can use this procedure in a multiplex environment. See *Using Sybase IQ Multiplex*.

Example

The following output is from the **sp_iqspaceinfo** stored procedure run on the *iqdemo* database. Output for some tables and indexes have been removed in this example.

Name	NBlocks	dbspace_name
Contacts	19	IQ_SYSTEM_MAIN
SalesOrderItems.DBA.ASIQ_IDX_T205_C5_FP	56	IQ_SYSTEM_MAIN
Contacts.DBA.ASIQ_IDX_T206_C10_FP	55	IQ_SYSTEM_MAIN
Contacts.DBA.ASIQ_IDX_T206_C1_FP	61	IQ_SYSTEM_MAIN
...		
Contacts.DBA.ASIQ_IDX_T206_C9_FP	55	IQ_SYSTEM_MAIN
Contacts.DBA.ASIQ_IDX_T206_I11_HG	19	IQ_SYSTEM_MAIN
Customers	20	IQ_SYSTEM_MAIN
Customers.DBA.ASIQ_IDX_T207_C1_FP	61	IQ_SYSTEM_MAIN
Customers.DBA.ASIQ_IDX_T207_C2_FP	55	IQ_SYSTEM_MAIN
...		
Customers.DBA.ASIQ_IDX_T207_I10_HG	19	IQ_SYSTEM_MAIN
...		

See also

- *sp_iqindexinfo Procedure* on page 440
- *sp_iqdbspace Procedure* on page 402
- *sp_iqdbspaceinfo Procedure* on page 405

sp_iqspaceused Procedure

Shows information about space available and space used in the IQ store, IQ temporary store, and IQ global and local shared temporary stores.

Syntax

```
sp_iqspaceused(out mainKB unsigned bigint,
               out mainKBUsed unsigned bigint,
               out tempKB unsigned bigint,
               out tempKBUsed unsigned bigint,
out shTempTotalKB unsigned bigint,
               out shTempTotalKBUsed unsigned bigint,
out shTempLocalKB unsigned bigint,
               out shTempLocalKBUsed unsigned
bigint)
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage

sp_iqspaceused returns eight values as unsigned bigint out parameters. This system stored procedure can be called by user-defined stored procedures to determine the amount of main and temporary IQ store space in use.

Description

sp_iqspaceused returns a subset of the information provided by **sp_iqstatus**, but allows the user to return the information in SQL variables to be used in calculations.

Table 187. sp_iqspaceused columns

Column name	Description
mainKB	The total IQ main store space in kilobytes.
mainKBUsed	The number of kilobytes of IQ main store space used by the database. (Secondary multiplex nodes return '(Null)'.)
tempKB	The total IQ temporary store space in kilobytes.
tempKBUsed	The total IQ temporary store space in kilobytes.

Column name	Description
shTempTotalKB	The total IQ global shared temporary store space in kilobytes.
shTempTotalKBUsed	The total IQ global shared temporary store space in kilobytes. (Secondary multiplex nodes return '(Null)').
shTempLocalKB	The total IQ local shared temporary store space in kilobytes.
shTempLocalKBUsed	The number of kilobytes of IQ local shared temporary store space in use by the database.

Example

sp_iqspaceused requires eight output parameters. The following example shows the creation of a user-defined stored procedure **myspace** that declares the eight output parameters and then calls **sp_iqspaceused**:

```

create procedure dbo.myspace()
begin
  declare mt unsigned bigint;
  declare mu unsigned bigint;
  declare tt unsigned bigint;
  declare tu unsigned bigint;
  declare gt unsigned bigint;
  declare gu unsigned bigint;
  declare lt unsigned bigint;
  declare lu unsigned bigint;
  call sp_iqspaceused(mt,mu,tt,tu,gt,gu,lt,lu);
  select cast(mt/1024 as unsigned bigint) as mainMB,
         cast(mu/1024 as unsigned bigint) as mainusedMB,
         mu*100/mt as mainPerCent,
         cast(tt/1024 as unsigned bigint) as tempMB,
         cast(tu/1024 as unsigned bigint) as tempusedMB,
         tu*100/tt as tempPerCent;
         cast(gt/1024 as unsigned bigint) as shTempTotalKB,
         cast(gu/1024 as unsigned bigint) as shTempTotalKBUsed,
         gu*100/gt as globalshTempPerCent;
         cast(lt/1024 as unsigned bigint) as shTempLocalMB,
         cast(lu/1024 as unsigned bigint) as shTempLocalKBUsed,
         lu*100/lt as localshTempPerCent;
end

```

To display the output of **sp_iqspaceused**, run the procedure **myspace**:

```

myspace

```

sp_iqstatistics Procedure

Returns serial number, name, description, value, and unit specifier for each available statistic, or a specified statistic.

Syntax

```
sp_iqstatistics [ stat_name ]
```

Usage

Parameter	Description
stat_name	(Optional) VARCHAR parameter specifying the name of a statistic.

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

When **stat_name** is provided, **sp_iqstatistics** returns one row for the given statistic, or zero rows if the name is invalid. When invoked without any parameter, **sp_iqstatistics** returns all statistics.

Result Set

Column name	Data type	Description
stat_num	UNSIGNED INTEGER	Serial number of a statistic
stat_name	VARCHAR(255)	Name of statistic
stat_desc	VARCHAR(255)	Description of statistic
stat_value	LONG VARCHAR	Value of statistic
stat_unit	VARCHAR(128)	Unit specifier

The following statistics may be returned:

stat_num	stat_name	stat_desc	stat_unit
0	CpuTotalTime	Total CPU time in seconds consumed by the IQ server since last server start-up	Second

stat_num	stat_name	stat_desc	stat_unit
1	CpuUserTime	CPU user time in seconds consumed by the IQ server since last server startup	Second
2	CpuSystemTime	CPU system time in seconds consumed by the IQ server since last server start-up	Second
3	ThreadsFree	Number of IQ threads free	N/A
4	ThreadsInUse	Number of IQ threads in use	N/A
5	MemoryAllocated	Allocated memory in megabytes	MB
6	MemoryMaxAllocated	Max allocated memory in megabytes	MB
7	MainCacheCurrentSize	Main cache current size in megabytes	MB
8	MainCacheFinds	Main cache total number of lookup requests	N/A
9	MainCacheHits	Main cache total number of hits	N/A
10	MainCachePagesPinned	Main cache number of pages pinned	Page
11	MainCachePagesPinnedPercentage	Percentage of main cache pages pinned	%
12	MainCachePagesDirtyPercentage	Percentage of main cache pages dirtied	%
13	MainCachePagesInUsePercentage	Percentage of main cache pages in use	%
14	TempCacheCurrentSize	Temporary cache current size in megabytes	MB
15	TempCacheFinds	Temporary cache total number of lookup requests	N/A
16	TempCacheHits	Temporary cache total number of hits	N/A
17	TempCachePagesPinned	Temporary cache number of pages pinned	Page
18	TempCachePagesPinnedPercentage	Percentage of temporary cache pages pinned	%
19	TempCachePagesDirtyPercentage	Percentage of temporary cache pages dirtied	%

stat_num	stat_name	stat_desc	stat_unit
20	TempCachePagesInUsePercentage	Percentage of temporary cache pages in use	%
21	MainStoreDiskReads	Number of kilobytes read from main store	KB
22	MainStoreDiskWrites	Number of kilobytes written to main store	KB
23	TempStoreDiskReads	Number of kilobytes read from main store	KB
24	TempStoreDiskWrites	Number of kilobytes written to main store	KB
25	ConnectionsTotalConnections	Total number of connections since server startup	N/A
26	ConnectionsTotalDisconnections	Total number of disconnections since server startup	N/A
27	ConnectionsActive	Number of active connections	N/A
28	OperationsWaiting	Number of operations waiting for IQ resource governor	N/A
29	OperationsActive	Number of active concurrent operations admitted by IQ resource governor	N/A
30	OperationsActiveLoadTableStatements	Number of active LOAD TABLE statements	N/A

Examples

Displays a single statistic, the total CPU time:

```
sp_iqstatistics 'CPUTotalTime'
```

Displays all statistics for MainCache%:

```
SELECT * from sp_iqstatistics() WHERE stat_name LIKE 'MainCache%'
```

sp_iqstatus Procedure

Displays a variety of Sybase IQ status information about the current database.

Syntax

```
sp_iqstatus
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Shows status information about the current database, including the database name, creation date, page size, number of dbspace segments, block usage, buffer usage, I/O, backup information, and so on.

sp_iqstatus displays an out-of-space status for main and temporary stores. If a store runs into an out-of-space condition, **sp_iqstatus** shows Y in the store's out-of-space status display value.

sp_iqspaceused returns a subset of the same information as provided by **sp_iqstatus**, but allows the user to return the information in SQL variables to be used in calculations.

To display space that can be reclaimed by dropping connections, use **sp_iqstatus** and add the results from the two returned rows:

```
(DBA)> select * from sp_iqstatus() where name like '%Versions:%'
Execution time: 6.25 seconds
Name           Value
-----
Other Versions: 2 = 1968Mb
Active Txn Versions: 1 = C:2175Mb/D:2850Mb

(First 2 rows)
```

The above example output shows that one active write transaction created 2175MB and destroyed 2850 MB of data. The total data consumed in transactions and not yet released is 4818MB, or 1968MB + 2850MB = 4818MB.

sp_iqstatus does not show blocks that will be deallocated at the next checkpoint. These blocks do however, appear in **sp_iqdbspace** output as type X.

You can use this procedure in a multiplex environment. See *Using Sybase IQ Multiplex*.

Example

Note: The following example shows objects in the iqdemo database to better illustrate output. Note that iqdemo includes a sample user dbspace named iq_main that may not be present in your own databases.

The following output is from the **sp_iqstatus** stored procedure:

```
Sybase IQ (TM)           Copyright (c) 1992-2010 by Sybase, Inc.
                          All rights reserved.
Version:                 15.3.0/090416/P/MS/Windows/2010/
                          32bit/2010-04-16 02:11:41
Time Now:                2010-04-21 13:48:22.319
```

```

Build Time:                2010-04-16 02:15:39
File Format:               23 on 03/18/1999
Server mode:              IQ Server
Catalog Format:            2
Stored Procedure Revision: 1
Page Size:                131072/8192blksz/16bpp
Number of Main DB Files : 2
Main Store Out Of Space:  N
Number of Local Temp DB Files : 1
Local Temp Store Out Of Space: N
DB Blocks: 1-3200         IQ_SYSTEM_MAIN
DB Blocks: 1045440-1055039 iq_main
Local Temp Blocks: 1-1600 IQ_SYSTEM_TEMP
Create Time:              2009-04-03 11:30:20.674
Update Time:              2009-04-03 11:34:33.040
Main IQ Buffers:          255, 32Mb
Temporary IQ Buffers:     191, 24Mb
Main IQ Blocks Used:      5915 of 11200, 52%=46Mb, Max
Block#:105278
Local Temporary IQ Blocks Used: 65 of 800, 8%=0Mb, Max
Block#: 0
Main Reserved Blocks Available: 1600 of 1600, 100%=6Mb
Local Temporary Reserved Blocks Available: 6400 of 6400,
100%=50Mb
IQ Dynamic Memory:        Current: 69mb, Max: 70mb
Main IQ Buffers:          Used: 17, Locked: 0
Temporary IQ Buffers:     Used: 4, Locked: 0
Main IQ I/O:              I: L1581/P14 O: C3/D163/P161
D:34 C:97.1
Temporary IQ I/O:        I: L6627/P0 O: C1086/D1166/P83
D:1082 C:100.0
Other Versions:           0 = 0Mb
Active Txn Versions:      0 = C:0Mb/D:0Mb
Last Full Backup ID:      0
Last Full Backup Time:
Last Backup ID:           0
Last Backup Type:         None
Last Backup Time:
DB Updated:               1
Blocks in next ISF Backup: 0 Blocks: =0Mb
Blocks in next ISI Backup: 0 Blocks: =0Mb
DB File Encryption Status: OFF

```

The following is a key to understanding the Main IQ I/O and Temporary IQ I/O output codes:

- I: Input
- L: Logical pages read (“Finds”)
- P: Physical pages read
- O: Output
- C: Pages created
- D: Pages dirtied

System Procedures

- P: Physically written
- D: Pages destroyed
- C: Compression ratio

See also

- *sp_iqtransaction Procedure* on page 506
- *sp_iqversionuse Procedure* on page 514

sp_iqsysmon Procedure

Monitors multiple components of Sybase IQ, including the management of buffer cache, memory, threads, locks, I/O functions, and CPU utilization.

Batch Mode Syntax

```
sp_iqsysmon start_monitor  
sp_iqsysmon stop_monitor [ , "section(s)" ]  
or  
sp_iqsysmon "time-period" [ , "section(s)" ]
```

File Mode Syntax

```
sp_iqsysmon start_monitor, 'filemode' [ , "monitor-options" ]  
sp_iqsysmon stop_monitor
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Batch Mode Usage

Parameter	Description
start_monitor	Starts monitoring.
stop_monitor	Stops monitoring and displays the report.
time-period	The time period for monitoring. Must be in the form HH:MM:SS.

Parameter	Description
section(s)	<p>The abbreviation for one or more sections to be displayed by sp_iqsysmon. When more than one section is specified, the section abbreviations must be separated by spaces and the list must be enclosed in single or double quotes. The default is to display all sections.</p> <p>For the sections related to IQ store, you can specify main or temporary store by prefixing the section abbreviation with “m” or “t”, respectively. Without the prefix, both stores are monitored. For example, if you specify “mbufman”, only the IQ main store buffer manager is monitored. If you specify “mbufman tbufman” or “bufman”, both the main and temporary store buffer managers are monitored.</p>

Table 188. sp_iqsysmon report section abbreviations

Report section or IQ component	Abbreviation
Buffer manager	(m/t)bufman
Buffer pool	(m/t)bufpool
Prefetch management	(m/t)prefetch
Free list management	(m/t)freelist
Buffer allocation	(m/t)bufalloc
Memory management	memory
Thread management	threads
CPU utilization	cpu
Transaction management	txn
Server context statistics	server
Catalog statistics	catalog

Note: The Sybase IQ components Disk I/O and lock manager are not currently supported by **sp_iqsysmon**.

*File Mode Usage***Table 189. Parameters**

Parameter	Description
start_monitor	Starts monitoring.
stop_monitor	Stops monitoring and writes the remaining output to the log file.
filemode	Specifies that sp_iqsysmon is running in file mode. In file mode, a sample of statistics is displayed for every interval in the monitoring period. By default, the output is written to a log file named <i>dbname.connid-iqmon</i> . Use the file_suffix option to change the suffix of the output file. See the <i>monitor_options</i> parameter for a description of the file_suffix option.
monitor_options	The <i>monitor_options</i> string

The *monitor_options* string can include one or more options:

Table 190. monitor_options string options

monitor_options string option	Description
-interval <i>seconds</i>	<p>Specifies the reporting interval in seconds. A sample of monitor statistics is output to the log file after every interval. The default is every 60 seconds, if the -interval option is not specified. The minimum reporting interval is 2 seconds. If the interval specified for this option is invalid or less than 2 seconds, the interval is set to 2 seconds.</p> <p>The first display shows the counters from the start of the server. Subsequent displays show the difference from the previous display. You can usually obtain useful results by running the monitor at the default interval of 60 seconds during a query with performance problems or during a time of day with performance problems. A very short interval may not provide meaningful results. The interval should be proportional to the job time; 60 seconds is usually more than enough time.</p>
-file_suffix <i>suffix</i>	<p>Creates a monitor output file named <code>dbname.connid-suffix</code>. If you do not specify the -file_suffix option, the suffix defaults to <code>iqmon</code>. If you specify the -file_suffix option and do not provide a suffix or provide a blank string as a suffix, no suffix is used.</p>
-append or -truncate	<p>Directs sp_iqsysmon to append to the existing output file or truncate the existing output file, respectively. Truncate is the default. If both options are specified, the option specified later in the string is effective.</p>

monitor_options string option	Description
<p>-section <i>section(s)</i></p>	<p>Specifies the abbreviation of one or more sections to write to the monitor log file. The default is to write all sections. The abbreviations specified in the sections list in file mode are the same abbreviations used in batch mode. When more than one section is specified, spaces must separate the section abbreviations.</p> <p>If the -section option is specified with no sections, none of the sections are monitored. An invalid section abbreviation is ignored and a warning is displayed in the IQ message file.</p>

Usage Syntax Examples

Table 191. sp_iqsysmon usage examples

Syntax	Result
<p>sp_iqsysmon start_monitor sp_iqsysmon stop_monitor</p>	<p>Starts the monitor in batch mode and displays all sections for main and temporary store</p>
<p>sp_iqsysmon start_monitor sp_iqsysmon stop_monitor "mbufman mbufpool"</p>	<p>Starts the monitor in batch mode and displays the Buffer Manager and Buffer Pool statistics for main store</p>
<p>sp_iqsysmon "00:00:10", "mbufpool memory"</p>	<p>Runs the monitor in batch mode for 10 seconds and displays the consolidated statistics at the end of the time period</p>
<p>sp_iqsysmon start_monitor, 'filemode', "-interval 5 -sections mbufpool memory" sp_iqsysmon stop_monitor</p>	<p>Starts the monitor in file mode and writes to the log file every 5 seconds the statistics for Main Buffer Pool and Memory Manager</p>

Description

The **sp_iqsysmon** stored procedure monitors multiple components of Sybase IQ, including the management of buffer cache, memory, threads, locks, I/O functions, and CPU utilization.

The **sp_iqsysmon** procedure supports two modes of monitoring:

- Batch mode
In batch mode, **sp_iqsysmon** collects the monitor statistics for the period between starting and stopping the monitor or for the time period specified in the *time-period* parameter. At the end of the monitoring period, **sp_iqsysmon** displays a list of consolidated statistics.

sp_iqsysmon in batch mode is similar to the Adaptive Server Enterprise procedure **sp_sysmon**.

- File mode

In file mode, **sp_iqsysmon** writes the sample statistics in a log file for every interval period between starting and stopping the monitor.

The first display in file mode shows the counters from the start of the server. Subsequent displays show the difference from the previous display.

sp_iqsysmon in file mode is similar to the **IQ UTILITIES** command **START MONITOR** and **STOP MONITOR** interface.

Batch Mode Examples

Prints monitor information after 10 minutes:

```
sp_iqsysmon "00:10:00"
```

Prints only the Memory Manager section of the **sp_iqsysmon** report after 5 minutes:

```
sp_iqsysmon "00:05:00", memory
```

Starts the monitor, executes two procedures and a query, stops the monitor, then prints only the Buffer Manager section of the report:

```
sp_iqsysmon start_monitor
go
execute proc1
go
execute proc2
go
select sum(total_sales) from titles
go
sp_iqsysmon stop_monitor, bufman
go
```

Prints only the Main Buffer Manager and Main Buffer Pool sections of the report after 20 minutes:

```
sp_iqsysmon "00:02:00", "mbufman mbufpool"
```

File Mode Examples

Truncates and writes information to the log file every 2 seconds between starting the monitor and stopping the monitor:

```
sp_iqsysmon start_monitor, 'filemode', '-interval 2'
.
.
.
sp_iqsysmon stop_monitor
```

Appends output for only the Main Buffer Manager and Memory Manager sections to an ASCII file with the name `dbname.connid-testmon`. For the database `iqdemo`, writes results in the file `iqdemo.2-testmon`:

```
sp_iqsysmon start_monitor, 'filemode',
"-file_suffix testmon -append -section mbufman memory"
```

System Procedures

```
.
.
.
sp_iqsysmon stop_monitor
```

Example

Run the monitor in batch mode for 10 seconds and display the consolidated statistics at the end of the time period

```
sp_iqsysmon "00:00:10", "mbufpool memory"

=====
Buffer Pool (Main)
=====
STATS-
NAME          TOTAL  NONE  BTREEV  BTREEF  BV  VDO  DBEXT  DBID  SORT
MovedToMRU    0      0     0       0       0   0   0     0     0     0
MovedToWash   0      0     0       0       0   0   0     0     0     0
RemovedFromLRU 0      0     0       0       0   0   0     0     0     0
RemovedFromWash 0     0     0       0       0   0   0     0     0     0
RemovedInScanMode 0     0     0       0       0   0   0     0     0     0

STORE  GARRAY  BARRAY  BLKMAP  HASH  CKPT  BM  TEST  CMID  RIDCA  LOB
0      0       0       0       0     0     0   0     0     0     0
0      0       0       0       0     0     0   0     0     0     0
0      0       0       0       0     0     0   0     0     0     0
0      0       0       0       0     0     0   0     0     0     0
0      0       0       0       0     0     0   0     0     0     0

STATS-NAME          VALUE
Pages                127  ( 100.0 %)
InUse                 4   (  3.1 %)
Dirty                 1   (  0.8 %)
Pinned                0   (  0.0 %)
Flushes               0
FlushedBufferCount   0
GetPageFrame         0
GetPageFrameFailure  0
GotEmptyFrame        0
Washed                0
TimesSweepersWoken  0

washTeamSize         0
WashMaxSize          26  ( 20.5 %)
washNBuffers          4   (  3.1 %)
washNDirtyBuffers    1   (  0.8
%)
                    washSignalThreshold          3   (  2.4 %)
washNActiveSweepers  0
washIntensity         1

=====
Memory Manager
=====
STATS-NAME          VALUE
```

MemAllocated	43616536	(42594 KB)
MemAllocatedMax	43735080	(42710 KB)
MemAllocatedEver	0	(0 KB)
MemNAllocated	67079	
MemNAllocatedEver	0	
MemNTimesLocked	0	
MemNTimesWaited	0	(0.0 %)

sp_iqtable Procedure

Displays information about tables in the database.

Syntax1

```
sp_iqtable ( [ table_name ],[table_owner ],[table_type ] )
```

Syntax2

```
sp_iqtable [table_name='tablename' ],  
[table_owner='tableowner' ],[table_type='tabletype' ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage: Syntax1

If you do not specify either of the first two parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. For example, `sp_iqtable NULL, NULL, TEMP` and `sp_iqtable NULL, dbo, SYSTEM`.

Note: The *table_type* values ALL and VIEW must be enclosed in single quotes in Syntax 1.

Usage: Syntax2

The parameters can be specified in any order. Enclose them in single quotes.

The allowed values for the *table_type* parameter:

Table 192. sp_iqtable table_type values

table_type value	Information displayed
SYSTEM	System tables
TEMP	Global temporary tables
VIEW	Views
ALL	IQ tables, system tables, and views
any other value	IQ tables

Description

Specifying one parameter returns only the tables that match that parameter. Specifying more than one parameter filters the results by all of the parameters specified. Specifying no parameters returns all Sybase IQ tables in the database. There is no method for returning the names of local temporary tables.

Table 193. sp_iqtable columns

Column name	Description
table_name	The name of the table.
table_type	BASE – a base table. MAT VIEW - a materialized view. (SA tables only) GBL TEMP - a global temporary table. PARTITION - a table partition (this table is for internal use only and cannot be used by Sybase IQ users). VIEW – a view. JVT – a join virtual table.
table_owner	The owner of the table
server_type	IQ – an object created in the IQ store. SA – an object created in the SA store. All views are created in the SA store.
location	TEMP – IQ temporary store. MAIN – IQ store. SYSTEM – catalog store.
dbspace_id	Name of the dbspace where the table resides.
isPartitioned	'Y' if the column belongs to a partitioned table and has one or more partitions whose dbspace is different from the table partition's dbspace, 'N' if the column's table is not partitioned or each partition of the column resides in the same dbspace as the table partition.
remarks	User comments added with the COMMENT statement.
table_constraints	Constraints against the table.

See also

- *sp_iqcolumn Procedure* on page 382

- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447
- *sp_iqpkeys Procedure* on page 467
- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqview Procedure* on page 515

sp_iqtable Procedure Examples

Use the examples as reference for **sp_iqtable** usage.

The following variations in syntax both return information about the table Departments:

```
sp_iqtable ('Departments')
```

```
sp_iqtable table_name='Departments'
```

Table_name	Table_type	Table_owner	Server_type	location
Departments	BASE	GRUPO	IQ	Main

dbspace_id	isPartitioned	Remarks	table_constraints
16387	N	(NULL)	(NULL)

The following variations in syntax both return all tables that are owned by table owner GRUPO:

```
sp_iqtable NULL,GRUPO
```

```
sp_iqtable table_owner='GRUPO'
```

Table_name	Table_type	Table_owner	Server_type	location
Contacts	BASE	GRUPO	IQ	Main
Customers	BASE	GRUPO	IQ	Main
Departments	BASE	GRUPO	IQ	Main
Employees	BASE	GRUPO	IQ	Main
FinancialCodes	BASE	GRUPO	IQ	Main
FinancialData	BASE	GRUPO	IQ	Main

Table_name	Table_type	Table_owner	Server_type	location
Products	BASE	GROUPO	IQ	Main
SalesOrders	BASE	GROUPO	IQ	Main
SalesOrderItems	BASE	GROUPO	IQ	Main

dbspace_id	isPartitioned	Remarks	table_constraints
16387	N	(NULL)	(NULL)
16387	N	(NULL)	(NULL)
16387	N	(NULL)	(NULL)
16387	N	(NULL)	(NULL)
16387	N	(NULL)	(NULL)
16387	N	(NULL)	(NULL)
16387	N	(NULL)	(NULL)
16387	N	(NULL)	(NULL)
16387	N	(NULL)	(NULL)

sp_iqtablesize Procedure

Returns the size of the specified table.

Syntax

```
sp_iqtablesize ( table_owner.table_name )
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Returns the total size of the table in KBytes and NBlocks (IQ blocks). Also returns the number of pages required to hold the table in memory, and the number of IQ pages that are compressed when the table is compressed (on disk). You must specify the *table_name* parameter with this procedure. If you are the owner of *table_name*, then you do not have to specify the *table_owner* parameter.

Table 194. sp_iqtablesize columns

Column name	Description
Ownername	Name of owner
Tablename	Name of table
Columns	Number of columns in the table
KBytes	Physical table size in KB
Pages	Number of IQ pages needed to hold the table in memory
CompressedPages	Number of IQ pages that are compressed, when the table is compressed (on disk)
NBlocks	Number of IQ blocks

`Pages` is the total number of IQ pages for the table. The unit of measurement for pages is IQ page size. All in-memory buffers (buffers in the IQ buffer cache) are the same size.

IQ pages on disk are compressed. Each IQ page on disk uses 1 to 16 blocks. If the IQ page size is 128KB, then the IQ block size is 8KB. In this case, an individual on-disk page could be 8, 16, 24, 32, 40, 48, 56, 64, 72, 80, 88, 96, 104, 112, 120, or 128 KB.

If you divide the `KBytes` value by page size, you see the average on-disk page size.

Note: Sybase IQ always reads and writes an entire page, not blocks. For example, if an individual page compresses to 88K, then IQ reads and writes the 88K in one I/O. The average page is compressed by a factor of 2 to 3.

`NBlocks` is `Kbytes` divided by IQ block size.

`CompressedPages` is the number of pages that are compressed. For example, if `Pages` is 1000 and `CompressedPages` is 992, this means that 992 of the 1000 pages are compressed. `CompressedPages` divided by `Pages` is usually near 100%, because most pages compress. An empty page is not compressed, since Sybase IQ does not write empty pages. IQ pages compress well, regardless of the fullness of the page.

Example

```
call sp_iqtablesize ('dba.empl')
```

Owner-name	Table-name	Col-umns	KBytes	Pages	Compressed-Pages	NBloc ks
DBA	empl	4	1504	24	14	188

sp_iqtableuse Procedure

Reports detailed usage information for tables accessed by the workload.

Syntax

sp_iqtableuse

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Tables created in SYSTEM are not reported.

Table 195. sp_iqtableuse columns

Column name	Description
TableName	Table name
Owner	User name of table owner
UID**	Table unique identifier
LastDT	Date/time of last access
NRef	Number of query references

**UID is a number assigned by the system that uniquely identifies the instance of the table (where instance is defined when an object is created).

See also

- *sp_iqcolumnuse Procedure* on page 385
- *sp_iqindexadvice Procedure* on page 437
- *sp_iqindexuse Procedure* on page 446
- *sp_iqunusedcolumn Procedure* on page 510
- *sp_iqunusedindex Procedure* on page 511
- *sp_iqunusedtable Procedure* on page 513
- *sp_iqworkmon Procedure* on page 521

sp_iqtransaction Procedure

Shows information about transactions and versions.

Syntax

sp_iqtransaction

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

sp_iqtransaction returns a row for each transaction control block in the Sybase IQ transaction manager. The columns Name, Userid, and ConnHandle are the connection properties **Name**, **Userid**, and **Number**, respectively. Rows are ordered by TxnID.

sp_iqtransaction output does not contain rows for connections that do not have a transaction started. To see all connections, use **sp_iqconnection**.

Note: Although you can use **sp_iqtransaction** to identify users who are blocking other users from writing to a table, **sp_iqlocks** is a better choice for this purpose.

You can use this procedure in a multiplex environment. See *Using Sybase IQ Multiplex*.

Table 196. sp_iqtransaction columns

Column name	Description
Name	The name of the application.
Userid	The user ID for the connection.
TxnID	The transaction ID of this transaction control block. The transaction ID is assigned during begin transaction. This is the same as the transaction ID displayed in the .iqmsg file by the BeginTxn, CmtTxn and PostCmtTxn messages as well as the Txn ID Seq logged when the database is opened.
CmtID	The ID assigned by the transaction manager when the transaction commits. It is zero for active transactions.
VersionID	In simplex databases, the VersionID is displayed as zero. For the multiplex coordinator, the VersionID is the same as the TxnID of the active transaction and VersionID is the same as the CmtID of a committed transaction. In multiplex secondary servers, the VersionID is the CmtID of the transaction that created the database version on the multiplex coordinator. It is used internally by the Sybase IQ in-memory catalog and the IQ transaction manager to uniquely identify a database version to all nodes within a multiplex database.

Column name	Description
State	<p>The state of the transaction control block. This variable reflects internal Sybase IQ implementation details and is subject to change in the future. At the time of this writing, transaction states are NONE, ACTIVE, ROLLING_BACK, ROLLED_BACK, COMMITTING, COMMITTED, and APPLIED.</p> <p>NONE, ROLLING_BACK, ROLLED_BACK, COMMITTING and APPLIED are transient states with a very small life-span.</p> <p>ACTIVE indicates that the transaction is active.</p> <p>COMMITTED indicates that the transaction has completed and is waiting to be APPLIED at which point a version that is not visible to any transaction will be garbage collected.</p> <p>Once the transaction state is ROLLED_BACK, COMMITTED, or APPLIED, it will cease to own any locks other than those held by open cursors.</p>
ConnHandle	The ID number of the connection.
IQConnID	The ten-digit connection ID displayed as part of all messages in the .iqmsg file. This is a monotonically increasing integer unique within a server session.
MainTableKBCr	The number of kilobytes of IQ store space created by this transaction.
MainTableKBDr	The number of kilobytes of IQ store space dropped by this transaction, but which persist on disk in the store because the space is visible in other database versions or other savepoints of this transaction.
TempTableKBCr	The number of kilobytes of IQ temporary store space created by this transaction for storage of IQ temporary table data.
TempTableKBDr	The number of kilobytes of IQ temporary table space dropped by this transaction, but which persist on disk in the IQ temporary store because the space is visible to IQ cursors or is owned by other savepoints of this transaction.
TempWorkSpaceKB	<p>For ACTIVE transactions, this is a snapshot of the work space in use at this instant by this transaction, such as sorts, hashes, and temporary bitmaps. The number varies depending on when you run sp_iqtransaction. For example, the query engine might create 60MB in the temporary cache but release most of it quickly, even though query processing continues. If you run sp_iqtransaction after the query finishes, this column shows a much smaller number. When the transaction is no longer active, this column is zero.</p> <p>For ACTIVE transactions, this column is the same as the TempWorkSpaceKB column of sp_iqconnection.</p>

Column name	Description
TxnCreateTime	The time the transaction began. All Sybase IQ transactions begin implicitly as soon as an active connection is established or when the previous transaction commits or rolls back.
CursorCount	The number of open Sybase IQ cursors that reference this transaction control block. If the transaction is ACTIVE, it indicates the number of open cursors created within the transaction. If the transaction is COMMITTED, it indicates the number of HOLD cursors that reference a database version owned by this transaction control block.
SpCount	The number of savepoint structures that exist within the transaction control block. Savepoints may be created and released implicitly. Therefore, this number does not indicate the number of user-created savepoints within the transaction.
SpNumber	The active savepoint number of the transaction. This is an implementation detail and might not reflect a user-created savepoint.
MPXServerName	The value indicates if an active transaction is from an inter-node communication (INC) connection. If from INC connection, the value is the name of the multiplex server where the transaction originates. NULL if not from an INC connection. Always NULL if the transaction is not active.
GlobalTxnID	The value indicates the global transaction ID associated with the current transaction. Zero if there is no associated global transaction.

Example

Here is an example of **sp_iqtransaction** output:

Name	Userid	TxnID	CmtID	VersionID	State	ConnHandle	IQConnID
red2	DBA	10058	10700	10058	COMMITTED	419740283	14
blue1	DBA	10568	0	10568	ACTIVE	640038605	17
	DBA	10604	0	10604	ACTIVE	2094200996	18
fromSCJ	DBA	10619	0	10619	ACTIVE	954498130	20
blue2	DBA	10634	10677	10634	COMMITTED	167015670	21
ntJava2	DBA	10676	0	10676	ACTIVE	1779741471	24
blue2	DBA	10678	0	10678	ACTIVE	167015670	21
nt1	DBA	10699	0	10699	ACTIVE	710225777	28
red2	DBA	10701	0	10701	ACTIVE	419740283	14
	DBA	16687	0	16687	ACTIVE	1306718536	23
MainTableKBCr		MainTableKBDr		TempTableKBCr		TempTableKBDr	
0		0		65824		0	
0		0		0		0	
0		0		0		0	

System Procedures

0	0	0	0
3960	152	0	0
0	0	0	0
2400	1992	0	0
0	0	0	0
0	0	2912	22096
0	0	0	0
TempWorkSpaceKB	TxnCreateTime	CursorCount	SpCount
SpNumber			
=====	=====	=====	=====
0	2009-06-26 13:17:27.612	1	3
102592	2009-06-26 13:27:28.491	1	1
0	2009-06-26 13:30:27.548	0	1
0	2009-06-26 13:31:27.151	0	24
0	2009-06-26 13:35:02.128	0	0
0	2009-06-26 13:43:58.805	0	39
128	2009-06-26 13:45:28.379	0	1
0	2009-06-26 14:05:15.759	0	42
680	2009-06-26 14:57:51.104	1	2
0	2009-06-26 15:09:30.319	0	1
MPXServerName	GlobalTxnID		
=====	=====		
(NULL)	0		
(NULL)	0		
(NULL)	0		
(NULL)	0		
(NULL)	0		
(NULL)	0		
(NULL)	0		
(NULL)	0		
(NULL)	0		
(NULL)	0		

See also

- *sp_iqstatus Procedure* on page 491
- *sp_iqversionuse Procedure* on page 514

sp_iqunusedcolumn Procedure

Reports IQ columns that were not referenced by the workload.

Syntax

```
sp_iqunusedcolumn
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Columns from tables created in SYSTEM or local temporary tables are not reported.

Table 197. sp_iqunusedcolumn columns

Column name	Description
TableName	Table name
ColumnName	Column name
Owner	User name of column owner

Example

Sample output from the **sp_iqunusedcolumn** procedure:

TableName	ColumnName	Owner
SalesOrders	ID	GROUPO
SalesOrders	CustomerID	GROUPO
SalesOrders	OrderDate	GROUPO
SalesOrders	FinancialCode	GROUPO
SalesOrders	Region	GROUPO
SalesOrders	SalesRepresentative	GROUPO
SalesOrderItems	ID	GROUPO
SalesOrderItems	LineID	GROUPO
SalesOrderItems	ProductID	GROUPO
SalesOrderItems	Quantity	GROUPO
SalesOrderItems	ShipDate	GROUPO
Contacts	ID	GROUPO
Contacts	Surname	GROUPO
Contacts	GivenName	GROUPO ...

See also

- *sp_iqcolumnuse Procedure* on page 385
- *sp_iqindexadvice Procedure* on page 437
- *sp_iqindexuse Procedure* on page 446
- *sp_iqtableuse Procedure* on page 506
- *sp_iqunusedindex Procedure* on page 511
- *sp_iqunusedtable Procedure* on page 513
- *sp_iqworkmon Procedure* on page 521

sp_iqunusedindex Procedure

Reports IQ secondary (non-FP) indexes that were not referenced by the workload.

Syntax

sp_iqunusedindex

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Indexes from tables created in SYSTEM or local temporary tables are not reported.

Table 198. sp_iqunusedindex columns

Column name	Description
IndexName	Index name
TableName	Table name
Owner	User name of index owner
IndexType	Index type

Example

Sample output from the **sp_iqunusedindex** procedure:

IndexName	TableName	Owner	IndexType
ASIQ_IDX_T450_I7_HG	SalesOrders	GROUPO	HG
ASIQ_IDX_T450_C6_HG	SalesOrders	GROUPO	HG
ASIQ_IDX_T450_C4_HG	SalesOrders	GROUPO	HG
ASIQ_IDX_T450_C2_HG	SalesOrders	GROUPO	HG
ASIQ_IDX_T451_I6_HG	SalesOrderItems	GROUPO	HG
ASIQ_IDX_T451_C3_HG	SalesOrderItems	GROUPO	HG
ASIQ_IDX_T451_C1_HG	SalesOrderItems	GROUPO	HG
ASIQ_IDX_T452_I11_HG	Contacts	GROUPO	HG
ASIQ_IDX_T453_I10_HG	Contacts	GROUPO	HG
ASIQ_IDX_T454_I4_HG	FinancialCodes	GROUPO	HG
ASIQ_IDX_T455_I5_HG	FinancialData	GROUPO	HG
ASIQ_IDX_T455_C3_HG	FinancialData	GROUPO	HG
ASIQ_IDX_T456_I8_HG	Products	GROUPO	HG
ASIQ_IDX_T457_I4_HG	Departments	GROUPO	HG
ASIQ_IDX_T457_C3_HG	Departments	GROUPO	HG
ASIQ_IDX_T458_I21_HG	Departments	GROUPO	HG
ASIQ_IDX_T458_C5_HG	Departments	GROUPO	HG

See also

- *sp_iqcolumnuse Procedure* on page 385
- *sp_iqindexadvice Procedure* on page 437
- *sp_iqindexuse Procedure* on page 446

- *sp_iqtableuse Procedure* on page 506
- *sp_iqunusedcolumn Procedure* on page 510
- *sp_iqunusedtable Procedure* on page 513
- *sp_iqworkmon Procedure* on page 521

sp_iqunusedtable Procedure

Reports IQ tables that were not referenced by the workload.

Syntax

sp_iqunusedtable

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

Tables created in SYSTEM and local temporary tables are not reported.

Table 199. sp_iqunusedtable columns

Column name	Description
TableName	Table name
Owner	User name of table owner

Example

The following table illustrates sample output from the **sp_iqunusedtable** procedure.

TableName	Owner	FinancialCodes	GROUPO
Contacts	GROUPO	Employees	GROUPO
empl	DBA	SalesOrders	GROUPO
FinancialData	GROUPO	Departments	GROUPO
SalesOrderItems	GROUPO	Products	GROUP
iq_dummy	DBA	Customers	GROUPO
sale	DBA		

See also

- *sp_iqcolumnuse Procedure* on page 385
- *sp_iqindexadvice Procedure* on page 437
- *sp_iqindexuse Procedure* on page 446
- *sp_iqtableuse Procedure* on page 506
- *sp_iqunusedcolumn Procedure* on page 510
- *sp_iqunusedindex Procedure* on page 511

- *sp_iqworkmon Procedure* on page 521

sp_iqversionuse Procedure

Displays version usage for the IQ main store.

Syntax

sp_iqversionuse

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

The **sp_iqversionuse** system stored procedure helps troubleshoot situations where the databases uses excessive storage space due to multiple table versions.

If out-of-space conditions occur or **sp_iqstatus** shows a high percentage of main blocks in use on a multiplex server, run **sp_iqversionuse** to find out which versions are being used and the amount of space that can be recovered by releasing versions.

The amount of space is expressed as a range because the actual amount typically depends on which other versions are released. The actual amount of space released can be anywhere between the values of MinKBRelease and MaxKBRelease. The oldest version always has MinKBRelease equal to MaxKBRelease.

WasReported indicates whether version usage information has been sent from the secondary server to the coordinator. WasReported is 0 initially on a coordinator for new versions. WasReported changes to 1 once the database server replicates version usage information back to the coordinator.

Note: The WasReported column is used in a multiplex setting. For more information on multiplex, see *Using Sybase IQ Multiplex*.

Table 200. sp_iqversionuse columns

Column name	Description
VersionID	In simplex databases, the VersionID is displayed as zero. For the multiplex coordinator, the VersionID is the same as the TxnID of the active transaction and VersionID is the same as the CmtID of a committed transaction. In multiplex secondary servers, the VersionID is the CmtID of the transaction that created the database version on the multiplex coordinator. It is used internally by the Sybase IQ in-memory catalog and the IQ transaction manager to uniquely identify a database version to all nodes within a multiplex database.

Column name	Description
Server	The server to which users of this version are connected
IQConnID	The connection ID using this version
WasReported	Indicates whether the server has received usage information for this version
MinKBRelease	The minimum amount of space returned once this version is no longer in use
MaxKBRelease	The maximum amount of space returned once this version is no longer in use

Example

The following table illustrates sample output from the **sp_iqversionuse** system procedure:

```

VersionID  Server          IQConnID  WasReported
=====  =====
          0  ab2ab_iqdemo          9          0

MinKBRelease  MaxKBRelease
=====  =====
          0          0

```

See also

- *sp_iqstatus Procedure* on page 491
- *sp_iqtransaction Procedure* on page 506

sp_iqview Procedure

Displays information about views in a database.

Syntax1

```
sp_iqview ([view_name],[view_owner],[view_type])
```

Syntax2

```
sp_iqview [view_name='viewname' ],
[view_owner='viewowner' ],[view_type='viewtype' ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Usage: Syntax1

```
sp_iqview NULL,NULL,SYSTEM
```

If you do not specify either of the first two parameters, but specify the next parameter in the sequence, you must substitute NULL for the omitted parameters. For example: `sp_iqview NULL, NULL, SYSTEM` and `sp_iqview deptview, NULL, 'ALL'`.

Note: The *view_type* value ALL must be enclosed in single quotes in Syntax1.

Usage: Syntax2

The parameters can be specified in any order. Enclose them in single quotes.

The allowed values for the *view_type* parameter:

Table 201. sp_iqview view_type values

view_type value	Information displayed
SYSTEM	System views
ALL	User and system views
any other value	User views

Description

Specifying one of the parameters returns only the views with the specified view name or views that are owned by the specified user. Specifying more than one parameter filters the results by all of the parameters specified. Specifying no parameters returns all user views in a database.

Table 202. sp_iqview columns

Column name	Description
view_name	The name of the view
view_owner	The owner of the view
view_def	The view definition as specified in the CREATE VIEW statement
remarks	User comments added with the COMMENT statement

sp_iqview returns a view definition greater than 32K characters without truncation.

See also

- *sp_iqcolumn Procedure* on page 382
- *sp_iqconstraint Procedure* on page 390
- *sp_iqdatatype Procedure* on page 398
- *sp_iqevent Procedure* on page 420
- *sp_iqhelp Procedure* on page 426
- *sp_iqindex and sp_iqindex_alt Procedures* on page 433
- *sp_iqjoinindex Procedure* on page 447

- *sp_iqpkkeys Procedure* on page 467
- *sp_iqprocparm Procedure* on page 472
- *sp_iq_reset_identity Procedure* on page 480
- *sp_iqtable Procedure* on page 501

sp_iqview Procedure Examples

Use the examples as a reference for **sp_iqview** usage.

The following variations in syntax both return information about the view deptview:

```
call sp_iqview('ViewSalesOrders')
```

```
sp_iqview view_name='ViewSalesOrders'
```

The following variations in syntax both return all views that are owned by view owner GROUPO:

```
sp_iqview NULL,GROUPO
```

```
sp_iqview view_owner='GROUPO'
```

view_name	view_owner	view_def	remarks
ViewSalesOrders	GROUPO	Create views GROUPO , ViewSalesOrders(ID, LineID, ProductID, Quantity, OrderDate, ShipDate, Region, SalesRepresentativeName	(NULL)

sp_iqwho Procedure

Displays information about all current users and connections, or about a particular user or connection.

Syntax

```
sp_iqwho [ { connhandle | user-name } [ , arg-type ] ]
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

Description

The **sp_iqwho** stored procedure displays information about all current users and connections, or about a particular user or connection.

Table 203. sp_iqwho columns

Column name	Description
ConnHandle	The SA connection handle

Column name	Description
IQConnID	The Sybase IQ specific connection ID
Userid	The name of the user that opened the connection “ConnHandle”
BlockedOn	The connection on which a particular connection is blocked; 0 if not blocked on any connection
BlockUserid	The owner of the blocking connection; NULL if there is no blocking connection
ReqType	The type of the request made through the connection; DO_NOTHING if no command is issued
IQCmdType	The type of Sybase IQ command issued from the connection; NONE if no command is issued
IQIdle	The time in seconds since the last Sybase IQ command was issued through the connection; in case of no last Sybase IQ command, the time since ‘01-01-2000’ is displayed
SAIdle	The time in seconds since the last SA request was issued through the connection; in case of no last SA command, the time since ‘01-01-2000’ is displayed
IQCursors	The number of active cursors in the connection; 0 if no cursors
IQThreads	The number of threads with the connection. At least one thread is started as soon as the connection is opened, so the minimum value for IQThreads is 1.
TempTableSpaceKB	The size of temporary table space in kilobytes; 0 if no temporary table space is used
TempWorkSpaceKB	The size of temporary workspace in kilobytes; 0 if no temporary workspace is used

Table 204. Mapping of sp_who and sp_iqwho columns

sp_who column	sp_iqwho column
fid	Family to which a lock belongs; omitted, as not applicable to Sybase IQ
spid	ConnHandle, IQConnID
status	IQIdle, SAIdle
loginame	Userid
origname	User alias; omitted, as not applicable to Sybase IQ

sp_who column	sp_iqwho column
hostname	Name of the host on which the server is running; currently not supported
blk_spid	BlockedOn
dbname	Omitted, as there is one server and one database for Sybase IQ and they are the same for every connection
cmd	ReqType, IQCmdType
block_xloid	BlockUserid

Usage

Table 205. Parameters

Parameter	Description
connhandle	An integer representing the connection ID. If this parameter is specified, sp_iqwho returns information only about the specified connection. If the specified connection is not open, no rows are displayed in the output.
user-name	A char(255) parameter representing a user login name. If this parameter is specified, sp_iqwho returns information only about the specified user. If the specified user has not opened any connections, no rows are displayed in the output. If the specified user name does not exist in the database, sp_iqwho returns the error message "User <i>user-name</i> does not exist"
arg-type	The <i>arg-type</i> parameter is optional and can be specified only when the first parameter has been specified. The only value for <i>arg-type</i> is "user". If the <i>arg-type</i> value is specified as "user", sp_iqwho interprets the first parameter as a user name, even if the first parameter is numeric. If any value other than "user" is specified for <i>arg-type</i> , sp_iqwho returns the error "Invalid parameter." Enclose the <i>arg-type</i> value in double quotes.

If no parameters are specified, **sp_iqwho** displays information about all currently active connections and users.

Either a connection handle or a user name can be specified as the first **sp_iqwho** parameter. The parameters *connhandle* and *user-name* are exclusive and optional. Only one of these parameters can be specified at a time. By default, if the first parameter is numeric, the parameter is assumed to be a connection handle. If the first parameter is not numeric, it is assumed to be a user name.

Sybase IQ allows numeric user names. The *arg-type* parameter directs **sp_iqwho** to interpret a numeric value in the first parameter as a user name. For example:

```
sp_iqwho 1, "user"
```

When the *arg-type* "user" is specified, **sp_iqwho** interprets the first parameter 1 as a user name, not as a connection ID. If a user named 1 exists in the database, **sp_iqwho** displays information about connections opened by user 1.

Table 206. sp_iqwho usage examples

Syntax	Output
sp_iqwho	Displays all active connections
sp_iqwho 3	Displays information about connection 3
sp_iqwho "DBA"	Displays connections opened by user DBA
sp_iqwho 3, "user"	Interprets 3 as a user name and displays connections opened by user 3. If user 3 does not exist, returns the error "User 3 does not exist"
sp_iqwho non-existing-user	Returns error "User non-existing-user does not exist"
sp_iqwho 3, "xyz"	Returns the error "Invalid parameter: xyz"

sp_iqwho Procedure Example

Use the example as a reference for **sp_iqwho** usage.

Display all active connections:

ConnHandle	IQConnID	Userid	ReqType	IQCmdType	Blocked
On	BlockUserid	IQCursors			
IQThreads	IQIdle	SAIdle	TempTableSpaceKB	TempWorkSpaceKB	
12	118	DBA	CURSOR_OPEN	IQUTILITYOPENCURSOR	0
(NULL)	0				
1	1	0	0	0	
13	119	shweta	DO_NOTHING	NONE	0
(NULL)	0				
1	16238757	470	0	0	

sp_iqwho compatibility with Adaptive Server Enterprise

The Sybase IQ **sp_iqwho** stored procedure incorporates the Sybase IQ equivalents of columns displayed by the Adaptive Server Enterprise **sp_who** procedure.

Some Adaptive Server Enterprise columns are omitted, as they are not applicable to Sybase IQ.

sp_iqworkmon Procedure

Controls collection of workload monitor usage information, and reports monitoring collection status. **sp_iqworkmon** collects information only for queries (SQL statements containing a **FROM** clause). You cannot use **sp_iqworkmon** for **INSERT** or **LOAD** statements.

Syntax

```
sp_iqworkmon [ 'action' ] [ , 'mode' ]
```

```
action = 'start' , 'stop' , 'status' , 'reset'
```

```
mode = 'index' , 'table' , 'column' , 'all'
```

For example:

```
sp_iqworkmon 'start' , 'all'
```

If one argument is specified, it can only be *action*. For example:

```
sp_iqworkmon 'stop'
```

Permissions

DBA authority required. Users without DBA authority must be granted EXECUTE permission to run the stored procedure.

*Usage***Table 207. Parameters**

Parameter	Description
action	<p>Specifies the control action to apply. A value of <i>start</i> starts monitoring for the specified mode immediately. A value of <i>stop</i> stops monitoring immediately. A value of <i>status</i> (the default) displays the current status without changing state.</p> <p>The statistics are persisted until they are cleared with the <i>reset</i> argument, or until the server is restarted. Statistics collection does not automatically resume after a server restart, and it needs to be restarted using <i>start</i>.</p>

Parameter	Description
mode	Specifies the type of monitoring to control. The INDEX, TABLE, and COLUMN keywords individually control monitoring of index usage, table usage, and column usage respectively. The default ALL keyword controls monitoring of all usage monitoring features simultaneously.

There is always a result set when you execute **sp_iqworkmon**. If you specify a specific mode (such as index), only the row for that mode appears.

Usage is collected only for SQL statements containing a **FROM** clause; for example, **SELECT**, **UPDATE**, and **DELETE**.

Table 208. sp_iqworkmon columns

Column name	Description
MonMode	Table, index, or column
Status	Started or stopped
Rowcount	Current number of rows collected

Example

Sample output from the **sp_iqworkmon** procedure:

```
MonMode      Status      Rowcount index      started      15
table        started    10 column      started      31
```

See also

- *sp_iqcolumnuse Procedure* on page 385
- *sp_iqindexadvice Procedure* on page 437
- *sp_iqindexuse Procedure* on page 446
- *sp_iqtableuse Procedure* on page 506
- *sp_iqunusedcolumn Procedure* on page 510
- *sp_iqunusedindex Procedure* on page 511
- *sp_iqunusedtable Procedure* on page 513

Catalog Stored Procedures

Catalog store stored procedures return result sets displaying database server, database, and connection properties in tabular form.

These procedures are owned by the dbo user ID. The PUBLIC group has EXECUTE permission on them.

sa_ansi_standard_packages system procedure

Returns information about the non-core SQL extensions used in a SQL statement.

Syntax

```
sa_ansi_standard_packages( standard, statement )
```

Arguments

- ***standard*** – The standard to use for the core extensions. One of SQL:1999 or SQL:2003.
- ***statement*** – The SQL statement to evaluate.

Remarks

If there are no non-core extensions used for the statement, the result set is empty.

Permissions

None

Side effects

None

Example

Following is an example call to the sa_ansi_standard_packages system procedure:

```
CALL sa_ansi_standard_packages( 'SQL:2003',
'SELECT *
    FROM ( SELECT o.SalesRepresentative,
                o.Region,
                SUM( s.Quantity * p.UnitPrice ) AS total_sales,
                DENSE_RANK() OVER ( PARTITION BY o.Region,
                                    GROUPING( o.SalesRepresentative )
                                    ORDER BY total_sales DESC ) AS
sales_rank
    FROM Product p, SalesOrderItems s, SalesOrders o
    WHERE p.ID = s.ProductID AND s.ID = o.ID
    GROUP BY GROUPING SETS( ( o.SalesRepresentative,
                              o.Region ), o.Region ) ) AS DT
    WHERE sales_rank <= 3
    ORDER BY Region, sales_rank');
```

System Procedures

The query generates the following result set:

package_id	package_name
T612	Advanced OLAP operations
T611	Elementary OLAP operations
F591	Derived tables
T431	Extended grouping capabilities

sa_audit_string system procedure

Adds a string to the transaction log.

Syntax

```
sa_audit_string( string )
```

Arguments

- **string** – A string of characters to add to the transaction log.

Remarks

If auditing is turned on, this system procedure adds a comment to the auditing information stored in the transaction log. The string can be a maximum of 200 bytes.

Permissions

DBA authority

Side effects

None

Example

The following example uses `sa_audit_string` to add a comment to the transaction log:

```
CALL sa_audit_string( 'Auditing test' );
```

sa_checkpoint_execute System Procedure

Allows the execution of shell commands during a checkpoint.

Syntax

```
sa_checkpoint_execute 'shell_commands'
```

Parameters

Parameter	Description
shell_commands	One or more user commands to be executed in a system shell. The shell commands are specific to the system shell. Commands are separated by a semicolon (;).

Permissions

None.

Description

Allows users to execute shell commands to copy a running database from the middle of a checkpoint operation, when the server is quiescent. The copied database can be started and goes through normal recovery, similar to recovery following a system failure.

sa_checkpoint_execute initiates a checkpoint, and then executes a system shell from the middle of the checkpoint, passing the user commands to the shell. The server then waits for the shell to complete, creating an arbitrary size time window during which to copy database files. Most database activity stops while the checkpoint is executing, so the duration of the shell commands should be limited to acceptable user response time.

If the shell commands return a nonzero status, **sa_checkpoint_execute** returns an error.

Do not use the **sa_checkpoint_execute** with interactive commands, as the server must wait until the interactive command is killed. Supply override flags to disable prompting for any shell commands that might become interactive; in other words, the **COPY**, **MOVE**, and **DELETE** commands might prompt for confirmation.

The intended use of **sa_checkpoint_execute** is with disk mirroring, to split mirrored devices.

When using `sa_checkpoint_execute` to copy `iqdemo.*` files to another directory, all files are copied except the `.db` and `.log` files. Error -910 is returned.

This error not a product defect but a Windows limitation; the Windows copy command cannot copy catalog files while they are open by the database.

Example

Assuming you have created a subdirectory named `backup`, the following statement issues a checkpoint, copies all of the **iqdemo** database files to the `backup` subdirectory, and completes the checkpoint:

```
sa_checkpoint_execute 'cp iqdemo.* backup/'
```

sa_conn_activity system procedure

Returns the most recently-prepared SQL statement for each connection to the specified database on the server.

Syntax

```
sa_conn_activity( [ connidparm ] )
```

Arguments

- **connidparm** – Use this optional INTEGER parameter to specify the connection ID number.

Result set

Column name	Data type	Description
Number	INT	Returns the ID number of the connection.
Name	VARCHAR(255)	Returns the name of the current connection. Temporary connection names have INT : prepended to the connection name.
Userid	VARCHAR(255)	Returns the user ID for the connection.
DBNumber	INT	Returns the ID number of the database.
LastReqTime	VARCHAR(255)	Returns the time at which the last request for the specified connection started. This property can return an empty string for internal connections, such as events.
LastStatement	LONG VARCHAR	Returns the most recently prepared SQL statement for the current connection.

Remarks

The sa_conn_activity system procedure returns a result set consisting of the most recently-prepared SQL statement for each connection, if the server has been told to collect the information. Recording of statements must be enabled for the database server before calling

sa_conn_activity. To do this, specify the -zl option when starting the database server, or execute the following:

```
CALL sa_server_option('RememberLastStatement','ON');
```

This procedure is useful when the database server is busy and you want to obtain information about the last SQL statement prepared for each connection. This feature can be used as an alternative to request logging.

If *connidparm* is not specified, then information is returned for all connections to all databases running on the database server. If *connidparm* is less than zero, option values for the current connection are returned.

Because of tenant database isolation rules, when this system procedure is run in the cloud it returns only information about the current tenant database.

Permissions

None

Side effects

None

sa_conn_info system procedure

Reports connection property information.

Syntax

```
sa_conn_info( [ connidparm ] )
```

Arguments

- *connidparm* – This optional INTEGER parameter specifies the connection ID number.

Result set

Column name	Data type	Description
Number	INTEGER	Returns the ID number of the connection.
Name	VARCHAR(255)	Returns the ID number of the connection. Temporary connection names have INT : prepended to the connection name.
Userid	VARCHAR(255)	Returns the user ID for the connection.

Column name	Data type	Description
DBNumber	INTEGER	Returns the ID number of the database.
LastReqTime	VARCHAR(255)	Returns the time at which the last request for the specified connection started. This property can return an empty string for internal connections, such as events.
ReqType	VARCHAR(255)	Returns the type of the last request. If a connection has been cached by connection pooling, its ReqType value is CONNECT_POOL_CACHE.
CommLink	VARCHAR(255)	Returns the communication link for the connection. This is one of the network protocols supported by SQL Anywhere, or local for a same-computer connection.
NodeAddr	VARCHAR(255)	Returns the address of the client in a client/server connection.
ClientPort	INTEGER	Returns the client's TCP/IP port number or 0 if the connection isn't a TCP/IP connection.
ServerPort	INTEGER	Returns the database server's TCP/IP port number or 0.
BlockedOn	INTEGER	Returns zero if the current connection isn't blocked, or if it is blocked, the connection number on which the connection is blocked because of a locking conflict.

Column name	Data type	Description
LockRowID	UNSIGNED BIGINT	Returns the identifier of the locked row. LockRowID is NULL if the connection is not waiting on a lock associated with a row (that is, it is not waiting on a lock, or it is waiting on a lock that has no associated row).
LockIndexID	INTEGER	Returns the identifier of the locked index. LockIndexID is -1 if the lock is associated with all indexes on the table in LockTable. LockIndexID is NULL if the connection is not waiting on a lock associated with an index (that is, it is not waiting on a lock, or it is waiting on a lock that has no associated index).
LockTable	VARCHAR(255)	Returns the name of the table associated with a lock if the connection is currently waiting for a lock. Otherwise, LockTable returns an empty string.
UncommitOps	INTEGER	Returns the number of uncommitted operations.
ParentConnection	INTEGER	Returns the connection ID of the connection that created a temporary connection to perform a database operation (such as performing a backup or creating a database). For other types of connections, this property returns NULL.

Remarks

If you specify the connection ID number, the `sa_conn_info` system procedure returns a result set consisting of connection properties for the supplied connection. If no `connidparm` is supplied, this system procedure returns information for all current connections to databases on

System Procedures

the server. If *connidparm* is less than zero, option values for the current connection are returned.

In a block situation, the BlockedOn value returned by this procedure allows you to check which users are blocked, and who they are blocked on. The sa_locks system procedure can be used to display the locks held by the blocking connection.

For more information based on any of these properties, you can execute something similar to the following:

```
SELECT *, DB_NAME( DBNumber ),  
        CONNECTION_PROPERTY( 'LastStatement', Number )  
FROM sa_conn_info( );
```

The value of LockRowID can be used to look up a lock in the output of the sa_locks procedure.

The value in LockIndexID can be used to look up a lock in the output of the sa_locks procedure. Also, the value in LockIndexID corresponds to the primary key of the ISYSIDX system table, which can be viewed using the SYSIDX system view.

Every lock has an associated table, so the value of LockTable can be used to unambiguously determine whether a connection is waiting on a lock.

Because of tenant database isolation rules, when this system procedure is run in the cloud it returns only information about the current tenant database.

Permissions

None

Side effects

None

Examples

The following example uses the sa_conn_info system procedure to return a result set summarizing connection properties for all connections to the server.

```
CALL sa_conn_info( );
```

Number	Name	Userid	DBNumber	...
79		DBA	0	...
46	Sybase Central 1	DBA	0	...
...

The following example uses the sa_conn_info system procedure to return a result set showing which connection created a temporary connection.

```
SELECT Number, Name, ParentConnection FROM sa_conn_info();
```

Connection 8 created the temporary connection that executed a CREATE DATABASE statement.

Number	Name	ParentConnection
1000000048	INT: CreateDB	8
9	SQL_DBC_14675af8	(NULL)
8	SQL_DBA_152d5ac0	(NULL)

sa_conn_list System Procedure

Returns a result set containing connection IDs.

Syntax

```
sa_conn_list ( [ connidparm ] [ ,dbidparm ] )
```

Arguments

- *connidparm* Use this optional INTEGER parameter to specify the connection ID number.
- *dbidparm* Use this optional INTEGER parameter to specify the database ID number.

Result Set

Column Name	Data Type	Description
Number	INTEGER	The connection ID number.

See Also

SQL Anywhere Server - SQL Reference > System procedures > Alphabetical list of system procedures > sa_conn_list system procedure

Note: This reference points to SQL Anywhere documentation.

sa_conn_properties System Procedure

Reports connection property information.

Syntax

```
sa_conn_properties ( [ connidparm ] )
```

See *SQL Anywhere Server - SQL Reference > System procedures > Alphabetical list of system procedures > sa_conn_properties system procedure.*

Note: This reference points to SQL Anywhere documentation.

sa_db_info system procedure

Reports database property information.

Syntax

```
sa_db_info( [ dbidparm ] )
```

Arguments

- **dbidparm** – Use this optional INTEGER parameter to specify the database ID number.

Result set

Column name	Data type	Description
Number	INTEGER	Returns the ID number of the connection.
Alias	VARCHAR(255)	Returns the database name.
File	VARCHAR(255)	Returns the file name of the database root file, including path.
ConnCount	INTEGER	Returns the number of connections to the database. The property value does not include connections used for internal operations, but it does include connections used for events and external environment support.
PageSize	INTEGER	Returns the page size of the database, in bytes.
LogName	VARCHAR(255)	Returns the file name of the transaction log, including path.

Remarks

If you specify a database ID, sa_db_info returns a single row containing the Number, Alias, File, ConnCount, PageSize, and LogName for the specified database.

If *dbidparm* is not supplied, properties for all databases are returned.

Because of tenant database isolation rules, when this system procedure is run in the cloud it returns only information about the current tenant database.

Permissions

None

Side effects

None

Example

The following statement returns a row for each database that is running on the server:

```
CALL sa_db_info( );
```

Property	Value
Number	0
Alias	demo
File	
ConnCount	1
PageSize	4096
LogName	C:\Documents and Settings\All Users\Program Data\SybaseIQ\demo\iqdemo.log

sa_db_properties system procedure

Reports database property information.

Syntax

```
sa_db_properties( [ dbidparm ] )
```

Arguments

- **dbidparm** – Use this optional INTEGER parameter to specify the database ID number.

Result set

Column name	Data type	Description
Number	INTEGER	The database ID number.
PropNum	INTEGER	The database property number.
PropName	VARCHAR(255)	The database property name.
PropDescription	VARCHAR(255)	The database property description.
Value	LONG VARCHAR	The database property value.

Remarks

If you specify a database ID, the `sa_db_properties` system procedure returns the database ID number and the `PropNum`, `PropName`, `PropDescription`, and `Value` for each available database property. Values are returned for all database properties and statistics related to databases. Valid properties with NULL values are also returned.

If `dbidparm` is not specified, properties for all databases are returned.

Because of tenant database isolation rules, when this system procedure is run in the cloud it returns only information about the current tenant database.

Permissions

None

Side effects

None

Example

The following example uses the `sa_db_properties` system procedure to return a result set summarizing database property information for all databases.

```
CALL sa_db_properties( );
```

Number	PropNum	PropName	...
0	0	ConnCount	...
0	1	IdleCheck	...
0	2	IdleWrite	...
...

sa_disable_auditing_type system procedure

Disables auditing of specific events.

Syntax

```
sa_disable_auditing_type(' types ')
```

Arguments

- **types** – Use this VARCHAR(128) parameter to specify a comma-delimited string containing one or more of the following values:
 - **all** – disables all types of auditing.
 - **connect** – disables auditing of both successful and failed connection attempts.

- **connectFailed** – disables auditing of failed connection attempts.
- **DDL** – disables auditing of DDL statements.
- **options** – disables auditing of public options.
- **permission** – disables auditing of permission checks, user checks, and SETUSER statements.
- **permissionDenied** – disables auditing of failed permission and user checks.
- **triggers** – disables auditing in response to trigger events.

Remarks

You can use the `sa_disable_auditing_type` system procedure to disable auditing of one or more categories of information.

Setting this option to all disables all auditing. You can also disable auditing by setting the `PUBLIC.auditing` option to Off.

Permissions

DBA authority

Side effects

None

Example

To disable all auditing:

```
CALL sa_disable_auditing_type( 'all' );
```

sa_disk_free_space system procedure

Reports information about space available for a dbspace, transaction log, transaction log mirror, and/or temporary file.

Syntax

```
sa_disk_free_space( [ p_dbspace_name ] )
```

Arguments

- **p_dbspace_name** – Use this VARCHAR(128) parameter to specify the name of a dbspace, transaction log file, transaction log mirror file, or temporary file.

If there is a dbspace called `log`, `mirror`, or `temp`, you can prefix the keyword with an underscore. For example, use `_log` to get information about the log file if a dbspace called `log` exists.

Specify `SYSTEM` to get information about the main database file, `TEMPORARY` or `TEMP` to get information about the temporary file, `TRANSLOG` to get information about the transaction log, or `TRANSLOGMIRROR` to get information about the transaction log mirror.

Result set

Column name	Data type	Description
dbspace_name	VARCHAR(128)	This is the dbspace name, transaction log file, transaction log mirror file, or temporary file.
free_space	UNSIGNED BIGINT	The number of free bytes on the volume.
total_space	UNSIGNED BIGINT	The total amount of disk space available on the drive where the dbspace resides.

Remarks

If the *p_dbspace_name* parameter is not specified or is NULL, then the result set contains one row for each dbspace, plus one row for each of the transaction log, transaction log mirror, and temporary file, if they exist. If *p_dbspace_name* is specified, then exactly one or zero rows are returned (zero if no such dbspace exists, or if log or mirror is specified and there is no log or mirror file).

Permissions

DBA authority

Side effects

None

Example

The following example uses the `sa_disk_free_space` system procedure to return a result set containing information about available space.

```
CALL sa_disk_free_space( );
```

dbspace_name	free_space	total_space
system	10952101888	21410402304
translog	10952101888	21410402304
temporary	10952101888	21410402304

sa_enable_auditing_type system procedure

Enables auditing and specifies which events to audit.

Syntax

```
sa_enable_auditing_type( 'types' )
```

Arguments

- **types** – Use this VARCHAR(128) parameter to specify a comma-delimited string containing one or more of the following values:
 - **all** – enables all types of auditing.
 - **connect** – enables auditing of both successful and failed connection attempts.
 - **connectFailed** – enables auditing of failed connection attempts.
 - **DDL** – enables auditing of DDL statements.
 - **options** – enables auditing of public options.
 - **permission** – enables auditing of permission checks, user checks, and SETUSER statements.
 - **permissionDenied** – enables auditing of failed permission and user checks.
 - **triggers** – enables auditing after a trigger event.

Remarks

sa_enable_auditing_type works in conjunction with the PUBLIC.auditing option to enable auditing of specific types of information.

If you set the PUBLIC.auditing option to On, and do not specify which type of information to audit, the default setting (all) takes effect. In this case, all types of auditing information are recorded.

If you set the PUBLIC.auditing option to On, and disable all types of auditing using sa_disable_auditing_type, no auditing information is recorded. To re-establish auditing, you must use sa_enable_auditing_type to specify which type of information you want to audit.

If you set the PUBLIC.auditing option to Off, then no auditing information is recorded, regardless of the sa_enable_auditing_type setting.

Permissions

DBA authority

Side effects

None

Example

To enable only option auditing:

```
CALL sa_enable_auditing_type( 'options' );
```

sa_eng_properties system procedure

Reports database server property information.

Syntax

```
sa_eng_properties( )
```

Result set

Column name	Data type	Description
PropNum	INTEGER	The database server property number.
PropName	VARCHAR(255)	The database server property name.
PropDescription	VARCHAR(255)	The database server property description.
Value	LONG VARCHAR	The database server property value.

Remarks

Returns the PropNum, PropName, PropDescription, and Value for each available server property. Values are returned for all database server properties and statistics related to database servers.

Permissions

None

Side effects

None

Example

The following statement returns a set of available server properties

```
CALL sa_eng_properties( );
```

PropNum	PropName	...
1	IdleWrite	...
2	IdleChkPt	...
...

sa_flush_cache system procedure

Empties all pages for the current database in the database server cache.

Syntax

```
sa_flush_cache( )
```

Remarks

Database administrators can use this procedure to empty the contents of the database server cache for the current database. This is useful in performance measurement to ensure repeatable results.

Permissions

DBA authority

Side effects

None

sa_make_object system procedure

Ensures that a skeletal instance of an object exists before executing an ALTER statement.

Syntax

```
sa_make_object(  
    objtype,  
    objname  
    [, owner  
    [, tabname ] ]  
)
```

objtype:

```
'procedure'  
| 'function'  
| 'view'  
| 'trigger'  
| 'service'  
| 'event'
```

Arguments

- **objtype** – Use this CHAR(30) parameter to specify the type of object being created. If objtype is 'trigger', this argument specifies the owner of the table on which the trigger is to be created.
- **objname** – Use this CHAR(128) parameter to specify the name of the object to be created.
- **owner** – Use this optional CHAR(128) parameter to specify the owner of the object to be created. The default value is CURRENT USER.

System Procedures

- **tablename** – This CHAR(128) parameter is required only if objtype is 'trigger', in which case you use it to specify the name of the table on which the trigger is to be created.

Remarks

This procedure is useful in scripts that are run repeatedly to create or modify a database schema. A common problem in such scripts is that the first time they are run, a CREATE statement must be executed, but subsequent times an ALTER statement must be executed. This procedure avoids the necessity of querying the system views to find out whether the object exists.

To use the procedure, follow it by an ALTER statement that contains the entire object definition.

Permissions

Resource authority is required to create or modify database objects

Side effects

Automatic commit

Examples

The following statements ensure that a skeleton procedure definition is created, define the procedure, and grant permissions on it. A script file containing these instructions could be run repeatedly against a database without error.

```
CALL sa_make_object( 'procedure','myproc' );
ALTER PROCEDURE myproc( in p1 INT, in p2 CHAR(30) )
BEGIN
    // ...
END;
GRANT EXECUTE ON myproc TO public;
```

The following example uses the sa_make_object system procedure to add a skeleton web service.

```
CALL sa_make_object( 'service','my_web_service' );
```

sa_rowgenerator system procedure

Returns a result set with rows between a specified start and end value.

Syntax

```
sa_rowgenerator(
  [ rstart
  [, rend
  [, rstep ] ] ]
)
```

Arguments

- **rstart** – Use this optional INTEGER parameter to specify the starting value. The default value is 0.
- **rend** – Use this optional INTEGER parameter to specify the ending value that is greater than or equal to *rstart*. The default value is 100.
- **rstep** – Use this optional INTEGER parameter to specify the increment by which the sequence values are increased. The default value is 1.

Result set

Column name	Data type	Description
row_num	INTEGER	Sequence number.

Remarks

The `sa_rowgenerator` procedure can be used in the FROM clause of a query to generate a sequence of numbers. This procedure is an alternative to using the RowGenerator system table. You can use `sa_rowgenerator` for such tasks as:

- generating test data for a known number of rows in a result set.
- generating a result set with rows for values in every range. For example, you can generate a row for every day of the month, or you can generate ranges of zip codes.
- generating a query that has a specified number of rows in the result set. This may be useful for testing the performance of queries.

No rows are returned if you do not specify correct start and end values and a positive non-zero step value.

You can emulate the behavior of the RowGenerator table with the following statement:

```
SELECT row_num FROM sa_rowgenerator( 1, 255 );
```

Permissions

None

Side effects

None

Example

The following query returns a result set containing one row for each day of the current month.

```
SELECT DATEADD( day, row_num-1,
              YMD( DATEPART( year, CURRENT DATE ),
                  DATEPART( month, CURRENT DATE ), 1 ) )
   AS day_of_month
FROM sa_rowgenerator( 1, 31, 1 )
WHERE DATEPART( month, day_of_month ) =
```

System Procedures

```
DATEPART( month, CURRENT DATE )
ORDER BY row_num;
```

The following query shows how many employees live in zip code ranges (0-9999), (10000-19999), ..., (90000-99999). Some of these ranges have no employees, which causes a warning.

The `sa_rowgenerator` procedure can be used to generate these ranges, even though no employees have a zip code in the range.

```
SELECT row_num AS r1, row_num+9999
       AS r2, COUNT( PostalCode ) AS zips_in_range
FROM sa_rowgenerator( 0, 99999, 10000 ) D LEFT JOIN Employees
     ON PostalCode BETWEEN r1 AND r2
GROUP BY r1, r2
ORDER BY 1;
```

The following example generates 10 rows of data and inserts them into the `NewEmployees` table:

```
INSERT INTO NewEmployees ( ID, Salary, Name )
SELECT row_num,
       CAST( RAND() * 1000 AS INTEGER ),
       'Mary'
FROM sa_rowgenerator( 1, 10 );
```

The following example uses the `sa_rowgenerator` system procedure to create a view containing all integers. The value 2147483647 in this example represents the maximum signed integer that is supported.

```
CREATE VIEW Integers AS
SELECT row_num AS n
FROM sa_rowgenerator( 0, 2147483647, 1 );
```

This example uses the `sa_rowgenerator` system procedure to create a view containing dates from 0001-01-01 to 9999-12-31. The value 3652058 in this example represents the number of days between 0001-01-01 and 9999-12-31, the earliest and latest dates that are supported.

```
CREATE VIEW Dates AS
SELECT DATEADD( day, row_num, '0001-01-01' ) AS d
FROM sa_rowgenerator( 0, 3652058, 1 );
```

The following query returns all years between 1900 and 2058 that have 54 weeks.

```
SELECT DATEADD ( day, row_num, '1900-01-01' ) AS d, DATEPART ( week,
d ) w
FROM sa_rowgenerator ( 0, 63919, 1 )
WHERE w = 54;
```


sa_server_option system procedure

Overrides a server option while the server is running.

Syntax

```
sa_server_option(
opt,
val
)
```

Arguments

- **opt** – Use this CHAR(128) parameter to specify a server option name.
- **val** – Use this CHAR(128) parameter to specify the new value for the server option.

Remarks

Database administrators can use this procedure to override some database server options temporarily, without restarting the database server.

The option values that are changed using this procedure are reset to their default values when the server shuts down. If you want to change an option value every time the server is started, you can specify the corresponding database server option when the database server is started if one exists (these are listed in the rightmost column in the table below).

The following option settings can be changed:

Option name	Values	Default
AutoMultiProgrammingLevel	YES, NO	YES
AutoMultiProgrammingLevel-Statistics	YES, NO	NO
CacheSizingStatistics	YES, NO	NO
CollectStatistics	YES, NO	YES
ConnsDisabled	YES, NO	NO
ConnsDisabledForDB	YES, NO	NO
ConsoleLogFile	<i>filename</i>	
ConsoleLogMaxSize	<i>file-size</i> , in bytes	
CurrentMultiProgrammingLevel	Integer	20
DatabaseCleaner	ON, OFF	ON

System Procedures

Option name	Values	Default
DeadlockLogging	ON, OFF, RESET, CLEAR	OFF
DebuggingInformation	YES, NO	NO
DropBadStatistics	YES, NO	YES
DropUnusedStatistics	YES, NO	YES
IdleTimeout	Integer, in minutes	240
IPAddressMonitorPeriod	Integer, in seconds	120 for portable devices, 0 otherwise
LivenessTimeout	Integer, in seconds	120
MaxMultiProgrammingLevel	Integer	Four times the CurrentMultiProgrammingLevel value
MessageCategoryLimit	Integer	400
MinMultiProgrammingLevel	Integer	The minimum of the value of the -gtc server option and the number of logical CPUs on the computer
OptionWatchAction	MESSAGE, ERROR	MESSAGE
OptionWatchList	comma-separated list of database options	
ProcedureProfiling	YES, NO, RESET, CLEAR	NO
ProfileFilterConn	<i>connection-id</i>	
ProfileFilterUser	<i>user-id</i>	
QuittingTime	valid date and time	
RememberLastPlan	YES, NO	NO
RememberLastStatement	YES, NO	NO
RequestFilterConn	<i>connection-id, -1</i>	
RequestFilterDB	<i>database-id, -1</i>	
RequestLogFile	<i>filename</i>	

Option name	Values	Default
RequestLogging	SQL, HOSTVARS, PLAN, PROCEDURES, TRIGGERS, OTHER, BLOCKS, REPLACE, ALL, YES, NONE, NO	NONE
RequestLogMaxSize	<i>file-size</i> , in bytes	
RequestLogNumFiles	Integer	
RequestTiming	YES, NO	NO
SecureFeatures	<i>feature-list</i>	
StatisticsCleaner	ON, OFF	ON
WebClientLogFile	<i>filename</i>	
WebClientLogging	ON, OFF	OFF

- **AutoMultiProgrammingLevel** – When set to YES, the database server automatically adjusts its multiprogramming level, which controls the maximum number of tasks that can be active at a time. If you choose to control the multiprogramming level manually by setting this option to NO, you can still set the initial, minimum, and maximum values for the multiprogramming level.
- **AutoMultiProgrammingLevelStatistics** – When set to YES, statistics for automatic multiprogramming level adjustments appear in the database server message log.
- **CacheSizingStatistics** – When set to YES, display cache information in the database server messages window whenever the cache size changes.
- **CollectStatistics** – When set to YES, the database server collects Performance Monitor statistics.
- **ConnsDisabled** – When set to YES, no other connections are allowed to any databases on the database server.
- **ConnsDisabledForDB** – When set to YES, no other connections are allowed to the current database.
- **ConsoleLogFile** – The name of the file used to record database server message log information. Specifying an empty string stops logging to the file. Any backslash characters in the path must be doubled because this is a SQL string.
- **ConsoleLogMaxSize** – The maximum size, in bytes, of the file used to record database server message log information. When the database server message log file reaches the size specified by either this property or the -on server option, the file is renamed with the extension `.old` appended (replacing an existing file with the same name if one exists). The database server message log file is then restarted.

- **CurrentMultiProgrammingLevel** – Sets the multiprogramming level of the database server.
- **DatabaseCleaner** – Do not change the setting of this option except on the recommendation of iAnywhere Technical Support.
- **DeadlockLogging** – Controls deadlock logging. The value `deadlock_logging` is also supported. Deadlock logging options can also be configured in the **Database Properties** window in Sybase Central. The following values are supported:
 - **ON** – Enables deadlock logging.
 - **OFF** – Disables deadlock logging and leaves the deadlock data available for viewing.
 - **RESET** – Clears the logged deadlock data, if any exists, and then enables deadlock logging.
 - **CLEAR** – Clears the logged deadlock data, if any exists, and then disables deadlock logging.

Once deadlock logging is enabled, you can use the `sa_report_deadlocks` system procedure to retrieve deadlock information from the database.

- **DebuggingInformation** – Displays diagnostic messages and other messages for troubleshooting purposes. The messages appear in the database server messages window.
- **DropBadStatistics** – Allows automatic statistics management to drop statistics that return bad estimates from the database.
- **DropUnusedStatistics** – Allows automatic statistics management to drop statistics that have not been used for 90 consecutive days from the database.
- **IdleTimeout** – Disconnects TCP/IP connections that have not submitted a request for the specified number of minutes. This prevents inactive connections from holding locks indefinitely.
- **IPAddressMonitorPeriod** – Sets the time to check for new IP addresses in seconds. The minimum value is 10 and the default is 0. For portable devices, the default value is 120 seconds.
- **LivenessTimeout** – A liveness packet is sent periodically across a client/server TCP/IP network to confirm that a connection is intact. If the network server runs for a `LivenessTimeout` period without detecting a liveness packet, the communication is severed.
- **MaxMultiProgrammingLevel** – Sets the maximum database server multiprogramming level.
- **MessageCategoryLimit** – Sets the minimum number of messages of each severity and category that can be retrieved using the `sa_server_messages` system procedure.
- **MinMultiProgrammingLevel** – Sets the minimum database server multiprogramming level.
- **OptionWatchAction** – Specifies the action the database server should take when an attempt is made to set an option in the list. The supported values are `MESSAGE` and `ERROR`. When `OptionWatchAction` is set to `MESSAGE`, and an option specified by `OptionWatchList` is set, a message appears in the database server messages window indicating that the option being set is on the options watch list.

When `OptionWatchAction` is set to `ERROR`, an error is returned indicating that the option cannot be set because it is on the options watch list.

You can view the current setting for this property by executing the following query:

```
SELECT DB_PROPERTY( 'OptionWatchAction' );
```

- **OptionWatchList** – Specifies a comma-separated list of database options that you want to be notified about, or have the database server return an error for, when they are set. The string length is limited to 128 bytes. By default, it is an empty string. For example, the following command adds the `automatic_timestamp`, `float_as_double`, and `tsql_hex_constant` option to the list of options being watched:

```
CALL sa_server_option( 'OptionWatchList', 'automatic_timestamp,  
float_as_double,tsql_hex_constant' )
```

You can view the current setting for this property by executing the following query:

```
SELECT DB_PROPERTY( 'OptionWatchList' );
```

- **ProcedureProfiling** – Controls procedure profiling for stored procedures, functions, events, and triggers. Procedure profiling shows you how long it takes your stored procedures, functions, events, and triggers to execute. You can also set procedure profiling options on the **Database Properties** window in Sybase Central
 - **YES** – enables procedure profiling for the database you are currently connected to.
 - **NO** – disables procedure profiling and leaves the profiling data available for viewing.
 - **RESET** – returns the profiling counters to zero, without changing the YES or NO setting.
 - **CLEAR** – returns the profiling counters to zero and disables procedure profiling.

Once profiling is enabled, you can use the `sa_procedure_profile_summary` and `sa_procedure_profile` system procedures to retrieve profiling information from the database.

- **ProfileFilterConn** – Instructs the database server to capture profiling information for a specific connection ID, without preventing other connections from using the database. When connection filtering is enabled, the value returned for `SELECT PROPERTY('ProfileFilterConn')` is the connection ID of the connection being monitored. If no ID has been specified, or if connection filtering is disabled, the value returned is -1.
- **ProfileFilterUser** – Instructs the database server to capture profiling information for a specific user ID.
- **QuittingTime** – Instructs the database server to shut down at the specified time.
- **RememberLastPlan** – Instructs the database server to capture the long text plan of the last query executed on the connection. This setting is also controlled by the `-zp` server option.

When `RememberLastPlan` is turned on, you can obtain the textual representation of the plan of the last query executed on the connection by querying the value of the `LastPlanText` connection property:

```
SELECT CONNECTION_PROPERTY( 'LastPlanText' );
```

- **RememberLastStatement** – Instructs the database server to capture the most recently prepared SQL statement for each database running on the server. For stored procedure calls, only the outermost procedure call appears, not the statements within the procedure.

When RememberLastStatement is turned on, you can obtain the current value of the LastStatement for a connection by querying the value of the LastStatement connection property:

```
SELECT CONNECTION_PROPERTY( 'LastStatement' );
```

When client statement caching is enabled, and a cached statement is reused, this property returns an empty string.

When RememberLastStatement is turned on, the following statement returns the most recently-prepared statement for the specified connection:

```
SELECT CONNECTION_PROPERTY( 'LastStatement', connection-id );
```

The sa_conn_activity system procedure returns this same information for all connections.

Warning! When -zl is specified, or when the RememberLastStatement server setting is turned on, any user can call the sa_conn_activity system procedure or obtain the value of the LastStatement connection property to find out the most recently-prepared SQL statement for any other user. This option should be used with caution and turned off when it is not required.

- **RequestFilterConn** – Filter the request logging information so that only information for a particular connection is logged. This can help reduce the size of the request log file when monitoring a database server with many active connections or multiple databases. You can obtain the connection ID by executing the following:

```
CALL sa_conn_info( );
```

To specify a specific connection to be logged once you have obtained the connection ID, execute the following:

```
CALL sa_server_option( 'RequestFilterConn', connection-id );
```

Filtering remains in effect until it is explicitly reset, or until the database server is shut down. To reset filtering, use the following statement:

```
CALL sa_server_option( 'RequestFilterConn', -1 );
```

- **RequestFilterDB** – Filter the request logging information so that only information for a particular database is logged. This can help reduce the size of the request log file when monitoring a server with multiple databases. You can obtain the database ID by executing the following statement when you are connected to the desired database:

```
SELECT CONNECTION_PROPERTY( 'DBNumber' );
```

To specify that only information for a particular database is to be logged, execute the following:

```
CALL sa_server_option( 'RequestFilterDB', database-id );
```

Filtering remains in effect until it is explicitly reset, or until the database server is shut down. To reset filtering, use the following statement:

```
CALL sa_server_option( 'RequestFilterDB', -1 );
```

- **RequestLogFile** – The name of the file used to record request information. Specifying an empty string stops logging to the request log file. If request logging is enabled, but the request log file was not specified or has been set to an empty string, the server logs requests to the database server messages window. Any backslash characters in the path must be doubled as this is a SQL string.

When client statement caching is enabled, the `max_client_statements_cached` option should be set to 0 to disable client statement caching while the request log is captured, if the log will be analyzed using the `tracetime.pl` Perl script.

- **RequestLogging** – This call turns on logging of individual SQL statements sent to the database server for use in troubleshooting, in conjunction with the database server `-zr` and `-zo` options. Values can be combinations of the following, separated by either a plus sign (+), or a comma:
 - **PLAN** – enables logging of execution plans (short form). Execution plans for procedures are also recorded if logging of procedures (PROCEDURES) is enabled.
 - **HOSTVARS** – enables logging of host variable values. If you specify HOSTVARS, the information listed for SQL is also logged.
 - **PROCEDURES** – enables logging of statements executed from within procedures.
 - **TRIGGERS** – enables logging of statements executed from within triggers.
 - **OTHER** – enables logging of additional request types not included by SQL, such as FETCH and PREFETCH. However, if you specify OTHER but do not specify SQL, it is the equivalent of specifying SQL+OTHER. Including OTHER can cause the log file to grow rapidly and could negatively impact server performance.
 - **BLOCKS** – enables logging of details showing when a connection is blocked and unblocked on another connection.
 - **REPLACE** – at the start of logging, the existing request log is replaced with a new (empty) one of the same name. Otherwise, the existing request log is opened and new entries are appended to the end of the file.
 - **ALL** – logs all supported information. This is equivalent to specifying SQL+PLAN+HOSTVARS+PROCEDURES+TRIGGERS+OTHER+BLOCKS. This setting can cause the log file to grow rapidly and could negatively impact server performance.
 - **NO or NONE** – turns off logging to the request log.

You can view the current setting for this property by executing the following query:

```
SELECT PROPERTY( 'RequestLogging' );
```

- **RequestLogMaxSize** – The maximum size of the file used to record request logging information, in bytes. If you specify 0, then there is no maximum size for the request logging file, and the file is never renamed. This is the default value.

When the request log file reaches the size specified by either the `sa_server_option` system procedure or the `-zs` server option, the file is renamed with the extension `.old` appended

(replacing an existing file with the same name if one exists). The request log file is then restarted.

- **RequestLogNumFiles** – The number of request log file copies to retain.

If request logging is enabled over a long period of time, the request log file can become large. The `-zn` option allows you to specify the number of request log file copies to retain.

- **RequestTiming** – Instructs the database server to maintain timing information for each connection. This feature is turned off by default. When it is turned on, the database server maintains cumulative timers for each connection that indicate how much time the connection spent in the server in each of several states. You can use the `sa_performance_diagnostics` system procedure to obtain a summary of this timing information, or you can retrieve individual values by inspecting the following connection properties:

- ReqCountUnscheduled
- ReqTimeUnscheduled
- ReqCountActive
- ReqTimeActive
- ReqCountBlockIO
- ReqTimeBlockIO
- ReqCountBlockLock
- ReqTimeBlockLock
- ReqCountBlockContention
- ReqTimeBlockContention

When the `RequestTiming` server property is on, there is a small overhead for each request to maintain the additional counters.

- **SecureFeatures** – Allows you to enable or disable secure features of a database server that is already running. The *feature-list* is a comma-separated list of feature names or feature sets. By adding a feature to the list, you limit its availability. To remove items from the list of secure features, specify a minus sign (-) before the secure feature name.

Any changes you make to enable or disable features take effect immediately for the connection. The settings do not affect the connection that executes the `sa_server_option` system procedure; you must disconnect and reconnect to see the change.

Note: To use the `sa_server_option` system procedure to enable or disable features, you must have specified a key with the `-sk` option when starting the database server, and set the value of the `secure_feature_key` database option to the key you specified for `-sk` (for example, `SET TEMPORARY OPTION secure_feature_key = 'j978k1s12'`). Setting the `secure_feature_key` database option to the `-sk` value allows you to change the setting for secure features.

For example, to disable two features and enable a third, you would use this syntax:

```
CALL sa_server_option('SecureFeatures',  
'CONSOLE_LOG,WEBCLIENT_LOG,-REQUEST_LOG');
```


After executing this statement, `CONSOLE_LOG`, and `WEBCLIENT_LOG` are added to the list of secure features, and `REQUEST_LOG` is removed from the list.

- **StatisticsCleaner** – The statistics cleaner fixes statistics that give bad estimates by performing scans on tables. By default the statistics cleaner runs in the background and has a minimal impact on performance.

Turning off the statistics cleaner does not disable the statistic governor, but when the statistics cleaner is turned off, statistics are only created or fixed when a query is run.

- **WebClientLogFile** – The name of the web service client log file. The web service client log file is truncated each time you use the `-zoc` server option or the `WebClientLogFile` property to set or reset the file name. Any backslash characters in the path must be doubled because this is a string.
- **WebClientLogging** – This option enables and disables logging of web service clients. The information that is logged includes HTTP requests and response data. Specify `ON` to start logging to the web service client log file, and specify `OFF` to stop logging to the file.

Permissions

The following options, which are related to application profiling or request logging, require either `DBA` or `PROFILE` authority:

ProcedureProfiling
 ProfileFilterConn
 ProfileFilterUser
 RequestFilterConn
 RequestFilterDB
 RequestLogFile
 RequestLogging
 RequestLogMaxSize
 RequestLogNumFiles

All other options require `DBA` authority.

Side effects

None

Example

The following statement disallows new connections to the database server:

```
CALL sa_server_option( 'ConnsDisabled', 'YES' );
```

The following statement disallows new connections to the current database:

```
CALL sa_server_option( 'ConnsDisabledForDB', 'YES' );
```

The following statement enables logging of all SQL statements, procedure calls, plans, blocking and unblocking events, and specifies that a new request log be started:

```
CALL sa_server_option( 'RequestLogging', 'SQL+PROCEDURES+BLOCKS+PLAN
+REPLACE' );
```

sa_set_http_header system procedure

Permits a web service to set an HTTP response header.

Syntax

```
sa_set_http_header(
fldname,
val
)
```

Arguments

- **fldname** – Use this CHAR(128) parameter to specify a string containing the name of one of the HTTP header fields.
- **val** – Use this LONG VARCHAR parameter to specify the value to which the named parameter should be set. Setting a response header to NULL, effectively removes it.

Remarks

Setting the special header field @HttpStatus sets the status code returned with the request. The status code is also known as the response code. For example, the following script sets the status code to 404 Not Found:

```
CALL sa_set_http_header( '@HttpStatus', '404' );
```

You can create a user-defined status message by specifying a three digit status code with an optional colon-delimited text message. For example, the following script outputs a status code with the message "999 User Code":

```
CALL sa_set_http_header( '@HttpStatus', '999:User Code' );
```

Note: A user defined status text message is not translated into a database character-set when logged using the LogOptions protocol option.

The body of the error message is inserted automatically. Only valid HTTP error codes can be used. Setting the status to an invalid code causes a SQL error.

The sa_set_http_header procedure always overwrites the existing header value of the header field when called.

Response headers generated automatically by the database server can be removed. For example, the following command removes the Expires response header:

```
CALL sa_set_http_header( 'Expires', NULL );
```

Permissions

None

Side effects

None

Example

The following example sets the Content-Type header field to text/html.

```
CALL sa_set_http_header( 'Content-Type', 'text/html' );
```

sa_set_http_option system procedure

Permits a web service to set an HTTP option for process control.

Syntax

```
sa_set_http_option(  
  optname,  
  val  
)
```

Arguments

- **optname** – Use this CHAR(128) parameter to specify a string containing the name of one of the HTTP options.

The supported options are:

- **CharsetConversion** – Use this option to control whether the result set is to be automatically converted from the character set encoding of the database to the character set encoding of the client. The only permitted values are ON and OFF. The default value is ON.
- **AcceptCharset** – Use this option to specify the web server's preferences for a response character set encoding. One or more character set encodings may be specified in order of preference. The syntax for this option conforms to the syntax used for the HTTP Accept-Charset request-header field specification in RFC2616 Hypertext Transfer Protocol.

An HTTP client such as a web browser may provide an Accept-Charset request header which specifies a list of character set encodings ordered by preference. Optionally, each encoding may be given an associated quality value (*q=quality*) which represents the client's preference for that encoding. By default, the quality value is 1 (*q=1*). Here is an example:

```
Accept-Charset: iso-8859-5, utf-8;q=0.8
```

A plus sign (+) in the AcceptCharset HTTP option value may be used as a shortcut to represent the current database character set encoding. The plus sign also indicates that the database character set encoding should take precedence if the client also specifies the encoding in its list, regardless of the quality value assigned by the client.

An asterisk (*) in the AcceptCharset HTTP option may be used to indicate that the web service should use a character set encoding preferred by the client, as long as it is also supported by the server, when client and server do not have an intersecting list.

When sending the response, the first character set encoding preferred by both client and web service is used. The client's order of preference takes precedence. If no mutual encoding preference exists, then the web service's most preferred encoding is used, unless an asterisk (*) appears in the web service list in which case the client's most preferred encoding is used.

If the AcceptCharset HTTP option is not used, the most preferred character set encoding specified by the client and supported by the server is used. If none of the encodings specified by the client are supported (or the client does not send an Accept-Charset request header) then the database character set encoding is used.

If a client does not send an Accept-Charset header then one of the following actions are taken:

- If the AcceptCharset HTTP option has not been specified then the web server will use the database character set encoding.
- If the AcceptCharset HTTP option has been specified then the web server will use its most preferred character set encoding.

If a client does send an Accept-Charset header then one of the following actions are taken:

- If the AcceptCharset HTTP option has not been specified then the web server will attempt to use one of the client's preferred character set encodings, starting with the most preferred encoding. If the web server does not support any of the client's preferred encodings, it will use the database character set encoding.
- If the AcceptCharset HTTP option has been specified then the web server will attempt to use the first preferred character set encoding common to both lists, starting with the client's most preferred encoding. For example, if the client sends an Accept-Charset header listing, in order of preference, encodings iso-a, iso-b, and iso-c and the web server prefers iso-b, then iso-a, and finally iso-c, then iso-a will be selected.

```
Web client: iso-a, iso-b, iso-c
Web server: iso-b, iso-a, iso-c
```

If the intersection of the two lists is empty, then the web server's first preferred character set is used. From the following example, encoding iso-d will be used.

```
Web client: iso-a, iso-b, iso-c
Web server: iso-d, iso-e, iso-f
```

If an asterisk (*) was included in the AcceptCharset HTTP option, then emphasis would be placed on the client's choice of encodings, resulting in iso-a being used. Essentially, the use of an asterisk guarantees that the intersection of the two lists will not be empty.

The ideal situation occurs when both client and web service use the database character set encoding since this eliminates the need for character set translation and improves the response time of the web server.

Note that if the `CharsetConversion` option has been set to `OFF`, then `AcceptCharset` processing is not performed.

- **SessionID** – Use this option to create, delete or rename an HTTP session. The database connection is persisted when a web service sets this option to create an HTTP session but sessions are not persisted across server restarts. If already within a session context, this call will rename the session to the new session ID. When called with a `NULL` value, the session will be deleted when the web service terminates.

The generated session keys are limited to 128 characters in length and unique across databases if multiple databases are loaded.

- **SessionTimeout** – Use this option to specify the amount of time, in minutes, that the HTTP session persists during inactivity. This time-out period is reset whenever an HTTP request uses the given session. The session is automatically deleted when the `SessionTimeout` is exceeded.
- *val* – Use this `LONG VARCHAR` parameter to specify the value to which the named option should be set.

Remarks

Use this procedure within statements or procedures that handle web services to set options.

When `sa_set_http_option` is called from within a procedure invoked through a web service, and either the option or option value is invalid, an error is returned.

Permissions

None

Side effects

None

Examples

The following example illustrates the use of `sa_set_http_option` to indicate the web service's preference for database character set encoding. The UTF-8 encoding is specified as a second choice. The asterisk (*) indicates that the web service is willing to use the character set encoding most preferred by the client, provided that it is supported by the web server.

```
CALL sa_set_http_option( 'AcceptCharset', '+,UTF-8,*');
```

The following example illustrates the use of `sa_set_http_option` to correctly identify the character encoding in use by the web service. In this example, the web server is connected to a 1251CYR database and is prepared to serve HTML documents containing the Cyrillic alphabet to any web browser.

```
CREATE PROCEDURE cyrillic_html()  
RESULT (html_doc XML)
```

System Procedures

```
BEGIN
  DECLARE pos INT;
  DECLARE charset VARCHAR(30);
  CALL sa_set_http_option( 'AcceptCharset', 'iso-8859-5, utf-8' );
  SET charset = CONNECTION_PROPERTY( 'CharSet' );
  -- Change any IANA labels like ISO_8859-5:1988
  -- to ISO_8859-5 for Firefox.
  SET pos = LOCATE( charset, ':' );
  IF pos > 0 THEN
    SET charset = LEFT( charset, pos - 1 );
  END IF;
  CALL sa_set_http_header( 'Content-Type', 'text/html; charset=' ||
    charset );
  SELECT '<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN">' ||
    XMLCONCAT(
      XMLELEMENT('HTML',
        XMLELEMENT('HEAD',
          XMLELEMENT('TITLE', 'Cyrillic characters')
        ),
        XMLELEMENT('BODY',
          XMLELEMENT('H1', 'First 5 lowercase Russian letters'),
          XMLELEMENT('P', UNISTR('\u0430\u0431\u0432\u0433\u0434'))
        )
      )
    );
END;
CREATE SERVICE cyrillic
TYPE 'RAW'
AUTHORIZATION OFF
USER DBA
AS CALL cyrillic_html();
```

To illustrate the process of establishing the correct character set encoding to use, consider the following Accept-Charset header delivered by a web browser such as Firefox to the web service. It indicates that the browser prefers ISO-8859-1 and UTF-8 encodings but is willing to accept others.

```
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
```

The web service will not accept the ISO-8859-1 character set encoding since the web page to be transmitted contains Cyrillic characters. The web service prefers ISO-8859-5 or UTF-8 encodings as indicated by the call to `sa_set_http_option`. In this example, the UTF-8 encoding will be chosen since it is agreeable to both parties. The database connection property 'CharSet' indicates which encoding has been selected by the web service. The `sa_set_http_header` procedure is used to indicate the HTML document's encoding to the web browser.

```
Content-Type: text/html; charset=UTF-8
```

If the web browser does not specify an Accept-Charset, then the web service defaults to its first preference, ISO-8859-5. The `sa_set_http_header` procedure is used to indicate the HTML document's encoding.

```
Content-Type: text/html; charset=ISO_8859-5
```

The following example sets a unique HTTP session identifier:

```

DECLARE sessionid VARCHAR(30);
DECLARE tm TIMESTAMP;
SET tm = NOW(*);
SET sessionid = 'MySessions_' ||
  CONVERT( VARCHAR, SECONDS(tm)*1000 + DATEPART(millisecond,tm));
CALL sa_set_http_option('SessionID', sessionid);

```

The following example sets the time-out for an HTTP session to 5 minutes:

```
CALL sa_set_http_option('SessionTimeout', '5');
```

sa_table_page_usage system procedure

Reports information about the page usage of database tables.

Syntax

```
sa_table_page_usage( )
```

Result set

Column name	Data type	Description
TableId	UNSIGNED INTEGER	The table ID.
TablePages	INTEGER	The number of table pages used by the table.
PctUsedT	INTEGER	The percentage of used table page space.
IndexPages	INTEGER	The number of index pages used by the table.
PctUsedI	INTEGER	The percentage of used index page space.
PctOfFile	INTEGER	The percentage of the total database file the table occupies.
TableName	CHAR(128)	The table name.

Remarks

The results include the same information provided by the Information utility. When the progress_messages database option is set to Raw or Formatted, progress messages are sent from the database server to the client while the sa_table_page_usage system procedure is running.

Permissions

DBA authority

Side effects

None

sa_validate system procedure

Validates all, or parts, of a database.

Syntax

```
sa_validate(
  [ tbl_name [ , owner_name ] ]
)
```

Arguments

- **tbl_name** – Use this optional VARCHAR(128) parameter to specify the name of a table or materialized view to validate.
- **owner_name** – Use this optional VARCHAR(128) parameter to specify an owner. When specified by itself, all tables and materialized views owned by the owner are validated.

Permissions

DBA authority

Side effects

None

Remarks

Argument specified	Type of validation
None	All tables, materialized views, and indexes in the database are validated. The database itself is also validated, including checksum validation.
<i>owner_name</i>	All tables, materialized views, and indexes owned by the specified user are validated.
<i>tbl_name</i>	The specified table, or materialized view, and all of its indexes, that are owned by the current user are validated.
<i>owner_name</i> and <i>tbl_name</i>	The specified table, or materialized view, and all of its indexes, that are owned by the specified user are validated.

The values for *tbl_name* and *owner_name* are strings and must be enclosed in quotes.

The procedure returns a single column named Messages. Errors returned during validation appear in the column. If validation succeeds without error, the column contains `No error detected`.

Warning! Validating a table or an entire database should be performed while no connections are making changes to the database; otherwise, errors may be reported indicating some form of database corruption even though no corruption actually exists.

Example

The following statement performs a validation of tables and materialized views owned by DBA:

```
CALL sa_validate( owner_name = 'DBA' );
```

sa_verify_password system procedure

Validates the password of the current user.

Syntax

```
sa_verify_password( curr_pswd )
```

Arguments

- **curr_pswd** – Use this CHAR(128) parameter to specify the password of the current database user.

Remarks

This procedure is used by `sp_password`. If the password matches, it is accepted. If the password does not match, an error is returned.

Permissions

None

Side effects

None

sp_login_environment system procedure

Sets connection options when users log in.

Syntax

```
sp_login_environment( )
```

Remarks

`sp_login_environment` is the default procedure called by the `login_procedure` database option.

It is recommended that you do not edit this procedure. Instead, to change the login environment, set the `login_procedure` option to point to a different procedure.

System Procedures

Here is the text of the `sp_login_environment` procedure:

```
CREATE PROCEDURE dbo.sp_login_environment( )
BEGIN
    IF connection_property( 'CommProtocol' ) = 'TDS' THEN
        CALL dbo.sp_tsql_environment( )
    END IF
END;
```

Permissions

None

Side effects

None

sp_remote_columns System Procedure

Produces a list of the columns in a remote table, and a description of their data types.

The server must be defined with the **CREATE SERVER** statement to use this system procedure.

Syntax

```
sp_remote_columns ( @servername , @tablename [ , @table_owner ] [ ,
@table_qualifier ] ] )
```

See *SQL Anywhere Server - SQL Reference > System procedures > Alphabetical list of system procedures > sp_remote_columns system procedure.*

Note: This reference points to SQL Anywhere documentation.

sp_remote_exported_keys system procedure

Provides information about tables with foreign keys on a specified primary table.

The server must be defined with the **CREATE SERVER** statement to use this system procedure.

Syntax

```
sp_remote_exported_keys(
    server_name
    , sp_name
    [ , sp_owner
    [ , sp_qualifier ] ]
)
```

Arguments

- **server_name** – Use this CHAR(128) parameter to specify the server the primary table is located on.

- *sp_name* – Use this CHAR(128) parameter to specify the table containing the primary key.
- *sp_owner* – Use this optional CHAR(128) parameter to specify the primary table's owner.
- *sp_qualifier* – Use this optional CHAR(128) parameter to specify the database containing the primary table.

Result set

Column name	Data type	Description
pk_database	CHAR(128)	The database containing the primary key table.
pk_owner	CHAR(128)	The owner of the primary key table.
pk_table	CHAR(128)	The primary key table.
pk_column	CHAR(128)	The name of the primary key column.
fk_database	CHAR(128)	The database containing the foreign key table.
fk_owner	CHAR(128)	The foreign key table's owner.
fk_table	CHAR(128)	The foreign key table.
fk_column	CHAR(128)	The name of the foreign key column.
key_seq	SMALLINT	The key sequence number.
fk_name	CHAR(128)	The foreign key name.
pk_name	CHAR(128)	The primary key name.

Remarks

This procedure provides information about the remote tables that have a foreign key on a particular primary table. The result set for the *sp_remote_exported_keys* system procedure includes the database, owner, table, column, and name for both the primary and the foreign key, and the foreign key sequence for the foreign key columns. The result set may vary because of the underlying ODBC and JDBC calls, but information about the table and column for a foreign key is always returned.

Permissions

None

Side effects

None

Example

This example uses named parameters when calling the system procedure, and returns information about the remote tables with foreign keys on the SYSOBJECTS table, in the production database, on a server named asetest:

```
CALL sp_remote_exported_keys(  
    @server_name='asetest',  
    @sp_name='sysobjects',  
    @sp_qualifier='production' );
```

sp_remote_imported_keys system procedure

Provides information about remote tables with primary keys that correspond to a specified foreign key.

The server must be defined with the CREATE SERVER statement to use this system procedure.

Syntax

```
sp_remote_imported_keys(  
server_name  
    , sp_name  
    [ , sp_owner  
    [ , sp_qualifier ] ]  
)
```

Arguments

- **server_name** – Use this optional CHAR(128) parameter to specify the server the foreign key table is located on. A value is required for this parameter.
- **sp_name** – Use this optional CHAR(128) parameter to specify the table containing the foreign key. A value is required for this parameter.
- **sp_owner** – Use this optional CHAR(128) parameter to specify the foreign key table's owner.
- **sp_qualifier** – Use this optional CHAR(128) parameter to specify the database containing the foreign key table.

Result set

Column name	Data type	Description
pk_database	CHAR(128)	The database containing the primary key table.

Column name	Data type	Description
pk_owner	CHAR(128)	The owner of the primary key table.
pk_table	CHAR(128)	The primary key table.
pk_column	CHAR(128)	The name of the primary key column.
fk_database	CHAR(128)	The database containing the foreign key table.
fk_owner	CHAR(128)	The foreign key table's owner.
fk_table	CHAR(128)	The foreign key table.
fk_column	CHAR(128)	The name of the foreign key column.
key_seq	SMALLINT	The key sequence number.
fk_name	CHAR(128)	The foreign key name.
pk_name	CHAR(128)	The primary key name.

Remarks

Foreign keys reference a row in a separate table that contains the corresponding primary key. This procedure allows you to obtain a list of the remote tables with primary keys that correspond to a particular foreign table. The `sp_remote_imported_keys` result set includes the database, owner, table, column, and name for both the primary and the foreign key, and the foreign key sequence for the foreign key columns. The result set may vary because of the underlying ODBC and JDBC calls, but information about the table and column for a primary key is always returned.

Permissions

None

Side effects

None

Example

The following example uses named parameters when calling the procedure, and returns the tables with primary keys that correspond to a foreign key on the `SYSOBJECTS` table in the `asetest` server:

```
CALL sp_remote_imported_keys(
    @server_name='asetest',
```

```
@sp_name='sysobjects',  
@sp_qualifier='production' );
```

sp_remote_primary_keys System Procedure

Provides primary key information about remote tables using remote data access.

Syntax

```
sp_remote_primary_keys( @server_name [, @table_name [, @table_owner  
[, @table_qualifier ] ] ] )
```

See *SQL Anywhere Server - SQL Reference > System procedures > Alphabetical list of system procedures > sp_remote_imported_keys system procedure.*

Note: This reference points to SQL Anywhere documentation.

sp_remote_tables system procedure

Returns a list of the tables on a server.

The server must be defined with the CREATE SERVER statement to use this system procedure.

Syntax

```
sp_remote_tables(  
server_name  
[, table_name  
[, table_owner  
[, table_qualifier  
[, with_table_type ] ] ] ]  
)
```

Arguments

- **server_name** – Use this CHAR(128) parameter to specify the server the remote table is located on.
- **table_name** – Use this optional CHAR(128) parameter to specify the remote table.
- **table_owner** – Use this optional CHAR(128) parameter to specify the owner of the remote table.
- **table_qualifier** – Use this optional CHAR(128) parameter to specify the database in which *table_name* is located.
- **with_table_type** – Use this optional BIT parameter to specify the type of remote table. This argument is a bit type and accepts two values, 0 (the default) and 1. You must enter the value 1 if you want the result set to include a column that lists table types.

Result set

Column name	Data type	Description
database	CHAR(128)	The name of the remote database.
owner	CHAR(128)	The name of the remote database owner.
table-name	CHAR(128)	The remote table.
table-type	CHAR(128)	Specifies the table type. The value depends on the type of remote server. For example, TABLE, VIEW, SYS, and GBL TEMP are possible values.

Remarks

It may be helpful when you are configuring your database server to get a list of the remote tables available on a particular server. This procedure returns a list of the tables on a server.

The procedure accepts five parameters. If a table, owner, or database name is given, the list of tables will be limited to only those that match the arguments.

Standards and compatibility

- **Sybase** – Supported by Open Client/Open Server.

Permissions

None

Side effects

None

Examples

To get a list of all the Microsoft Excel worksheets available from an ODBC data source referenced by a server named excel:

```
CALL sp_remote_tables( 'excel' );
```

To get a list of all the tables owned by fred in the production database in an Adaptive Server Enterprise server named asetest:

```
CALL sp_remote_tables( 'asetest', null, 'fred', 'production' );
```

sp_servercaps system procedure

Displays information about a remote server's capabilities.

The server must be defined with the CREATE SERVER statement to use this system procedure.

Syntax

```
sp_servercaps ( sname )
```

Arguments

- **sname** – Use this CHAR(64) parameter to specify a server defined with the CREATE SERVER statement. *sname* must be the same server name used in the CREATE SERVER statement.

Remarks

This procedure displays information about a remote server's capabilities. The capability information is used to determine how much of a SQL statement can be forwarded to a remote server. The ISYSCAPABILITY system table, which lists the server capabilities, is not populated until a connection is made to the first remote server.

Standards and compatibility

- **Sybase** – Supported by Open Client/Open Server.

Permissions

None

Side effects

None

Example

To display information about the remote server testasa:

```
CALL sp_servercaps( 'testasa' );
```

sp_tsql_environment system procedure

Sets connection options when users connect from jConnect or Open Client applications.

Syntax

```
sp_tsql_environment( )
```

Remarks

The sp_login_environment procedure is the default procedure specified by the login_procedure database option. For each new connection, the procedure specified by

login_procedure is called. If the connection uses the TDS communications protocol (that is, if it is an Open Client or jConnect connection), then sp_login_environment in turn calls sp_tsql_environment.

This procedure sets database options so that they are compatible with default Adaptive Server Enterprise behavior.

If you want to change the default behavior, it is recommended that you create new procedures and alter your login_procedure option to point to these new procedures.

Permissions

None

Side effects

None

Example

Here is the text of the sp_tsql_environment procedure:

```
CREATE PROCEDURE dbo.sp_tsql_environment()
BEGIN
    IF db_property( 'IQStore' ) = 'Off' THEN
        -- SQL Anywhere datastore
        SET TEMPORARY OPTION close_on_endtrans='OFF';
    END IF;
    SET TEMPORARY OPTION ansinull='OFF';
    SET TEMPORARY OPTION tsql_variables='ON';
    SET TEMPORARY OPTION ansi_blanks='ON';
    SET TEMPORARY OPTION chained='OFF';
    SET TEMPORARY OPTION quoted_identifier='OFF';
    SET TEMPORARY OPTION allow_nulls_by_default='OFF';
    SET TEMPORARY OPTION on_tsql_error='CONTINUE';
    SET TEMPORARY OPTION isolation_level='1';
    SET TEMPORARY OPTION date_format='YYYY-MM-DD';
    SET TEMPORARY OPTION timestamp_format='YYYY-MM-DD HH:NN:SS.SSS';
    SET TEMPORARY OPTION time_format='HH:NN:SS.SSS';
    SET TEMPORARY OPTION date_order='MDY';
    SET TEMPORARY OPTION escape_character='OFF';
END
```

Adaptive Server Enterprise System and Catalog Procedures

Adaptive Server Enterprise provides system and catalog procedures to carry out many administrative functions and to obtain system information. Sybase IQ supports some of these procedures.

System procedures are built-in stored procedures used for getting reports from and updating system tables. Catalog stored procedures retrieve information from the system tables in tabular form.

Note: While these procedures perform the same functions as they do in Adaptive Server Enterprise and pre-version 12 Sybase IQ, they are not identical. If you have preexisting scripts that use these procedures, you might want to examine the procedures. To see the text of a stored procedure, run:

```
sp_helptext 'owner.procedure_name'
```

For all system stored procedures delivered by Sybase, the owner is dbo. To see the text of a stored procedure of the same name owned by a different user, you must specify that user, for example:

```
sp_helptext 'myname.myprocedure'
```

Adaptive Server Enterprise System Procedures

Adaptive Server Enterprise system procedures provided in Sybase IQ:

Table 209. ASE system procedures provided in Sybase IQ

System procedure	Description	Permissions
sp_addgroup <i>group-name</i>	Adds a group to a database	Requires DBA or PERMS ADMIN to change an existing user to a group. Requires DBA or USER ADMIN and PERMS ADMIN to create a new user and change to a group.
sp_addlogin <i>userid, password[, defdb[, deflanguage[, fullname]]]</i>	Adds a new user account to a database	Requires DBA or USER ADMIN authority.
sp_addmessage <i>message-num, message_text [, language]</i>	Adds user-defined messages to SYSUSERMESSAGES for use by stored procedure PRINT and RAISERROR calls	Requires DBA or RESOURCE authority.
sp_addtype <i>typename, data-type, [, "identity" null-type]</i>	Creates a user-defined data type. Sybase IQ does not support IDENTITY columns.	Requires DBA or RESOURCE authority.
sp_adduser <i>userid [, name_in_db [, grpname]]</i>	Adds a new user to a database	Requires DBA or USER ADMIN authority to create a new user. Requires DBA or USER ADMIN and PERMS ADMIN authority to create a new user and add the user to the group specified.
sp_changegroup <i>new-group-name, userid</i>	Changes a user's group or adds a user to a group	Requires DBA or PERMS ADMIN authority.
sp_dboption <i>[dbname, optname, {true false}]</i>	Displays or changes database options	None required.

System procedure	Description	Permissions
sp_dropgroup <i>group-name</i>	Drops a group from a database	Requires DBA or PERMS ADMIN authority.
sp_droplogin <i>userid</i>	Drops a user from a database	Requires DBA or USER ADMIN authority.
sp_dropmessage <i>message-number</i> [, <i>language</i>]	Drops user-defined messages	Requires DBA or RESOURCE authority.
sp_droptype <i>typename</i>	Drops a user-defined data type	Requires DBA or RESOURCE authority.
sp_dropuser <i>userid</i>	Drops a user from a database	Requires DBA or USER ADMIN authority.
sp_getmessage <i>message-num</i> , <i>@msg-var output</i> [, <i>language</i>]	Retrieves stored message strings from SYSUSERMESSAGES for PRINT and RAISERROR statements.	None required.
sp_helptext ' <i>owner.object-name</i> '	Displays the text of a system procedure or view	None required.
sp_password <i>caller_passwd</i> , <i>new_passwd</i> [, <i>userid</i>]	Adds or changes a password for a user ID	Requires DBA authority or PERMS ADMIN authority for users without DBA authority to change the passwords of other users. You may change your own password without these authorities.

Note: Procedures like **sp_dropuser** provide minimal compatibility with Adaptive Server Enterprise stored procedures. If you are accustomed to Adaptive Server Enterprise (or Sybase IQ 11.x) stored procedures, compare their text with Sybase IQ 12 procedures before using the procedure in Interactive SQL. To compare, use the command:

```
sp_helptext 'owner.procedure_name'
```

For system stored procedures delivered by Sybase, the owner is always dbo. To see the text of a stored procedure of the same name owned by a different user, you must specify that user, for example:

```
sp_helptext 'myname.myprocedure'
```

See also

- *Users, Groups, and Permissions* on page 698

Adaptive Server Enterprise Catalog Procedures

Sybase IQ implements most of the Adaptive Server Enterprise catalog procedures with the exception of the `sp_column_privileges` procedure.

Sybase IQ also has similar customized stored procedures for some of these Adaptive Server Enterprise catalog procedures.

Table 210. ASE catalog procedures implemented in Sybase IQ

ASE catalog procedure	Description	IQ procedure
<code>sp_columns</code> <i>table-name</i> [, <i>table-owner</i>] [, <i>table-qualifier</i>] [, <i>column-name</i>]	Returns the data types of the specified column	
<code>sp_fkeys</code> <i>pktable_name</i> [, <i>pktable-owner</i>] [, <i>pktable-qualifier</i>] [, <i>fktable-name</i>] [, <i>fktable_owner</i>] [, <i>fktable-qualifier</i>]	Returns foreign-key information about the specified table	
<code>sp_pkeys</code> <i>table-name</i> [, <i>table_owner</i>] [, <i>table-qualifier</i>]	Returns primary-key information for a single table	<code>sp_iqpkeys</code>
<code>sp_special_columns</code> <i>table_name</i> [, <i>table-owner</i>] [, <i>table-qualifier</i>] [, <i>col-type</i>]	Returns the optimal set of columns that uniquely identify a row in a table	
<code>sp_sproc_columns</code> <i>proc-name</i> [, <i>proc_owner</i>] [, <i>proc-qualifier</i>] [, <i>column-name</i>]	Returns information about the input and return parameters of a stored procedure	<code>sp_iqprocparm</code>
<code>sp_stored_procedures</code> [<i>sp-name</i>] [, <i>sp-owner</i>] [, <i>sp-qualifier</i>]	Returns information about one or more stored procedures	<code>sp_iqprocedure</code>
<code>sp_tables</code> <i>table-name</i> [, <i>table-owner</i>] [, <i>table-qualifier</i>] [, <i>table-type</i>]	Returns a list of objects that can appear in a FROM clause	

The following Adaptive Server Enterprise catalog procedures are not supported:

- `sp_column_privileges`
- `sp_databases`
- `sp_datatype_info`
- `sp_server_info`

SQL Anywhere Supported Procedures

Sybase IQ supports the SQL Anywhere system procedures.

See *SQL Anywhere Server - SQL Reference > System Objects > System procedures > Alphabetical list of system procedures*.

The **sa_get_table_definition** procedure is only supported for SQL Anywhere tables. If run against an IQ table, the procedure returns the error `not implemented for IQ tables`.

System Tables and Views

These topics describes the system tables, system views, consolidated views, compatibility views, and ASE T-SQL compatibility views supported by Sybase IQ.

System Tables

The structure of every Sybase IQ database is described in a number of system tables. The system tables are designed for internal use.

The DUMMY system table is the only system table you are permitted to access directly. For all other system tables, you access their underlying data through their corresponding views.

Table 211. List of system tables

System table	Internal use only?
DUMMY	No
ISYSARTICLE	Yes
ISYSARTICLECOL	Yes
ISYSATTRIBUTE	Yes
ISYSATTRIBUTENAME	Yes
ISYSCAPABILITY	Yes
ISYSCHECK	Yes
ISYSCOLPERM	Yes
ISYSCOLSTAT	Yes
ISYSCONSTRAINT	Yes
ISYSDBFILE	Yes
ISYSDBSPACE	Yes
ISYSDBSPACEPERM	Yes
ISYSDEPENDENCY	Yes
ISYSDOMAIN	Yes
ISYSEVENT	Yes
ISYSEXTERNENV	Yes

System Tables and Views

System table	Internal use only?
ISYSEXTERNENVOBJECT	Yes
ISYSEXTERNLOGIN	Yes
ISYSFKEY	Yes
ISYSGROUP	Yes
ISYSHISTORY	Yes
ISYSIDX	Yes
ISYSIDXCOL	Yes
ISYSIQBACKUPHISTORY	Yes
ISYSIQBACKUPHISTORYDETAIL	Yes
ISYSIQDBFILE	Yes
ISYSIQDBSPACE	Yes
ISYSIQIDX	Yes
ISYSIQINFO	Yes
ISYSIQJOINIDX	Yes
ISYSIQJOINIXCOLUMN	Yes
ISYSIQLOGICALSERVER	Yes
ISYSIQLOGINPOLICYLSINFO	Yes
ISYSIQLSLOGINPOLICYOPTION	Yes
ISYSIQLSMEMBER	Yes
ISYSIQLSPOLICY	Yes
ISYSIQLSPOLICYOPTION	Yes
ISYSIQMPXSERVER	Yes
ISYSIQPARTITIONCOLUMN	Yes
ISYSIQTAB	Yes
ISYSIQTABCOL	Yes
ISYSJAR	Yes
ISYSJARCOMPONENT	Yes
ISYSJAVACLASS	Yes

System table	Internal use only?
ISYSLOGINMAP	Yes
ISYSLOGINPOLICY	Yes
ISYSLOGINPOLICYOPTION	Yes
ISYSMVOPTION	Yes
ISYSMVOPTIONNAME	Yes
ISYSOBJECT	Yes
ISYSOPTION	Yes
ISYSOPTSTAT	Yes
ISYSPARTITION	Yes
ISYSPARTITIONKEY	Yes
ISYSPARTITIONSCHEME	Yes
ISYSPHYSIDX	Yes
ISYSPROCEDURE	Yes
ISYSPROCPARM	Yes
ISYSPROCPERM	Yes
ISYSPROXYTAB	Yes
ISYSPUBLICATION	Yes
ISYSREMARK	Yes
ISYSREMOTEOPTION	Yes
ISYSREMOTEOPTIONTYPE	Yes
ISYSREMOTETYPE	Yes
ISYSREMOTEUSER	Yes
ISYSSCHEDULE	Yes
ISYSSERVER	Yes
ISYSSOURCE	Yes
ISYSSQLSERVERTYPE	Yes
ISYSSUBPARTITIONKEY	Yes
ISYSSUBSCRIPTION	Yes

System table	Internal use only?
ISYSSYNC	Yes
ISYSSYNCPROFILE	Yes
ISYSSYNCSRIPT	Yes
ISYSTAB	Yes
ISYSTABCOL	Yes
ISYSTABLEPERM	Yes
ISYSTECONFIG	Yes
ISYSTEINDEX	Yes
ISYSTEINDEXTAB	Yes
ISYSTRIGGER	Yes
ISYSTYPEMAP	Yes
ISYSUSER	Yes
ISYSUSERAUTHORITY	Yes
ISYSUSERMESSAGE	Yes
ISYSUSERTYPE	Yes
ISYSVIEW	Yes
ISYSWEBSERVICE	Yes

SYS.DUMMY Table Versus IQ_DUMMY Table

The DUMMY system table is provided as a table that always has exactly one row.

This can be useful for extracting information from the database, as in the following example that gets the current user ID and the current date from the database

```
SELECT USER, today(*) FROM SYS.DUMMY
```

Queries using the DUMMY table are run by SQL Anywhere (the catalog store), rather than by Sybase IQ. You can create a dummy table in the Sybase IQ database. For example:

```
CREATE TABLE iq_dummy (dummy_col INT NOT NULL);
```

and use this table explicitly:

```
SELECT NOW() FROM iq_dummy;
```

DUMMY system table

Column name	Column type	Column constraint	Table constraints
dummy_col	INTEGER	NOT NULL	

The DUMMY table is provided as a read-only table that always has exactly one row. This can be useful for extracting information from the database, as in the following example that gets the current user ID and the current date from the database.

```
SELECT USER, today(*) FROM IQ.DUMMY;
```

```
SELECT USER, today(*);
```

dummy_col – This column is not used. It is present because a table cannot be created with no columns.

The cost of reading from the IQ.DUMMY table is less than the cost of reading from a similar user-created table because there is no latch placed on the table page of IQ.DUMMY.

Access plans are not constructed with scans of the IQ.DUMMY table. Instead, references to IQ.DUMMY are replaced with a Row Constructor algorithm, which virtualizes the table reference. This eliminates contention associated with the use of IQ.DUMMY. Note that DUMMY still appears as the table and/or correlation name in short, long, and graphical plans.

System Views

Use the system views to view the contents of the system tables.

A number of predefined system views are provided that present the information in the system tables in a readable format.

The definitions for the system views are included with their descriptions. Some of these definitions are complicated, but you do not need to understand them to use the views.

Consolidated Views

Consolidated views provide data in a form more frequently required by users.

For example, consolidated views often provide commonly needed joins. Consolidated views differ from system views in that they are not just a straightforward view of raw data in an underlying system table. For example, many of the columns in the system views are unintelligible ID values, whereas in the consolidated views, they are readable names.

Consolidated views such as SYSCATALOG and SYSINDEXES are common to both Sybase IQ and SQL Anywhere. For definitions of these and other consolidated views, see *SQL Anywhere Server - SQL Reference > Views > Consolidated views*.

Compatibility Views

Compatibility views are deprecated views provided for compatibility with earlier versions of SQL Anywhere and Sybase IQ.

Where possible, use system views and consolidated views instead of compatibility views, as support for compatibility views may be eliminated in future versions of Sybase IQ.

For detailed information on compatibility views, see *SQL Anywhere Server - SQL Reference > Views > Compatibility views*.

ASE T-SQL Compatibility Views

Sybase IQ provides a set of views owned by the special user DBO, which correspond to the Adaptive Server Enterprise system tables and views.

See also

- *Transact-SQL Compatibility Views* on page 678

SYSARTICLE system view

Each row of the SYSARTICLE system view describes an article in a publication. The underlying system table for this view is ISYSARTICLE.

Column name	Data type	Description
publication_id	UNSIGNED INT	The publication of which the article is a part.
table_id	UNSIGNED INT	Each article consists of columns and rows from a single table. This column contains the table ID for this table.
where_expr	LONG VARCHAR	For articles that contain a subset of rows defined by a WHERE clause, this column contains the search condition.
subscribe_by_expr	LONG VARCHAR	For articles that contain a subset of rows defined by a SUBSCRIBE BY expression, this column contains the expression.
query	CHAR(1)	Indicates information about the article type to the database server.
alias	VARCHAR(256)	The alias for the article.

Column name	Data type	Description
schema_change_active	BIT	1 if the table and publication are part of a synchronization schema change.

Constraints on underlying system table

```
PRIMARY KEY (publication_id, table_id)
```

```
FOREIGN KEY (publication_id) references SYS.ISYSPUBLICATION  
(publication_id)
```

```
FOREIGN KEY (table_id) references SYS.ISYSTAB (table_id)
```

SYSARTICLECOL system view

Each row of the SYSARTICLECOL system view identifies a column in an article. The underlying system table for this view is ISYSARTICLECOL.

Column name	Data type	Description
publication_id	UNSIGNED INT	A unique identifier for the publication of which the column is a part.
table_id	UNSIGNED INT	The table to which the column belongs.
column_id	UNSIGNED INT	The column identifier, from the SYSTABCOL system view.

Constraints on underlying system table

```
PRIMARY KEY (publication_id, table_id, column_id)
```

```
FOREIGN KEY (publication_id, table_id) references SYS.ISYSARTICLE  
(publication_id, table_id)
```

```
FOREIGN KEY (table_id, column_id) references SYS.ISYSTABCOL  
(table_id, column_id)
```

SYSARTICLECOLS consolidated view

Each row in the SYSARTICLECOLS view identifies a column in an article.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSARTICLECOLS"  
as select p.publication_name,t.table_name,c.column_name  
from SYS.ISYSARTICLECOL as ac  
join SYS.ISYSPUBLICATION as p on p.publication_id =  
ac.publication_id
```

```

join SYS.ISYSTAB as t on t.table_id = ac.table_id
join SYS.ISYSTABCOL as c on c.table_id = ac.table_id
and c.column_id = ac.column_id
    
```

SYSARTICLES consolidated view

Each row in the SYSARTICLES view describes an article in a publication.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```

ALTER VIEW "SYS"."SYSARTICLES"
  as select u1.user_name as publication_owner,p.publication_name,
    u2.user_name as table_owner,t.table_name,
    a.where_expr,a.subscribe_by_expr,a.alias
  from SYS.ISYSARTICLE as a
    join SYS.ISYSPUBLICATION as p on(a.publication_id =
p.publication_id)
    join SYS.ISYSTAB as t on(a.table_id = t.table_id)
    join SYS.ISYSUSER as u1 on(p.creator = u1.user_id)
    join SYS.ISYSUSER as u2 on(t.creator = u2.user_id)
    
```

SYSCAPABILITIES consolidated view

Each row in the SYSCAPABILITIES view specifies the status of a capability for a remote database server. This view gets its data from the ISYSCAPABILITY and ISYSCAPABILITYNAME system tables.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```

ALTER VIEW "SYS"."SYSCAPABILITIES"
  as select
  ISYSCAPABILITY.capid,ISYSCAPABILITY.srvid,property('RemoteCapabilit
y',ISYSCAPABILITY.capid) as capname,ISYSCAPABILITY.capvalue
  from SYS.ISYSCAPABILITY
    
```

SYSCAPABILITY system view

Each row of the SYSCAPABILITY system view specifies the status of a capability on a remote database server. The underlying system table for this view is ISYSCAPABILITY.

Column name	Data type	Description
capid	INTEGER	The ID of the capability, as listed in the SYSCAPABILITY-NAME system view.
srvid	UNSIGNED INT	The server to which the capability applies, as listed in the SYS-SERVER system view.

Column name	Data type	Description
capvalue	CHAR(128)	The value of the capability.

Constraints on underlying system table

PRIMARY KEY (capid, srvid)

FOREIGN KEY (srvid) references SYS.ISYSSERVER (srvid)

SYSCAPABILITYNAME system view

Each row in the SYSCAPABILITYNAME system view provides a name for each capability ID in the SYSCAPABILITY system view.

Column name	Data type	Description
capid	INTEGER	A number uniquely identifying the capability.
capname	VARCHAR(32000)	The name of the capability.

Remarks

The SYSCAPABILITYNAME system view is defined using a combination of sa_rowgenerator and the following server properties:

RemoteCapability

MaxRemoteCapability

SYSCATALOG consolidated view

Each row in the SYSCATALOG view describes a system table.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSCATALOG"( creator,
    tname,dbspacename,tabletype,ncols,primary_key,"check",
    remarks )
as select u.user_name,tab.table_name,dbs.dbSPACE_name,
    if tab.table_type_str = 'BASE' then 'TABLE' else
tab.table_type_str endif,
    (select count() from SYS.ISYSTABCOL
    where ISYSTABCOL.table_id = tab.table_id),
    if ix.index_id is null then 'N' else 'Y' endif,
    null,
    rmk.remarks
from SYS.SYSTAB as tab
    join SYS.ISYSDBSPACE as dbs on(tab.dbSPACE_id = dbs.dbSPACE_id)
    join SYS.ISYSUSER as u on u.user_id = tab.creator
    left outer join SYS.ISYSIDX as ix on(tab.table_id = ix.table_id
and ix.index_id = 0)
```

```
left outer join SYS.ISYSREMARK as rmk on(tab.object_id =
rmk.object_id)
```

SYSCHECK system view

Each row in the SYSCHECK system view provides the definition for a named check constraint in a table. The underlying system table for this view is ISYSCHECK.

Column name	Data type	Description
check_id	UNSIGNED INT	A number that uniquely identifies the constraint in the database.
check_defn	LONG VARCHAR	The CHECK expression.

Constraints on underlying system table

```
PRIMARY KEY (check_id)
```

```
FOREIGN KEY (check_id) references SYS.ISYSCONSTRAINT (constraint_id)
```

SYSCOLAUTH consolidated view

Each row in the SYSCOLAUTH view describes the set of privileges (UPDATE, SELECT, or REFERENCES) granted on a column.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW
"SYS"."SYSCOLAUTH"( grantor,grantee,creator,tname,colname,
privilege_type,is_grantable )
as select u1.user_name,u2.user_name,u3.user_name,tab.table_name,
col.column_name,cp.privilege_type,cp.is_grantable
from SYS.ISYSCOLPERM as cp
join SYS.ISYSUSER as u1 on u1.user_id = cp.grantor
join SYS.ISYSUSER as u2 on u2.user_id = cp.grantee
join SYS.ISYSTAB as tab on tab.table_id = cp.table_id
join SYS.ISYSUSER as u3 on u3.user_id = tab.creator
join SYS.ISYSTABCOL as col on col.table_id = cp.table_id
and col.column_id = cp.column_id
```

SYSCOLLATION compatibility view (deprecated)

The SYSCOLLATION compatibility view contains the collation sequence information for the database. It is obtainable via built-in functions and is not kept in the catalog. Following is definition for this view:

```
ALTER VIEW "SYS"."SYSCOLLATION"
as select 1 as collation_id,
DB_PROPERTY('Collation') as collation_label,
DB_EXTENDED_PROPERTY('Collation','Description') as
collation_name,
```



```
cast(DB_EXTENDED_PROPERTY('Collation','LegacyData') as
binary(1280)) as collation_order
```

SYSCOLLATIONMAPPINGS compatibility view (deprecated)

The SYSCOLLATIONMAPPINGS compatibility view contains only one row with the database collation mapping. It is obtainable via built-in functions and is not kept in the catalog. Following is definition for this view:

```
ALTER VIEW "SYS"."SYSCOLLATIONMAPPINGS"
as select DB_PROPERTY('Collation') as collation_label,
DB_EXTENDED_PROPERTY('Collation','Description') as
collation_name,
DB_PROPERTY('Charset') as cs_label,
DB_EXTENDED_PROPERTY('Collation','ASEnsensitiveSortOrder') as
so_case_label,
DB_EXTENDED_PROPERTY('Collation','ASEnsensitiveSortOrder') as
so_caseless_label,
DB_EXTENDED_PROPERTY('Charset','java') as jdk_label
```

SYSCOLPERM system view

The GRANT statement can give UPDATE, SELECT, or REFERENCES permission to individual columns in a table. Each column with UPDATE, SELECT, or REFERENCES permission is recorded in one row of the SYSCOLPERM system view. The underlying system table for this view is ISYSCOLPERM.

Column name	Data type	Description
table_id	UNSIGNED INT	The table number for the table containing the column.
grantee	UNSIGNED INT	The user number of the user ID that is given permission on the column. If the grantee is the user number for the special PUBLIC user ID, the permission is given to all user IDs.
grantor	UNSIGNED INT	The user number of the user ID that grants the permission.
column_id	UNSIGNED INT	This column number, together with the table_id, identifies the column for which permission has been granted.
privilege_type	SMALLINT	The number in this column indicates the kind of column permission (16=REFERENCES, 1=SELECT, or 8=UPDATE).

Column name	Data type	Description
is_grantable	CHAR(1)	Indicates if the permission on the column was granted WITH GRANT OPTION.

Constraints on underlying system table

PRIMARY KEY (table_id, grantee, grantor, column_id, privilege_type)

FOREIGN KEY (table_id, column_id) references SYS.ISYSTABCOL (table_id, column_id)

FOREIGN KEY (grantor) references SYS.ISYSUSER (user_id)

FOREIGN KEY (grantee) references SYS.ISYSUSER (user_id)

SYSCOLSTAT system view

The SYSCOLSTAT system view contains the column statistics, including histograms, that are used by the optimizer. The contents of this view are best retrieved using the sa_get_histogram stored procedure or the Histogram utility. The underlying system table for this view is ISYSCOLSTAT.

Column name	Data type	Description
table_id	UNSIGNED INT	A number that uniquely identifies the table or materialized view to which the column belongs.
column_id	UNSIGNED INT	A number that, together with table_id, uniquely identifies the column.
format_id	SMALLINT	For system use only.
update_time	TIMESTAMP	The time of the last update of the column statistics.
density	FLOAT	An estimate of the average selectivity of a single value for the column, not counting the large single value selectivities stored in the row.
max_steps	SMALLINT	For system use only.
actual_steps	SMALLINT	For system use only.
step_values	LONG BINARY	For system use only.
frequencies	LONG BINARY	For system use only.

Constraints on underlying system table

```
PRIMARY KEY (table_id, column_id)
```

```
FOREIGN KEY (table_id, column_id) references SYS.ISYSTABCOL  
(table_id, column_id)
```

SYSCOLSTATS consolidated view

The SYSCOLSTATS view contains the column statistics that are stored as histograms and used by the optimizer.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSCOLSTATS"  
  as select u.user_name,t.table_name,c.column_name,  
           s.format_id,s.update_time,s.density,s.max_steps,  
           s.actual_steps,s.step_values,s.frequencies  
  from SYS.ISYSCOLSTAT as s  
       join SYS.ISYSTABCOL as c on(s.table_id = c.table_id  
       and s.column_id = c.column_id)  
       join SYS.ISYSTAB as t on(t.table_id = c.table_id)  
       join SYS.ISYSUSER as u on(u.user_id = t.creator)
```

SYSCOLUMN compatibility view (deprecated)

The SYSCOLUMN view is provided for compatibility with older versions of the software that offered a SYSCOLUMN system table. However, the previous SYSCOLUMN table has been replaced by the ISYSTABCOL system table, and its corresponding SYSTABCOL system view, which you should use instead.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSCOLUMN"  
  as select b.table_id,  
           b.column_id,  
           if c.sequence is null then 'N' else 'Y' endif as pkey,  
           b.domain_id,  
           b.nulls,  
           b.width,  
           b.scale,  
           b.object_id,  
           b.max_identity,  
           b.column_name,  
           r.remarks,  
           b."default",  
           b.user_type,  
           b.column_type  
  from SYS.SYSTABCOL as b  
       left outer join SYS.ISYSREMARK as r on(b.object_id =
```

```
r.object_id)
    left outer join SYS.ISYSIDXCOL as c on(b.table_id = c.table_id
and b.column_id = c.column_id and c.index_id = 0)
```

SYSCOLUMNS consolidated view

Each row in the SYSCOLUMNS view describes one column of each table and view in the catalog.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW
"SYS"."SYSCOLUMNS"( creator,cname,tname,coltype,nulls,length,
    syslength,in_primary_key,colno,default_value,
    column_kind,remarks )
as select
u.user_name,col.column_name,tab.table_name,dom.domain_name,
    col.nulls,col.width,col.scale,if ixcol.sequence is null then 'N'
else 'Y' endif,col.column_id,
    col."default",col.column_type,rmk.remarks
from SYS.SYSTABCOL as col
    left outer join SYS.ISYSIDXCOL as ixcol on(col.table_id =
ixcol.table_id and col.column_id = ixcol.column_id and
ixcol.index_id = 0)
    join SYS.ISYSTAB as tab on(tab.table_id = col.table_id)
    join SYS.ISYSDOMAIN as dom on(dom.domain_id = col.domain_id)
    join SYS.ISYSUSER as u on u.user_id = tab.creator
    left outer join SYS.ISYSREMARK as rmk on(col.object_id =
rmk.object_id)
```

SYSCOLUMNS ASE Compatibility View

This view is owned by user DBO. syscolumns contains one row for every column in every table and view, and a row for each parameter in a procedure.

See also

- *Tables in Each Adaptive Server Enterprise Database* on page 679
- *SYSCOMMENTS ASE Compatibility View* on page 587
- *SYSINDEXES ASE Compatibility View* on page 607
- *SYSIQOBJECTS ASE Compatibility View* on page 619
- *SYSIQVINDEXT ASE Compatibility View* on page 623
- *SYSOBJECTS ASE Compatibility View* on page 629
- *SYSTYPES ASE Compatibility View* on page 670
- *SYSUSERS ASE Compatibility View* on page 676

SYSCOMMENTS ASE Compatibility View

syscomments contains entries for each view, rule, default, trigger, table constraint, partition, procedure, computed column, function-based index key, and other forms of compiled objects.

This view is owned by user DBO.

The text column contains the original definition statements. If the text column is longer than 255 bytes, the entries span rows. Each object can occupy as many as 65,025 rows.

See also

- *Tables in Each Adaptive Server Enterprise Database* on page 679
- *SYSCOLUMNS ASE Compatibility View* on page 586
- *SYSINDEXES ASE Compatibility View* on page 607
- *SYSIQOBJECTS ASE Compatibility View* on page 619
- *SYSIQVINDEXT ASE Compatibility View* on page 623
- *SYSOBJECTS ASE Compatibility View* on page 629
- *SYSTYPES ASE Compatibility View* on page 670
- *SYSUSERS ASE Compatibility View* on page 676

SYSCONSTRAINT system view

Each row in the SYSCONSTRAINT system view describes a named constraint in the database. The underlying system table for this view is ISYSCONSTRAINT.

Column name	Data type	Description
constraint_id	UNSIGNED INT	The unique ID for the constraint.
constraint_type	CHAR(1)	The type of constraint: <ul style="list-style-type: none"> • C – column check constraint • T – table constraint • P – primary key • F – foreign key • U – unique constraint
ref_object_id	UNSIGNED BIGINT	The object ID of the column, table, or index to which the constraint applies.
table_object_id	UNSIGNED BIGINT	The object ID of the table to which the constraint applies.
constraint_name	CHAR(128)	The name of the constraint.

Constraints on underlying system table

```
PRIMARY KEY (constraint_id)
FOREIGN KEY (ref_object_id) references SYS.ISYSOBJECT (object_id)
FOREIGN KEY (table_object_id) references SYS.ISYSOBJECT (object_id)
UNIQUE Constraint (table_object_id, constraint_name)
```

SYSDBFIL system view

Each row in the SYSDBFIL system view describes a dbspace file. The underlying system table for this view is ISYSDBFIL.

Column name	Data type	Description
dbfile_id	SMALLINT	For internal use only.
dbspace_id	SMALLINT	Each dbspace file in a database is assigned a unique number. The system dbspace contains all system objects and has a dbspace_id of 0.
dbfile_name	CHAR(128)	The file name for the dbspace. For dbspaces other than system and TEMPORARY, the file name can be changed using the following statement: <code>ALTER DBSPACE <i>dbspace</i> RENAME 'new-file-name' ;</code>
file_name	LONG VARCHAR	A unique name for the dbspace. It is used in the CREATE TABLE command.
lob_map	LONG VARBIT	For internal use only.

Constraints on underlying system table

```
PRIMARY KEY (dbfile_id)
FOREIGN KEY (dbspace_id) references SYS.ISYSDBSpace (dbspace_id)
UNIQUE Index (file_name)
```

SYSDBSPACE system view

Each row in the SYSDBSPACE system view describes a dbspace file. The underlying system table for this view is ISYSDBSPACE.

Column name	Data type	Description
dbspace_id	SMALLINT	Unique number identifying the dbspace. The system dbspace contains all system objects and has a dbspace_id of 0.
object_id	UNSIGNED BIGINT	The file name for the dbspace. For the system dbspace, the value is the name of the database file when the database was created and is for informational purposes only; it cannot be changed. For other dbspaces, the file name can be changed using the following statement: <pre>ALTER DBSPACE <i>dbspace</i> RENAME '<i>new-file-name</i>' ;</pre>
dbspace_name	CHAR(128)	A unique name for the dbspace. It is used in the CREATE TABLE command.
store_type	TINYINT	For internal use only.

Constraints on underlying system table

PRIMARY KEY (dbspace_id)

FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH UNIQUE FULL

SYSDBSPACEPERM System View

Each row in the SYSDBSPACEPERM system view describes a permission on a dbspace file. The underlying system table for this view is ISYSDBSPACEPERM.

The SYSDBSPACEPERM view is a SQL Anywhere system view. See *SQL Anywhere Server - SQL Reference > Views > System views > SYSDBFILE system view*.

Note: This reference points to SQL Anywhere documentation.

SYSDEPENDENCY system view

Each row in the SYSDEPENDENCY system view describes a dependency between two database objects. The underlying system table for this view is ISYSDEPENDENCY.

A dependency exists between two database objects when one object references another object in its definition. For example, if the query specification for a view references a table, the view is said to be dependent on the table. The database server tracks dependencies of views on tables, views, materialized views, and columns.

Column name	Data type	Description
ref_object_id	UNSIGNED BIGINT	The object ID of the referenced object.
dep_object_id	UNSIGNED BIGINT	The ID of the referencing object.

Constraints on underlying system table

PRIMARY KEY (ref_object_id, dep_object_id)

FOREIGN KEY (ref_object_id) references SYS.ISYSOBJECT (object_id)

FOREIGN KEY (dep_object_id) references SYS.ISYSOBJECT (object_id)

SYSDOMAIN system view

The SYSDOMAIN system view records information about built-in data types (also called domains). The contents of this view does not change during normal operation. The underlying system table for this view is ISYSDOMAIN.

Column name	Data type	Description
domain_id	SMALLINT	The unique number assigned to each data type. These numbers cannot be changed.
domain_name	CHAR(128)	The name of the data type normally found in the CREATE TABLE command, such as CHAR or INTEGER.
type_id	SMALLINT	The ODBC data type. This value corresponds to the value for data_type in the Transact-SQL-compatibility dbo.SYSTYPES table.

Column name	Data type	Description
"precision"	SMALLINT	The number of significant digits that can be stored using this data type. The column value is NULL for non-numeric data types.

Constraints on underlying system table

PRIMARY KEY (domain_id)

SYSEVENT system view

Each row in the SYSEVENT system view describes an event created with CREATE EVENT. The underlying system table for this view is ISYSEVENT.

Column name	Data type	Description
event_id	UNSIGNED INT	The unique number assigned to each event.
object_id	UNSIGNED BIGINT	The internal ID for the event, uniquely identifying it in the database.
creator	UNSIGNED INT	The user number of the owner of the event. The name of the user can be found by looking in the SYSUSER system view.
event_name	VARCHAR(128)	The name of the event.
enabled	CHAR(1)	Indicates whether the event is allowed to fire.
location	CHAR(1)	The location where the event is to fire: C = consolidated R = remote A = all
event_type_id	UNSIGNED INT	For system events, the event type as listed in the SYSEVENTTYPE system view.
action	LONG VARCHAR	The event handler definition. An obfuscated value indicates a hidden event.
external_action	LONG VARCHAR	For system use only.

Column name	Data type	Description
condition	LONG VARCHAR	The condition used to control firing of the event handler.
remarks	LONG VARCHAR	Remarks for the event; this column comes from ISYSREMARK.
source	LONG VARCHAR	The original source for the event; this column comes from ISYSSOURCE.

Constraints on underlying system table

PRIMARY KEY (event_id)

FOREIGN KEY (creator) references SYS.ISYSUSER (user_id)

FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH UNIQUE FULL

UNIQUE Index (event_name)

SYSEVENTTYPE System View

The SYSEVENTTYPE system view defines the system event types that can be referenced by CREATE EVENT.

The SYSEVENTTYPE view is a SQL Anywhere system view. See *SQL Anywhere Server - SQL Reference > Views > System views > SYSEVENT system view.*

Note: This reference points to SQL Anywhere documentation.

SYSEXTERNENV system view

Many external runtime environments are supported, including embedded SQL and ODBC applications written in C/C++, and applications written in Java, Perl, PHP, or languages such as C# and Visual Basic that are based on the Microsoft .NET Framework Common Language Runtime (CLR).

Each row in the SYSEXTERNENV system view describes the information needed to identify and launch each of the external environments. The underlying system table for this view is ISYSEXTERNENV.

Column name	Data type	Description
object_id	unsigned bigint	A unique identifier for the external environment.
name	char(128)	The name of the external environment or language.

Column name	Data type	Description
scope	char(1)	Identifies if the external environment is launched as one-per-connection (C), or one-per-database (D).
support_result_sets	char(1)	Identifies the external environments that can return result sets to the user.
location	long varchar	Identifies the location where the main executable for the environment can be found.
options	long varchar	Identifies the options required on the command line to launch the external environment.
user_id	unsigned int	For internal use only.

Constraints on underlying system table

```
PRIMARY KEY (object_id)
```

```
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH  
UNIQUE FULL
```

```
FOREIGN KEY (user_id) references SYS.ISYSUSER (user_id)
```

```
UNIQUE Index (name)
```

SYSEXTERNENVOBJECT system view

Many external runtime environments are supported, including embedded SQL and ODBC applications written in C/C++, and applications written in Java, Perl, PHP, or languages such as C# and Visual Basic that are based on the Microsoft .NET Framework Common Language Runtime (CLR).

Each row in the SYSEXTERNENVOBJECT system view describes an installed external object. The underlying system table for this view is ISYSEXTERNENVOBJECT.

Column name	Data type	Description
object_id	unsigned bigint	A unique identifier for the external object.
extenv_id	unsigned bigint	The unique identifier for the external environment (SYSEXTERNENV.object_id).
owner	unsigned int	Identifies the creator/owner of the external object.

Column name	Data type	Description
name	long varchar	Identifies the name of the external object as specified in the INSTALL statement.
contents	long binary	The contents of the external object.
update_time	timestamp	Identifies the last time the object was modified (or installed).

Constraints on underlying system table

PRIMARY KEY (object_id)

FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH UNIQUE FULL

FOREIGN KEY (extenv_id) references SYS.ISYSEXTERNENV (object_id)

FOREIGN KEY (owner) references SYS.ISYSUSER (user_id)

UNIQUE Index (name)

SYSEXTERNLOGIN system view

Each row in the SYSEXTERNLOGIN system view describes an external login for remote data access. The underlying system table for this view is ISYSEXTERNLOGIN.

Note: Previous versions of the catalog contained a SYSEXTERNLOGINS system table. That table has been renamed to be ISYSEXTERNLOGIN (without an 'S'), and is the underlying table for this view.

Column name	Data type	Description
user_id	UNSIGNED INT	The user ID on the local database.
srvid	UNSIGNED INT	The remote server, as listed in the SYSSERVER system view.
remote_login	VARCHAR(128)	The login name for the user, for the remote server.
remote_password	VARBINARY(128)	The password for the user, for the remote server.

Constraints on underlying system table

PRIMARY KEY (user_id, srvid)

FOREIGN KEY (user_id) references SYS.ISYSUSER (user_id)

```
FOREIGN KEY (srvid) references SYS.ISYSSERVER (srvid)
```

SYSFILE compatibility view (deprecated)

Each row in the SYSFILE system view describes a dbspace for a database. Every database consists of one or more dbspaces; each dbspace corresponds to an operating system file.

dbspaces are automatically created for the main database file, temporary file, transaction log file, and transaction log mirror file. Information about the transaction log, and transaction log mirror dbspaces does not appear in the SYSFILE system view.

```
ALTER VIEW "SYS"."SYSFILE"
as select b.dbfile_id as file_id,
  if b.dbspace_id = 0 and b.dbfile_id = 0 then
    db_property('File')
  else
    if b.dbspace_id = 15 and b.dbfile_id = 15 then
      db_property('TempFileName')
    else
      b.file_name
    endif
  endif as file_name,
a.dbspace_name,
a.store_type,
b.lob_map,
b.dbspace_id
from SYS.ISYSDBSPACE as a
  join SYS.ISYSDBFILE as b on(a.dbspace_id = b.dbspace_id)
```

SYSFKCOL compatibility view (deprecated)

Each row of SYSFKCOL describes the association between a foreign column in the foreign table of a relationship and the primary column in the primary table. This view is deprecated; use the SYSIDX and SYSIDXCOL system views instead.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSFKCOL"
as select a.table_id as foreign_table_id,
  a.index_id as foreign_key_id,
  a.column_id as foreign_column_id,
  a.primary_column_id
from SYS.ISYSIDXCOL as a
  ,SYS.ISYSIDX as b
where a.table_id = b.table_id
and a.index_id = b.index_id
and b.index_category = 2
```

SYSFKEY system view

Each row in the SYSFKEY system view describes a foreign key constraint in the system. The underlying system table for this view is ISYSFKEY.

Column name	Data type	Description
foreign_table_id	UNSIGNED INT	The table number of the foreign table.
foreign_index_id	UNSIGNED INT	The index number for the foreign key.
primary_table_id	UNSIGNED INT	The table number of the primary table.
primary_index_id	UNSIGNED INT	The index number of the primary key.
match_type	TINYINT	The matching type for the constraint. Matching types include: <ul style="list-style-type: none"> • 0 – Use the default matching • 1 – SIMPLE • 2 – FULL • 129 – SIMPLE UNIQUE • 130 – FULL UNIQUE
check_on_commit	CHAR(1)	Indicates whether INSERT and UPDATE statements should wait until the COMMIT to check if foreign keys are still valid.
nulls	CHAR(1)	Indicates whether the columns in the foreign key are allowed to contain the NULL value. Note that this setting is independent of the nulls setting in the columns contained in the foreign key.

Constraints on underlying system table

```
PRIMARY KEY (foreign_table_id, foreign_index_id)
```

```
FOREIGN KEY (foreign_table_id, foreign_index_id) references  
SYS.ISYSIDX (table_id, index_id)
```

```
FOREIGN KEY (primary_table_id, primary_index_id) references  
SYS.ISYSIDX (table_id, index_id)
```

SYSFOREIGNKEY compatibility view (deprecated)

The SYSFOREIGNKEY view is provided for compatibility with older versions of the software that offered a SYSFOREIGNKEY system table. However, the previous SYSFOREIGNKEY system table has been replaced by the ISYSFKEY system table, and its corresponding SYSFKEY system view, which you should use instead.

A foreign key is a relationship between two tables—the foreign table and the primary table. Every foreign key is defined by one row in SYSFOREIGNKEY and one or more rows in SYSFKCOL. SYSFOREIGNKEY contains general information about the foreign key while SYSFKCOL identifies the columns in the foreign key and associates each column in the foreign key with a column in the primary key of the primary table.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSFOREIGNKEY"
  as select b.foreign_table_id,
    b.foreign_index_id as foreign_key_id,
    a.object_id,
    b.primary_table_id,
    p.root,
    b.check_on_commit,
    b.nulls,
    a.index_name as role,
    r.remarks,
    b.primary_index_id,
    a.not_enforced as fk_not_enforced,
    10 as hash_limit
  from(SYS.ISYSIDX as a left outer join SYS.ISYSPHYSIDX as p
on(a.table_id = p.table_id and a.phys_index_id = p.phys_index_id))
  left outer join SYS.ISYSREMARK as r on(a.object_id =
r.object_id)
  ,SYS.ISYSFKEY as b
  where a.table_id = b.foreign_table_id
  and a.index_id = b.foreign_index_id
```

SYSFOREIGNKEYS consolidated view

Each row in the SYSFOREIGNKEYS view describes one foreign key for each table in the catalog.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSFOREIGNKEYS"( foreign_creator,
  foreign_tname,
  primary_creator,primary_tname,role,columns )
  as select fk_up.user_name,fk_tab.table_name,pk_up.user_name,
    pk_tab.table_name,ix.index_name,
    (select list(string(fk_col.column_name,' IS '),
```

System Tables and Views

```
pk_col.column_name)
order by fkc.table_id,fkc.index_id,fkc."sequence")
from SYS.ISYSIDXCOL as fkc
  join SYS.ISYSTABCOL as fk_col on(
    fkc.table_id = fk_col.table_id
    and fkc.column_id = fk_col.column_id)
  ,SYS.ISYSTABCOL as pk_col
where fkc.table_id = fk.foreign_table_id
and fkc.index_id = fk.foreign_index_id
and pk_col.table_id = fk.primary_table_id
and pk_col.column_id = fkc.primary_column_id)
from SYS.ISYSFKEY as fk
  join SYS.ISYSTAB as fk_tab on fk_tab.table_id =
fk.foreign_table_id
  join SYS.ISYSUSER as fk_up on fk_up.user_id = fk_tab.creator
  join SYS.ISYSTAB as pk_tab on pk_tab.table_id =
fk.primary_table_id
  join SYS.ISYSUSER as pk_up on pk_up.user_id = pk_tab.creator
  join SYS.ISYSIDX as ix on ix.table_id = fk.foreign_table_id and
ix.index_id = fk.foreign_index_id
```

SYSGROUP system view

There is one row in the SYSGROUP system view for each member of each group. This view describes the many-to-many relationship between groups and members. A group may have many members, and a user may be a member of many groups. The underlying system table for this view is ISYSGROUP.

Column name	Data type	Description
group_id	UNSIGNED INT	The user number of the group.
group_member	UNSIGNED INT	The user number of a member.

Constraints on underlying system table

```
PRIMARY KEY (group_id, group_member)
```

```
FOREIGN KEY group_id (group_id) references SYS.ISYSUSER (user_id)
```

```
FOREIGN KEY group_member (group_member) references SYS.ISYSUSER
(user_id)
```

SYSGROUPS consolidated view

There is one row in the SYSGROUPS view for each member of each group. This view describes the many-to-many relationship between groups and members. A group may have many members, and a user may be a member of many groups.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSGROUPS"( group_name,
  member_name )
```



```
as select g.user_name,u.user_name
      from SYS.ISYSGROUP,SYS.ISYSUSER as g,SYS.ISYSUSER as u
      where ISYSGROUP.group_id = g.user_id and ISYSGROUP.group_member =
u.user_id
```

SYSHISTORY system view

Each row in the SYSHISTORY system view records a system operation on the database, such as a database start, a database calibration, and so on. The underlying system table for this view is ISYSHISTORY.

Column name	Data type	Description
operation	CHAR(128)	<p>The type of operation performed on the database file. The operation must be one of the following values:</p> <ul style="list-style-type: none"> • INIT – Information about when the database was created. • UPGRADE – Information about when the database was upgraded. • START – Information about when the database was started using a specific version of the database server on a particular operating system. • LAST_START – Information about the most recent time the database server was started. A LAST_START operation is converted to a START operation when the database is started with a different version of the database server and/or on a different operating system than those values currently stored in the LAST_START row. • DTT – Information about the <i>second to last</i> Disk Transfer Time (DTT) calibration operation performed on the dbspace. That is, information about the second to last execution of either an ALTER DATABASE CALIBRATE or ALTER DATABASE RESTORE DEFAULT CALIBRATION statement. • LAST_DTT – Information about the <i>most recent</i> DTT calibration operation performed on the dbspace. That

Column name	Data type	Description
		<p>is, information about the most recent execution of either an ALTER DATABASE CALIBRATE or ALTER DATABASE RESTORE DEFAULT CALIBRATION statement.</p> <ul style="list-style-type: none"> • LAST_BACKUP – Information about the last backup, including date and time of the backup, the backup type, the files that were backed up, and the version of database server that performed the backup.
object_id	UNSIGNED INT	For any operation other than DTT and LAST_DTT, the value in this column will be 0. For DTT and LAST_DTT operations, this is the dbspace_id of the dbspace as defined in the SYSDBSpace system view.
sub_operation	CHAR(128)	<p>For any operation other than DTT and LAST_DTT, the value in this column will be a set of empty single quotes (""). For DTT and LAST_DTT operations, this column contains the type of sub-operation performed on the dbspace. Values include:</p> <ul style="list-style-type: none"> • DTT_SET – The dbspace calibration has been set. • DTT_UNSET – The dbspace calibration has been restored to the default setting.
version	CHAR(128)	The version and build number of the database server used to perform the operation.
platform	CHAR(128)	The operating system on which the operation was carried out.

Column name	Data type	Description
first_time	TIMESTAMP	The date and time the database was first started on a particular operating system with a particular version of the software.
last_time	TIMESTAMP	The most recent date and time the database was started on a particular operating system with a particular version of the software.
details	LONG VARCHAR	This column stores information such as command line options used to start the database server or the capability bits enabled for the database. This information is for use by technical support.

Constraints on underlying system table

PRIMARY KEY (operation, object_id, version, platform)

SYSIDX system view

Each row in the SYSIDX system view defines a logical index in the database. The underlying system table for this view is ISYSIDX.

Column name	Data type	Description
table_id	UNSIGNED INT	Uniquely identifies the table to which this index applies.
index_id	UNSIGNED INT	A unique number identifying the index within its table.
object_id	UNSIGNED BIGINT	The internal ID for the index, uniquely identifying it in the database.
phys_index_id	UNSIGNED INT	Identifies the underlying physical index used to implement the logical index. This value is NULL for indexes on temporary tables or remote tables. Otherwise, the value corresponds to the object_id of a physical index in the SYSPHYSIDX system view.

Column name	Data type	Description
dbspace_id	SMALLINT	The ID of the file in which the index is contained. This value corresponds to an entry in the SYSDBSpace system view.
file_id	SMALLINT	DEPRECATED. This column is present in SYSVIEW, but not in the underlying system table ISYSIDX. The contents of this column is the same as dbspace_id and is provided for compatibility. Use dbspace_id instead.
index_category	TINYINT	The type of index. Values include: <ul style="list-style-type: none"> • 1 – Primary key • 2 – Foreign key • 3 – Secondary index (includes unique constraints) • 4 – Text indexes
"unique"	TINYINT	Indicates whether the index is a unique index (1), a non-unique index (4), or a unique constraint (2). A unique index prevents two rows in the indexed table from having the same values in the index columns.
index_name	CHAR(128)	The name of the index.
not_enforced	CHAR(1)	For system use only.
file_id	SMALLINT	For system use only.

Constraints on underlying system table

PRIMARY KEY (table_id, index_id)

FOREIGN KEY (table_id) references SYS.ISYSTAB (table_id)

FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH UNIQUE FULL

FOREIGN KEY (table_id, phys_index_id) references SYS.ISYSPHYSIDX (table_id, phys_index_id)

UNIQUE Index (index_name, table_id, index_category)

SYSIDXCOL system view

Each row in the SYSIDXCOL system view describes one column of an index described in the SYSIDX system view. The underlying system table for this view is ISYSIDXCOL.

Column name	Data type	Description
table_id	UNSIGNED INT	Identifies the table to which the index applies.
index_id	UNSIGNED INT	Identifies the index to which the column applies. Together, table_id and index_id identify one index described in the SYSIDX system view.
sequence	SMALLINT	Each column in an index is assigned a unique number starting at 0. The order of these numbers determines the relative significance of the columns in the index. The most important column has sequence number 0.
column_id	UNSIGNED INT	Identifies which column of the table is indexed. Together, table_id and column_id identify one column described in the SYSCOLUMN system view.
"order"	CHAR(1)	Indicates whether the column in the index is kept in ascending(A) or descending(D) order. This value is NULL for text indexes.
primary_column_id	UNSIGNED INT	The ID of the primary key column that corresponds to this foreign key column. The value is NULL for non foreign key columns.

Constraints on underlying system table

```
PRIMARY KEY (table_id, index_id, column_id)
```

```
FOREIGN KEY (table_id, index_id) references SYS.ISYSIDX (table_id, index_id)
```

```
FOREIGN KEY (table_id, column_id) references SYS.ISYSTABCOL (table_id, column_id)
```

SYSINDEX compatibility view (deprecated)

The SYSINDEX view is provided for compatibility with older versions of the software that offered a SYSINDEX system table. However, the SYSINDEX system table has been replaced by the ISYSIDX system table, and its corresponding SYSIDX system view, which you should use instead.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSINDEX"
  as select b.table_id,
    b.index_id,
    b.object_id,
    p.root,
    b.dbspace_id,
    case b."unique"
    when 1 then 'Y'
    when 2 then 'U'
    when 3 then 'M'
    when 4 then 'N'
    when 5 then 'Y'
    else 'I'
    end as "unique",
    t.creator,
    b.index_name,
    r.remarks,
    10 as hash_limit,
    b.dbspace_id as file_id
  from(SYS.ISYSIDX as b left outer join SYS.ISYSPHYSIDX as p
on(b.table_id = p.table_id and b.phys_index_id = p.phys_index_id))
    left outer join SYS.ISYSREMARK as r on(b.object_id =
r.object_id)
    ,SYS.ISYSTAB as t
    where t.table_id = b.table_id
    and b.index_category = 3
```

SYSINDEXES consolidated view

Each row in the SYSINDEXES view describes one index in the database. As an alternative to this view, you could also use the SYSIDX and SYSIDXCOL system views.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSINDEXES"( icreator,
  iname, fname, creator, tname, indextype,
  colnames, interval, level_num )
  as select u.user_name, idx.index_name, dbs.dbspace_name, u.user_name,
    tab.table_name,
    case idx.index_category
```



```

when 1 then 'Primary Key'
when 2 then 'Foreign Key'
when 3 then(
    if idx."unique" = 4 then 'Non-unique'
    else if idx."unique" = 2 then 'UNIQUE constraint'
        else if idx."unique" = 5 then 'UNIQUE NULLS NOT DISTINCT'
            else 'UNIQUE'
        endif
    endif
endif) when 4 then 'Text Index' end,(select
list(string(c.column_name,
if idx."order" = 'A' then ' ASC' else ' DESC' endif) order by
ixc.table_id asc,ixc.index_id asc,ixc.sequence asc)
from SYS.ISYSIDXCOL as ixc
    join SYS.ISYSTABCOL as c on(
        c.table_id = ixc.table_id
        and c.column_id = ixc.column_id)
where ixc.index_id = idx.index_id
and ixc.table_id = idx.table_id),
0,0
from SYS.ISYSTAB as tab
    join SYS.ISYSDBSPACE as dbs on(tab.dbpace_id = dbs.dbpace_id)
    join SYS.ISYSIDX as idx on(idx.table_id = tab.table_id)
    join SYS.ISYSUSER as u on u.user_id = tab.creator

```

SYSINDEXES ASE Compatibility View

sysindexes contains one row for each clustered index, one row for each nonclustered index, one row for each table that has no clustered index, and one row for each table that contains text or image columns.

This table also contains one row for each function-based index or index created on a computed column.

This view is owned by user DBO.

See also

- *Tables in Each Adaptive Server Enterprise Database* on page 679
- *SYSCOLUMNS ASE Compatibility View* on page 586
- *SYS COMMENTS ASE Compatibility View* on page 587
- *SYSIQOBJECTS ASE Compatibility View* on page 619
- *SYSIQVINDE X ASE Compatibility View* on page 623
- *SYSOBJECTS ASE Compatibility View* on page 629
- *SYSTYPES ASE Compatibility View* on page 670
- *SYSUSERS ASE Compatibility View* on page 676

SYSINFO compatibility view (deprecated)

The SYSINFO view indicates the database characteristics, as defined when the database was created. It always contains only one row. This view is obtainable via built-in functions and is not kept in the catalog. Following is the definition for the SYSINFO view:

```
ALTER VIEW "SYS"."SYSINFO"( page_size,
    encryption,
    blank_padding,
    case_sensitivity,
    default_collation,
    database_version )
as select db_property('PageSize'),
    if db_property('Encryption') <> 'None' then 'Y' else 'N' endif,
    if db_property('BlankPadding') = 'On' then 'Y' else 'N' endif,
    if db_property('CaseSensitive') = 'On' then 'Y' else 'N' endif,
    db_property('Collation'),
    NULL
```

SYSIQBACKUPHISTORY System View

This view presents group information from ISYSIQBACKUPHISTORY in a readable format. Each row in this view describes a particular backup operation that finished successfully.

The view SYSIQBACKUP projects equivalent string values for columns type, subtype, and bkp_virtual.

Column name	Column type	Column constraint	Description
bu_id	unsigned bigint	NOT NULL	Transaction identifier of the checkpoint of the operation. Backup ID for backup operations.
bu_time	timestamp	NOT NULL	Time of backup operation that is recorded in backup record.
type	tinyint	NOT NULL	Backup type: 0 = FULL 1 = INCREMENTAL 2 = INCREMENTAL SINCE FULL

Column name	Column type	Column constraint	Description
selective_type	tinyint	NOT NULL	Backup subtype: 0 = ALL (backs up all dbfiles) 1 = READ/WRITE ONLY (backs up all read-write files) 2 = READ ONLY (backs up a particular read-only file)
virtual_type	tinyint	NOT NULL	Backup virtual type: 0 = NONE 1 = DECOUPLED 2 = ENCAPSULATED
dependson_id	unsigned bigint	NULL	NULL for FULL backup
cmd	long varchar	NOT NULL	Full text of command
creator	char(128)	NOT NULL	User who issued backup command
version	unsigned int	NOT NULL	Backup version

Constraints on underlying system table:

Primary key (bu_id)

See also

- *sp_iqbackupdetails Procedure* on page 364

SYSIQBACKUPHISTORYDETAIL System View

This view describes all the dbfile records present in the database at backup time. Each row in this view describes a particular backup operation that finished successfully.

It presents group information from `ISYSIQBACKUPHISTORYDETAIL` in a readable format. The column constraint for each column is NOT NULL.

Column name	Column type	Description
bu_id	unsigned bigint	Transaction identifier of the checkpoint of the operation. Backup ID for backup operation.
dbspace_id	smallint	The dbspace ID of which this dbfile record is associated.

Column name	Column type	Description
dbfile_id	smallint	The dbfile ID present in dbspace during ongoing backup operation
dbspace_rwstatus	char(1)	T indicates read-write
dbspace_createid	unsigned bigint	The transaction ID of the transaction that created the dbspace
dbspace_alterid	unsigned bigint	Transaction ID that marked the dbspace RO. If not marked, then the create ID
dbspace_online	char(1)	T indicates online
dbfile_rwstatus	char(1)	T indicates read-write
dbfile_createid	unsigned bigint	The transaction ID of the transaction that created this dbfile
dbfile_alterid	unsigned bigint	The transaction ID of the transaction that last altered the read-write status of this dbfile
is_backed_up	char(1)	Indicates that the dbfile is backed up in this backup
start_block	unsigned bigint	Start block for the dbfile
num_blocks	unsigned bigint	Total number of blocks in dbfile
num_blocks_backed_up	unsigned bigint	Total number of blocks backed up
dbspace_name	char(128)	Dbspace name
dbfile_name	char(128)	Logical file name of the dbfile
dbfile_path	long varchar	Physical path of the file

Constraints on underlying system table:

Primary key (bu_id, dbfile_id)

Foreign key (txn_id) references SYS.ISYSBACKUPHISTORY

SYSIQCOLUMN System View (Deprecated)

SYSIQCOLUMN has been replaced by the SYSIQTABCOL system view.

See also

- *SYSIQTABCOL System View* on page 621

SYSIQDBFILE System View

Presents group information from ISYSIQDBFILE in a readable format.

Note: This view replaces the deprecated system view SYSIQFILE.

Column name	Column type	Description
dbfile_id	small int	Unique ID for the dbfile
start_block	rowid	Number of the first block
block_count	rowid	Number of blocks for this file (dbspace)
reserve_size	rowid	Pre-allocated file system space for the dbspace
allocated	char(1)	Defines whether the segment is pre-allocated (T) or auto-allocated (F)
data_offset	unsigned int	Identifies the byte location of where the Sybase IQ data starts, relative to the beginning of the raw partition
create_time	timestamp	Date and time the file was created
last_modified	timestamp	Date and time the file was last modified
read_write	char(1)	T indicates read-write
online	char(1)	T indicates online
create_txn_id	unsigned bigint	Transaction ID that created the dbfile
alter_txn_id	unsigned bigint	Transaction ID that last modified read_write status
server_id	unsigned int	Multiplex server name
file_name	text	the IQ dbspace name used by the multiplex server to open the IQ dbspace

Constraints on underlying system table:

Foreign key (server_id) references SYS.ISYSIQMPXSERVER

Unique (server_id, file_name)

See also

- *SYSIQFILE System View (Deprecated)* on page 612

SYSIQDBSPACE System View

Presents group information from ISYSIQDBSPACE in a readable format.

Column name	Column type	Description
dbspace_id	small int	Each dbspace in a database is assigned a unique number (dbspace ID)
last_modified	timestamp	Time at which the dbspace's read-write status was last modified
segment_type	char(8)	Segment type: Main, Temp or Msg
read_write	char(1)	'T' – read writable, 'F' – read only
online	char(1)	'T' – online; 'F' – offline
create_txn_id	unsigned bigint	Transaction ID that create the dbspace
alter_txn_id	unsigned bigint	Transaction ID that last modified read_write status
striping_on	char(1)	'T' – disk striping on; 'F' – disk striping off
stripe_size_kb	unsigned int	Number of kilobytes written to each file of the dbspace before the disk striping algorithm moves to the next dbfile

Constraints on underlying system table:

Primary key (dbspace_id)

Foreign key (dbspace_id) references SYS.ISYSDBSPACE(dbspace_id)

SYSIQFILE System View (Deprecated)

SYSIQFILE has been replaced by the SYSIQDBFILE system view.

See also

- *SYSIQDBFILE System View* on page 611

SYSIQIDX System View

Presents group information from ISYSIQIDX in a readable format. Each row in the SYSIQIDX view describes an IQ index.

Note: This view replaces the deprecated system view SYSIQINDEX.

Column name	Column type	Description
table_id	unsigned int	The table number uniquely identifies the table to which this index applies
index_id	unsigned int	Each index for one particular table is assigned a unique index number
index_type	char(4)	Index type
index_owner	char(4)	Index owner
max_key	unsigned int	For internal use
identity_location	hs_vdorecid	For internal use
identity_size	unsigned int	For internal use
identity_location_size	unsigned int	For internal use
link_table_id	unsigned int	For internal use
link_index_id	unsigned int	For internal use
delimited_by	varchar(1024)	(WD indexes only) List of separators used to parse a column's string into the words to be stored in that column's WD index
limit	unsigned int	(WD indexes only) Maximum word length for WD index

Constraints on underlying system table:

Primary key (table_id, index_id)

Foreign key (table_id, index_id) references SYS.ISYIDX

Foreign key (link_table_id, link_index_id, table_id, index_id) references SYS.ISYSIDX

SYSIQINFO System View

Presents group information from ISYSIQINFO in a readable format.

The ISYSIQINFO system table indicates the database characteristics as defined when the Sybase IQ database was created using **CREATE DATABASE**. It always contains only one row.

Column name	Column type	Description
create_time	TIMESTAMP NOT NULL	Date and time that the database was created
update_time	TIMESTAMP NOT NULL	Date and time of the last update
file_format_version	UNSIGNED INT NOT NULL	File format number of files for this database
cat_format_version	UNSIGNED INT NOT NULL	Catalog format number for this database
sp_format_version	UNSIGNED INT NOT NULL	Stored procedure format number for this database
block_size	UNSIGNED INT NOT NULL	Block size specified for the database
chunk_size	UNSIGNED INT NOT NULL	Number of blocks per page as determined by the block size and page size specified for the database
file_format_date	CHAR(10) NOT NULL	Date when file format number was last changed
dbsig	BINARY(136) NOT NULL	Used internally by catalog
commit_txn_id	unsigned bigint	For internal use
rd_commit_txn_id	unsigned bigint	For internal use
multiplex name	CHAR(128) NULL	Name of the multiplex that this database is a member of

Column name	Column type	Description
last_multiplex_mode	TINYINT NULL	(Column unused in Sybase IQ 15.3) Mode of the server that last opened the catalog read-write. One of the following values. <ul style="list-style-type: none"> • 0 – Single Node. • 1 – Reader. • 2 – Coordinator. • 3 – Writer.

SYSIQJOINIDX System View

Presents group information from ISYSIQJOINIDX in a readable format. Each row in the SYSIQJOINIDX view describes an IQ join index.

Note: This view replaces the deprecated system view SYSIQJOININDEX.

Column name	Column type	Description
joinindex_id	unsigned int	Each join index is assigned a unique number that is the primary key.
jvt_id	unsigned int	For internal use.
dbspace_id	smallint	ID of the dbspace.
joinindex_name	char(128)	Defines the name of the join index.
joinindex_type	char(12)	For internal use.
creator	unsigned int	The number of the user that created the join index.
join_info_location	hs_vdorecid	For internal use.
join_info_loc_size	unsigned int	For internal use.
join_info_size	unsigned int	For internal use.
block_map	hs_blockmapidentity	For internal use.
block_map_size	unsigned int	For internal use.
vdo	hs_vdoidentity	For internal use.
vdo_size	unsigned_int	For internal use.

Column name	Column type	Description
commit_txn_id	unsigned bigint	For internal use.
txn_id	unsigned bigint	For internal use.
valid	char(1)	Indicates whether this join index needs to be synchronized. Y indicates that it does not require synchronization, N indicates that it does require synchronization.

Constraints on underlying system table:

Primary key (joinindex_id)

Foreign key (jvt_id) references SYS.ISYSTAB

Foreign key (dbspace_id) references SYS.ISYSDBSPACE

Foreign key (creator) references SYS.ISYSUSER

Unique (jvt_id, commit_txn_id, txn_id)

See also

- *SYSIQJOININDEX System View (Deprecated)* on page 616

SYSIQJOININDEX System View (Deprecated)

SYSIQJOININDEX has been replaced by the SYSIQJOINIDX system view.

See also

- *SYSIQJOINIDX System View* on page 615

SYSIQJOINIXCOLUMN System View

Presents group information from ISYSIQJOINIXCOLUMN in a readable format. Each row in the SYSIQJOINIXCOLUMN view describes an IQ join index.

```
ALTER VIEW "SYS"."SYSIQJOINIXCOLUMN"
as select * from SYS.ISYSIQJOINIXCOLUMN
```

Column name	Column type	Description
joinindex_id	unsigned bigint	Corresponds to a join index value in SYSIQJOINIDX.

Column name	Column type	Description
left_table_id	unsigned int	Corresponds to a table value in SYSTAB that forms the left side of the join operation.
left_column_id	unsigned int	Corresponds to a column value in SYSTABCOL that is part of the left side of the join.
join_type	char(4)	Only value currently supported is “=”.
right_table_id	unsigned int	Corresponds to a table value in SYSTAB that forms the right side of the join operation
right_column_id	unsigned int	Corresponds to a column value in SYSTABCOL that is part of the right side of the join.
order_num	unsigned int	For internal use.
left_order_num	unsigned int	For internal use.
right_order_num	unsigned int	For internal use.
key_type	char(8)	Defines the type of join on the keys. ‘NATURAL’ is a natural join, ‘KEY’ is a key join, ‘ON’ is a left outer/right outer/full join.
coalesce	char(1)	Not used.

Constraints on underlying system table:

Primary key (joinindex_id, left_table_id, left_column_id, right_table_id, right_column_id)

Foreign key (joinindex_id) references SYS.ISYSIQJOINIDX

Foreign key (left_table_id, column_id) references SYS.ISYSTABCOL

Foreign key (right_table_id, column_id) references SYS.ISYSTABCOL

SYSIQJOINIXTABLE System View

Presents group information from ISYSIQJOINIXTABLE in a readable format. Each row in the SYSIQJOINIXTABLE view describes an IQ join index.

```
ALTER VIEW "SYS"."SYSIQJOINIXTABLE"
as select * from SYS.ISYSIQJOINIXTABLE
```

Column name	Column type	Description
table_id	unsigned int	Corresponds to a table value in SYSTAB that is included in a join operation.
joinindex_id	unsigned bigint	Corresponds to a join index value in SYSIQJOINIDX.
active	unsigned int	Defines the number of times the table is used in the join index.

Constraints on underlying system table:

Primary key (table_id, joinindex_id)

Foreign key (table_id) references SYS.ISYSTAB

Foreign key (joinindex_id) references SYS.ISYSIQJOINIDX

SYSIQLOGICALSERVER System View

Presents information from the ISYSIQLOGICALSERVER table, which stores the logical server and the correspondence between logical server and associated logical server policy information.

See *Using Sybase IQ Multiplex > Multiplex Reference > System Views > SYSIQLOGICALSERVER System View.*

SYSIQLOGINPOLICYLSINFO System View

Presents group information from the ISYSIQLOGINPOLICYLSINFO table, which stores the login policy logical server assignment information.

See *Using Sybase IQ Multiplex > Multiplex Reference > System Views > SYSIQLOGINPOLICYLSINFO System View.*

SYSIQLSLOGINPOLICYOPTION System View

Presents group information from the ISYSIQLOGINPOLICYLSINFO table, which stores the login policy logical server assignment information.

See *Using Sybase IQ Multiplex > Multiplex Reference > System Views > SYSIQLSLOGINPOLICYOPTION System View.*

SYSIQLSMEMBER System View

Presents group information from the ISYSIQLSMEMBER table, which stores the logical server membership information.

See *Using Sybase IQ Multiplex > Multiplex Reference > System Views > SYSIQLSMEMBER System View.*

SYSIQLSMEMBERS Consolidated View

This view describes all user-defined logical server memberships.

See *Using Sybase IQ Multiplex > Multiplex Reference > System Views > SYSIQLSMEMBERS Consolidated View*.

SYSIQLSLOGINPOLICIES Consolidated View

This view describes all the logical server assignments from the login policies.

See *Using Sybase IQ Multiplex > Multiplex Reference > System Views > SYSIQLSLOGINPOLICIES Consolidated View*

SYSIQLSPOLICY System View

Presents group information from the ISYSIQLSPOLICY table, which stores the logical server policies.

See *Using Sybase IQ Multiplex > Multiplex Reference > System Views > SYSIQLSPOLICY System View*.

SYSIQLSPOLICYOPTION System View

Presents group information from the ISYSIQLSPOLICYOPTION table, which stores the logical server policy options.

See *Using Sybase IQ Multiplex > Multiplex Reference > System Views > SYSIQLSPOLICYOPTION System View*.

SYSIQMPXSERVER System View

Presents a readable version of the table ISYSIQMPXSERVER, which stores membership properties and version status data for the given multiplex node.

Users must be licensed for the Multiplex Grid Option to run secondary nodes. See *Using Sybase IQ Multiplex*.

SYSIQOBJECTS ASE Compatibility View

sysiqobjects presents one row for each system table, user table, view, procedure, trigger, event, join index, constraint, domain (sysdomain), domain (sysusertype), column, and index. This view is owned by user DBO.

See also

- *Tables in Each Adaptive Server Enterprise Database* on page 679
- *SYSCOLUMNS ASE Compatibility View* on page 586
- *SYS COMMENTS ASE Compatibility View* on page 587
- *SYSINDEXES ASE Compatibility View* on page 607

System Tables and Views

- *SYSIQINDEX ASE Compatibility View* on page 623
- *SYSOBJECTS ASE Compatibility View* on page 629
- *SYSTYPES ASE Compatibility View* on page 670
- *SYSUSERS ASE Compatibility View* on page 676

SYSIQPARTITIONCOLUMN System View

Presents group information from ISYSIQPARTITIONCOLUMN in a readable format.

```
ALTER VIEW "SYS"."SYSIQPARTITIONCOLUMN"  
as select * from SYS.ISYSIQPARTITIONCOLUMN
```

Each row in the SYSIQPARTITIONCOLUMN view describes a column in a partition described in the SYSIQPARTITION view in a partitioned table described in the SYSPARTITIONSCHEME view. SYSIQPARTITIONCOLUMN only describes partitions of columns that are not stored on the dbspace of the partition.

Column name	Column type	Description
partitioned_object_id	unsigned bigint	Unique ID assigned to each partitioned object (table)
partition_id	unsigned int	Identifies a partition in a partitioned table.
column_id	unsigned int	The column ID of the column.
dbspace_id	smallint	The dbspace ID of the dbspace where this column of the partition is stored.

Constraints on underlying system table:

Primary key (partitioned_object_id, partition_id, column_id)

Foreign key (partitioned_object_id, partition_id) references SYS.ISYSPARTITION

Foreign key (dbspace_id) references SYS.ISYSDBSPACE

SYSIQTAB System View

Presents group information from ISYSIQTAB in a readable format. Each row in the SYSIQTAB view describes an IQ table.

```
ALTER VIEW "SYS"."SYSIQTAB"  
as select * from SYS.ISYSIQTAB
```

Note: This view replaces the deprecated system view SYSIQTABLE.

Column name	Column type	Description
table_id	unsigned int	Each table is assigned a unique number (the table number) that is the primary key.
block_map	hs_blockmapidentity	For internal use.
block_map_size	unsigned int	For internal use.
vdo	hs_vdoidentity	For internal use.
vdo_size	unsigned int	For internal use.
info_location	hs_vdorecid	Not used. Always zero.
info_recid_size	unsigned int	Not used. Always zero.
info_location_size	unsigned int	Not used. Always zero.
commit_txn_id	unsigned bigint	For internal use.
txn_id	unsigned bigint	For internal use.
join_id	unsigned int	For internal use.
update_time	timestamp	Last date and time the IQ table was modified.

Constraints on underlying system table:

Primary key (table_id)

See also

- *SYSIQTABLE System View (Deprecated)* on page 622

SYSIQTABCOL System View

Presents group information from ISYSIQTABCOL in a readable format. Each row in the SYSIQTABCOL view describes a column in an IQ table.

```
ALTER VIEW "SYS"."SYSIQTABCOL"
as select * from SYS.ISYSIQTABCOL
```

Note: This view replaces the deprecated system view SYSIQCOLUMN.

Column name	Column type	Description
table_id	unsigned int	The table number uniquely identifies the table to which this column belongs. It corresponds to the table_id column of SYS-TAB.

Column name	Column type	Description
column_id	unsigned int	Each table starts numbering columns at 1. The order of column numbers determines the order that columns are displayed in the command select * from table.
link_table_id	unsigned int	For internal use.
link_column_id	unsigned int	For internal use.
max_length	unsigned int	Indicates the maximum length allowed by the column.
approx_unique_count	rowid	Approximate number of unique values (cardinality) of this column.
cardinality	rowid	The actual number of unique values (cardinality) of this column.
has_data	char(1)	Indicates that the column contains data (T/F).
has_original	char(1)	Indicates the join index has the original data (T/F).
original_not_null	char(1)	Indicates the join index column with the original data was NOT NULL (T/F).
original_unique	char(1)	Indicates the join index column with the original data was UNIQUE (T/F).

Constraints on underlying system table:

Primary key (table_id)

See also

- *SYSIQCOLUMN System View (Deprecated)* on page 610

SYSIQTAB System View (Deprecated)

SYSIQTAB has been replaced by the SYSIQTAB system view.

See also

- *SYSIQTAB System View* on page 620

SYSIQVINDEXT ASE Compatibility View

sysiqvindex provides one row for each non-FP IQ index.

This view is owned by user DBO.

See also

- *Tables in Each Adaptive Server Enterprise Database* on page 679
- *SYSCOLUMNS ASE Compatibility View* on page 586
- *SYSCOMMENTS ASE Compatibility View* on page 587
- *SYSINDEXES ASE Compatibility View* on page 607
- *SYSIQOBJECTS ASE Compatibility View* on page 619
- *SYSOBJECTS ASE Compatibility View* on page 629
- *SYSTYPES ASE Compatibility View* on page 670
- *SYSUSERS ASE Compatibility View* on page 676

SYSIXCOL compatibility view (deprecated)

The SYSIXCOL view is provided for compatibility with older versions of the software that offered a SYSIXCOL system table. However, the SYSIXCOL system table has been replaced by the ISYSIDXCOLD system table, and its corresponding SYSIDXCOLD system view. You should switch to using the SYSIDXCOLD system view.

Each row of the SYSIXCOL describes a column in an index. The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSIXCOL"
as select a.table_id,
       a.index_id,
       a.sequence,
       a.column_id,
       a."order"
from SYS.ISYSIDXCOLD as a
     ,SYS.ISYSIDX as b
where a.table_id = b.table_id
and a.index_id = b.index_id
and b.index_category = 3
```

SYSJAR system view

Each row in the SYSJAR system view defines a JAR file stored in the database. The underlying system table for this view is ISYSJAR.

Column name	Data type	Description
jar_id	INTEGER	A unique number identifying the JAR file.

Column name	Data type	Description
object_id	UNSIGNED BIGINT	The internal ID for the JAR file, uniquely identifying it in the database.
creator	UNSIGNED INT	The user number of the creator of the JAR file.
jar_name	LONG VARCHAR	The name of the JAR file.
jar_file	LONG VARCHAR	The external file name of the JAR file within the database.
update_time	TIMESTAMP	The time the JAR file was last updated.

Constraints on underlying system table

PRIMARY KEY (jar_id)

FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH UNIQUE FULL

UNIQUE Index (jar_name)

SYSJARCOMPONENT system view

Each row in the SYSJAR system view defines a JAR file component. The underlying system table for this view is ISYSJARCOMPONENT.

Column name	Data type	Description
component_id	INTEGER	The primary key containing the id of the component.
jar_id	INTEGER	A field containing the ID number of the JAR.
component_name	LONG VARCHAR	The name of the component.
component_type	CHAR(1)	The type of the component.
contents	LONG BINARY	The byte code of the JAR file.

Constraints on underlying system table

PRIMARY KEY (component_id)

FOREIGN KEY (jar_id) references SYS.ISYSJAR (jar_id)

SYSJAVACLASS system view

Each row in the SYSJAVACLASS system view describes one Java class stored in the database. The underlying system table for this view is ISYSJAVACLASS.

Column name	Data type	Description
class_id	INTEGER	The unique number for the Java class. Also the primary key for the table.
object_id	UNSIGNED BIGINT	The internal ID for the Java class, uniquely identifying it in the database.
creator	UNSIGNED INT	The user number of the creator of the class.
jar_id	INTEGER	The id of the JAR file from which the class came.
class_name	LONG VARCHAR	The name of the Java class.
public	CHAR(1)	Indicates whether the class is public (Y) or private (N).
component_id	INTEGER	The id of the component in the SYSJARCOMPONENT system view.
update_time	TIMESTAMP	The last update time of the class.

Constraints on underlying system table

```
PRIMARY KEY (class_id)
```

```
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH  
UNIQUE FULL
```

```
FOREIGN KEY (creator) references SYS.ISYSUSER (user_id)
```

```
FOREIGN KEY (component_id) references SYS.ISYSJARCOMPONENT  
(component_id)
```

SYSLOGINMAP system view

The SYSLOGINMAP system view contains one row for each user that can connect to the database using either an integrated login, or Kerberos login. As a security measure, only users

System Tables and Views

with DBA authority can view the contents of this view. The underlying system table for this view is `ISYSLOGINMAP`.

Column name	Data type	Description
<code>login_mode</code>	<code>TINYINT</code>	The type of login: 1 for integrated logins, 2 for Kerberos logins.
<code>login_id</code>	<code>VARCHAR(1024)</code>	Either the integrated login user profile name, or the Kerberos principal that maps to <code>database_uid</code> .
<code>object_id</code>	<code>UNSIGNED BIGINT</code>	A unique identifier, one for each mapping between user ID and database user ID.
<code>database_uid</code>	<code>UNSIGNED INT</code>	The database user ID to which the login ID is mapped.

Constraints on underlying system table

```
PRIMARY KEY (login_mode, login_id)
```

```
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH  
UNIQUE FULL
```

```
FOREIGN KEY (database_uid) references SYS.ISYSUSER (user_id)
```

SYSLOGINPOLICY system view

The underlying system table for this view is `ISYSLOGINPOLICY`.

Column name	Data type	Description
<code>login_policy_id</code>	<code>UNSIGNED BIGINT</code>	A unique identifier for the login policy.
<code>login_policy_name</code>	<code>CHAR(128)</code>	The name of the login policy.

Constraints on underlying system table

```
PRIMARY KEY (login_policy_id)
```

```
FOREIGN KEY (login_policy_id) references SYS.ISYSOBJECT (object_id)
```

```
UNIQUE Index (login_policy_name)
```

SYSLOGINPOLICYOPTION system view

The underlying system table for this view is ISYSLOGINPOLICYOPTION.

Column name	Data type	Description
login_policy_id	UNSIGNED BIGINT	A unique identifier for the login policy.
login_option_name	CHAR(128)	The name of the login policy.
login_option_value	LONG VARCHAR	The value of the login policy at the time it was created.

Constraints on underlying system table

```
PRIMARY KEY (login_policy_id, login_option_name)
```

```
FOREIGN KEY (login_policy_id) references SYS.ISYSLOGINPOLICY  
(login_policy_id)
```

SYSLOGINS ASE Compatibility View

This view is owned by user DBO. syslogins contains one row for each valid Adaptive Server Enterprise user account.

SYSMVOPTION system view

Each row in the SYSMVOPTION system view describes the setting of one option value for a materialized view or text index at the time of its creation. The name of the option can be found in the SYSMVOPTIONNAME system view. The underlying system table for this view is ISYSMVOPTION.

Column name	Data type	Description
view_object_id	UNSIGNED BIGINT	The object ID of the materialized view.
option_id	UNSIGNED INT	A unique number identifying the option in the database. To see the option name, see the SYSMVOPTIONNAME system view.
option_value	LONG VARCHAR	The value of the option when the materialized view was created.

Constraints on underlying system table

```
PRIMARY KEY (view_object_id, option_id)
```

```
FOREIGN KEY (view_object_id) references SYS.ISYSOBJECT (object_id)
```

```
FOREIGN KEY (option_id) references SYS.ISYSMVOPTIONNAME (option_id)
```

SYSMVOPTIONNAME system view

Each row in the SYSMVOPTION system view gives the name option value for a materialized view or text index at the time of its creation. The value for the option can be found in the SYSMVOPTION system view. The underlying system table for this view is ISYSMVOPTIONNAME.

Column name	Data type	Description
option_id	UNSIGNED INT	A number uniquely identifying the option in the database.
option_name	CHAR(128)	The name of the option.

Constraints on underlying system table

```
PRIMARY KEY (option_id)
```

```
UNIQUE Index (option_name)
```

SYSOBJECT system view

Each row in the SYSOBJECT system view describes a database object. The underlying system table for this view is ISYSOBJECT.

Column name	Data type	Description
object_id	UNSIGNED BIGINT	The internal ID for the object, uniquely identifying it in the database.

Column name	Data type	Description
status	TINYINT	<p>The status of the object. Values include:</p> <ul style="list-style-type: none"> • 1 (valid) – The object is available for use by the database server. This status is synonymous with ENABLED. That is, if you ENABLE an object, the status changes to VALID. • 2 (invalid) – An attempt to recompile the object after an internal operation has failed, for example, after a schema-altering modification to an object on which it depends. The database server continues to try to recompile the object whenever it is referenced in a statement. • 4 (disabled) – The object has been explicitly disabled by the user, for example using an ALTER TABLE...DISABLE VIEW DEPENDENCIES statement.
object_type	TINYINT	Type of object.
creation_time	TIMESTAMP	The date and time when the object was created.
object_type_str	CHAR (128)	Type of object.

Constraints on underlying system table

PRIMARY KEY (object_id)

SYSOBJECTS ASE Compatibility View

sysobjects contains one row for each table, view, stored procedure, extended stored procedure, log, rule, default, trigger, check constraint, referential constraint, computed column, function-based index key, and temporary object, and other forms of compiled objects.

This view is owned by user DBO.

It also contains one row for each partition condition ID when object type is N.

See also

- *Tables in Each Adaptive Server Enterprise Database* on page 679
- *SYSCOLUMNS ASE Compatibility View* on page 586
- *SYSCOMMENTS ASE Compatibility View* on page 587
- *SYSINDEXES ASE Compatibility View* on page 607
- *SYSIQOBJECTS ASE Compatibility View* on page 619
- *SYSIQVINDEXT ASE Compatibility View* on page 623
- *SYSTYPES ASE Compatibility View* on page 670
- *SYSUSERS ASE Compatibility View* on page 676

SYSOPTION system view

The SYSOPTION system view contains the options one row for each option setting stored in the database. Each user can have their own setting for a given option. In addition, settings for the PUBLIC user ID define the default settings to be used for users that do not have their own setting. The underlying system table for this view is ISYSOPTION.

Column name	Data type	Description
user_id	UNSIGNED INT	The user number to whom the option setting applies.
"option"	CHAR(128)	The name of the option.
"setting"	LONG VARCHAR	The current setting for the option.

Constraints on underlying system table

PRIMARY KEY (user_id, "option")

FOREIGN KEY (user_id) references SYS.ISYSUSER (user_id)

SYSOPTIONS consolidated view

Each row in the SYSOPTIONS view describes one option created using the SET command. Each user can have their own setting for each option. In addition, settings for the PUBLIC user define the default settings to be used for users that do not have their own setting.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSOPTIONS"( user_name,"option",setting )
as select u.user_name,opt."option",opt.setting
from SYS.ISYSOPTION as opt
join SYS.ISYSUSER as u on opt.user_id = u.user_id
```

SYSOPTSTAT system view

The SYSOPTSTAT system view stores the cost model calibration information as computed by the ALTER DATABASE CALIBRATE statement. The contents of this view are for internal

use only and are best accessed via the `sa_get_dtt` system procedure. The underlying system table for this view is `ISYSOPTSTAT`.

Column name	Data type	Description
stat_id	UNSIGNED INT	For system use only.
group_id	UNSIGNED INT	For system use only.
format_id	SMALLINT	For system use only.
data	LONG BINARY	For system use only.

Constraints on underlying system table

```
PRIMARY KEY (stat_id, group_id, format_id)
```

SYSPARTITION System View

Presents group information from `ISYSPARTITION` in a readable format.

Each row in the `SYSPARTITION` view describes a partitioned object (table or index) in the database.

```
ALTER VIEW "SYS"."SYSPARTITION"
as select * from SYS.ISYSPARTITION
```

Column name	Column type	Description
partitioned_object_id	unsigned bigint	Unique number assigned to each partitioned object (table)
partition_id	unsigned int	Identifies a partition in a partitioned table.
partition_object_id	unsigned bigint	Each table partition is an object itself and is assigned a unique number from the table object or index object.
partition_values	long varchar	Contains the upper bound for this range partition.
position	unsigned int	Ordinal number of partition.
partition_name	char(128)	Name of partition

Constraints on underlying system table:

Primary key (partitioned_object_id, partition_id)

Unique (partition_object_id, position)

Foreign key (partition_object_id) references `SYS.ISYSOBJECT`

Foreign key (partitioned_object_id) references `SYS.ISYSOBJECT`

SYSPARTITIONKEY System View

Presents group information from ISYSPARTITIONKEY in a readable format.

Each row in the SYSPARTITIONKEY view describes a partitioned object (table or index) in the database.

```
ALTER VIEW "SYS"."SYSPARTITIONKEY"
as select * from SYS.ISYSPARTITIONKEY
```

Column name	Column type	Description
partitioned_object_id	unsigned bigint	Each partitioned object (table) is assigned a unique object number.
column_id	unsigned int	The column ID identifies the table column as part of the partitioning key.
position	smallint	Position of this column in the partitioning key. Position is 0 based. A position of 0 indicates the 1st column in the partitioning key.

Constraints on underlying system table:

Primary key (partitioned_object_id, column_id)

Foreign key (partitioned_object_id) references SYS.ISYSOBJECT

SYSPARTITIONSCHEME System View

Presents group information from ISYSPARTITIONSCHEME in a readable format.

Each row in the SYSPARTITIONSCHEME view describes a partitioned object (table or index) in the database.

```
ALTER VIEW "SYS"."SYSPARTITIONSCHEME"
as select * from SYS.ISYSPARTITIONSCHEME
```

Column name	Column type	Description
partitioned_object_id	unsigned bigint	Each partitioned object (table) is assigned a unique number
partition_method	tinyint	The partitioning method for this table. Valid values: 1– for range Sybase IQ supports only range partitioning.
subpartition_method	tinyint	Reserved for future use.

Constraints on underlying system table:

Primary key (partitioned_object_id)

Foreign key (partitioned_object_id) references SYS.ISYSOBJECT

SYSphysIDX system view

Each row in the SYSphysIDX system view defines a physical index in the database. The underlying system table for this view is ISYSphysIDX.

Column name	Data type	Description
table_id	UNSIGNED INT	The object ID of the table to which the index corresponds.
phys_index_id	UNSIGNED INT	The unique number of the physical index within its table.
root	INTEGER	Identifies the location of the root page of the physical index in the database file.
key_value_count	UNSIGNED INT	The number of distinct key values in the index.
leaf_page_count	UNSIGNED INT	The number of leaf index pages.
depth	UNSIGNED SMALLINT	The depth (number of levels) of the physical index.
max_key_distance	UNSIGNED INT	For system use only.
seq_transitions	UNSIGNED INT	For system use only.
rand_transitions	UNSIGNED INT	For system use only.
rand_distance	UNSIGNED INT	For system use only.
allocation_bitmap	LONG VARBIT	For system use only.
long_value_bitmap	LONG VARBIT	For system use only.

Constraints on underlying system table

PRIMARY KEY (table_id, phys_index_id)

SYSprocAUTH consolidated view

Each row in the SYSprocAUTH view describes a set of privileges granted on a procedure. As an alternative, you can also use the SYSprocPERM system view.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

System Tables and Views

```
ALTER VIEW "SYS"."SYSPROCAUTH"( grantee,
creator,procname )
as select u1.user_name,u2.user_name,p.proc_name
from SYS.ISYSPROCEDURE as p
join SYS.ISYSPROCPerm as pp on(p.proc_id = pp.proc_id)
join SYS.ISYSUSER as u1 on u1.user_id = pp.grantee
join SYS.ISYSUSER as u2 on u2.user_id = p.creator
```

SYSPROCEDURE system view

Each row in the SYSPROCEDURE system view describes one procedure in the database. The underlying system table for this view is ISYSPROCEDURE.

Column name	Data type	Description
proc_id	UNSIGNED INT	Each procedure is assigned a unique number (the procedure number).
creator	UNSIGNED INT	The owner of the procedure.
object_id	UNSIGNED BIGINT	The internal ID for the procedure, uniquely identifying it in the database.
proc_name	CHAR(128)	The name of the procedure. One creator cannot have two procedures with the same name.
proc_defn	LONG VARCHAR	The definition of the procedure.
remarks	LONG VARCHAR	Remarks about the procedure. This value is stored in the ISYS-REMARK system table.
replicate	CHAR(1)	This property is for internal use only.
srv_id	UNSIGNED INT	If the procedure is a proxy for a procedure on a remote database server, indicates the remote server.
source	LONG VARCHAR	The preserved source for the procedure. This value is stored in the ISYSSOURCE system table.
avg_num_rows	FLOAT	Information collected for use in query optimization when the procedure appears in the FROM clause.

Column name	Data type	Description
avg_cost	FLOAT	Information collected for use in query optimization when the procedure appears in the FROM clause.
stats	LONG BINARY	Information collected for use in query optimization when the procedure appears in the FROM clause.

Constraints on underlying system table

PRIMARY KEY (proc_id)

FOREIGN KEY (srvid) references SYS.ISYSSERVER (srvid)

FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH
UNIQUE FULL

FOREIGN KEY (creator) references SYS.ISYSUSER (user_id)

UNIQUE Index (proc_name, creator)

SYSPROCARM system view

Each row in the SYSPROCARM system view describes one parameter to a procedure in the database. The underlying system table for this view is ISYSPROCARM.

Column name	Data type	Description
proc_id	UNSIGNED INT	Uniquely identifies the procedure to which the parameter belongs.
parm_id	SMALLINT	Each procedure starts numbering parameters at 1. The order of parameter numbers corresponds to the order in which they were defined. For functions, the first parameter has the name of the function and represents the return value for the function.

Column name	Data type	Description
parm_type	SMALLINT	The type of parameter is one of the following: <ul style="list-style-type: none"> • 0 – Normal parameter (variable) • 1 – Result variable - used with a procedure that returns result sets • 2 – SQLSTATE error value • 3 – SQLCODE error value • 4 – Return value from function
parm_mode_in	CHAR(1)	Indicates whether the parameter supplies a value to the procedure (IN or INOUT parameters).
parm_mode_out	CHAR(1)	Indicates whether the parameter returns a value from the procedure (OUT or INOUT parameters) or columns in the RESULT clause.
domain_id	SMALLINT	Identifies the data type for the parameter, by the data type number listed in the SYSDOMAIN system view.
width	BIGINT	Contains the length of a string parameter, the precision of a numeric parameter, or the number of bytes of storage for any other data type.
scale	SMALLINT	For numeric data types, the number of digits after the decimal point. For all other data types, the value of this column is 1.
user_type	SMALLINT	The user type of the parameter, if applicable.
parm_name	CHAR(128)	The name of the procedure parameter.

Column name	Data type	Description
"default"	LONG VARCHAR	Default value of the parameter. Provided for informational purposes only.
remarks	LONG VARCHAR	Always returns NULL. Provided to allow the use of previous versions of ODBC drivers with newer personal database servers.
base_type_str	VARCHAR(32767)	The annotated type string representing the physical type of the parameter.

Constraints on underlying system table

PRIMARY KEY (proc_id, parm_id)

FOREIGN KEY (proc_id) references SYS.ISYSPROCEDURE (proc_id)

FOREIGN KEY (domain_id) references SYS.ISYSDOMAIN (domain_id)

FOREIGN KEY (user_type) references SYS.ISYSUSERTYPE (type_id)

SYSROCPARMS consolidated view

Each row in the SYSROCPARMS view describes a parameter to a procedure in the database.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSROCPARMS"( creator,
    procname,parmname,parm_id,parmtime,parmmode,parmdomain,
    length,scale,"default",user_type )
as select
up.user_name,p.proc_name,pp.parm_name,pp.parm_id,pp.parm_type,
    if pp.parm_mode_in = 'Y' and pp.parm_mode_out = 'N' then 'IN'
    else if pp.parm_mode_in = 'N' and pp.parm_mode_out = 'Y' then
'OUT'
    else 'INOUT'
    endif
endif,dom.domain_name,pp.width,pp.scale,pp."default",ut.type_name
from SYS.SYSROCPARM as pp
    join SYS.ISYSPROCEDURE as p on p.proc_id = pp.proc_id
    join SYS.ISYSUSER as up on up.user_id = p.creator
    join SYS.ISYSDOMAIN as dom on dom.domain_id = pp.domain_id
    left outer join SYS.ISYSUSERTYPE as ut on ut.type_id =
pp.user_type
```

SYSROCPERM system view

Each row of the SYSROCPERM system view describes a user granted permission to execute a procedure. Only users who have been granted permission can execute a procedure. The underlying system table for this view is ISYSROCPERM.

Column name	Data type	Description
proc_id	UNSIGNED INT	The procedure number uniquely identifies the procedure for which permission has been granted.
grantee	UNSIGNED INT	The user number of the user receiving the permission.

Constraints on underlying system table

PRIMARY KEY (proc_id, grantee)

FOREIGN KEY (grantee) references SYS.ISYSUSER (user_id)

FOREIGN KEY (proc_id) references SYS.ISYSPROCEDURE (proc_id)

SYSROCS consolidated view

The SYSROCS view shows the procedure or function name, the name of its creator and any comments recorded for the procedure or function.

The tables and columns that make up this view are provided in the ALTER VIEW statement below.

```
ALTER VIEW "SYS"."SYSROCS"( creator,
    procname,remarks )
as select u.user_name,p.proc_name,r.remarks
    from SYS.ISYSPROCEDURE as p
        join SYS.ISYSUSER as u on u.user_id = p.creator
        left outer join SYS.ISYSREMARK as r on(p.object_id =
r.object_id)
```

SYSROXYTAB system view

Each row of the SYSROXYTAB system view describes the remote parameters of one proxy table. The underlying system table for this view is ISYSROXYTAB.

Column name	Data type	Description
table_object_id	UNSIGNED BIGINT	The object ID of the proxy table.
existing_obj	CHAR(1)	Indicates whether the proxy table previously existed on the remote server .

Column name	Data type	Description
srvid	UNSIGNED INT	The unique ID for the remote server associated with the proxy table.
remote_location	LONG VARCHAR	The location of the proxy table on the remote server.

Constraints on underlying system table

PRIMARY KEY (table_object_id)

FOREIGN KEY (table_object_id) references ISYSOBJECT (object_id)
MATCH UNIQUE FULL

FOREIGN KEY (srvid) references SYS.ISYSSERVER (srvid)

SYSPUBLICATION system view

Each row in the SYSPUBLICATION system view describes a publication. The underlying system table for this view is ISYSPUBLICATION.

Column name	Data type	Description
publication_id	UNSIGNED INT	A number uniquely identifying the publication.
object_id	UNSIGNED BIGINT	The internal ID for the publication, uniquely identifying it in the database.
creator	UNSIGNED INT	The owner of the publication.
publication_name	CHAR(128)	The name of the publication.
remarks	LONG VARCHAR	Remarks about the publication. This value is stored in the ISYS-REMARK system table.
type	CHAR(1)	This column is deprecated.

Column name	Data type	Description
sync_type	UNSIGNED INT	<p>The type of synchronization for the publication. Values include:</p> <ul style="list-style-type: none"> • 0 (logscan) – This is a regular publication that uses the transaction log to upload all relevant data that has changed since the last upload. • 1 (scripted upload) – For this publication, the transaction log is ignored and the upload is defined by the user using stored procedures. Information about the stored procedures is stored in the ISYSSYNCSRIPT system table. • 2 (download only) – This is a download-only publication; no data is uploaded.

Constraints on underlying system table

```
PRIMARY KEY (publication_id)
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH
UNIQUE FULL
FOREIGN KEY (creator) references SYS.ISYSUSER (user_id)
UNIQUE Index (publication_name, creator)
```

SYSPUBLICATIONS consolidated view

Each row in the SYSPUBLICATIONS view describes a publication.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSPUBLICATIONS"
as select u.user_name as creator,
p.publication_name,
r.remarks,
p.type,
case p.sync_type
when 0 then 'logscan'
when 1 then 'scripted upload'
when 2 then 'download only'
```

```

else 'invalid'
end as sync_type
from SYS.ISYSPUBLICATION as p
  join SYS.ISYSUSER as u on u.user_id = p.creator
  left outer join SYS.ISYSREMARK as r on(p.object_id =
r.object_id)

```

SYSREMARK system view

Each row in the SYSREMARK system view describes a remark (or comment) for an object. The underlying system table for this view is ISYSREMARK.

Column	Data type	Description
object_id	UNSIGNED BIGINT	The internal ID for the object that has an associated remark.
remarks	LONG VARCHAR	The remark or comment associated with the object.

Constraints on underlying system table

```
PRIMARY KEY (object_id)
```

```
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH
UNIQUE FULL
```

SYSREMOTEOPTION System View

Each row in the SYSREMOTEOPTION system view describes the value of a SQL Remote message link parameter. The underlying system table for this view is ISYSREMOTEOPTION.

Some columns in this view contain potentially sensitive data. For that reason, access to this view is restricted to users with DBA authority. The SYSREMOTEOPTION2 view provides public access to the data in this view except for the potentially sensitive columns.

The SYSREMOTEOPTION view is a SQL Anywhere system view. See *SQL Anywhere Server - SQL Reference > Views > System views > SYSREMOTEOPTION system view*.

Note: This reference points to SQL Anywhere documentation.

SYSREMOTEOPTION2 consolidated view

Presents, in a more readable format, the columns from SYSREMOTEOPTION and SYSREMOTEOPTIONTYPE that do not contain sensitive data.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSREMOTEOPTION2"
  as select ISYSREMOTEOPTION.option_id,
```

```

    ISYSREMOTEOPTION.user_id,
    SYS.HIDE_FROM_NON_DBA(ISYSREMOTEOPTION.setting) as setting
from SYS.ISYSREMOTEOPTION

```

SYSREMOTEOPTIONS consolidated view

Each row of the SYSREMOTEOPTIONS view describes the values of a message link parameter. Some columns in this view contain potentially sensitive data. For that reason, access to this view is restricted to users with DBA authority. The SYSREMOTEOPTION2 view provides public access to the insensitive data.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```

ALTER VIEW "SYS"."SYSREMOTEOPTIONS"
as select srt.type_name,
    sup.user_name,
    srot."option",
    SYS.HIDE_FROM_NON_DBA(sro.setting) as setting
from SYS.ISYSREMOTETYPE as srt
    ,SYS.ISYSREMOTEOPTIONTYPE as srot
    ,SYS.ISYSREMOTEOPTION as sro
    ,SYS.ISYSUSER as sup
where srt.type_id = srot.type_id
and srot.option_id = sro.option_id
and sro.user_id = sup.user_id

```

SYSREMOTEOPTIONTYPE system view

Each row in the SYSREMOTEOPTIONTYPE system view describes one of the message link parameters. The underlying system table for this view is ISYSREMOTEOPTIONTYPE.

Column	Data type	Description
option_id	UNSIGNED INT	An identification number for the message link parameter.
type_id	SMALLINT	An identification number for the message type that uses the parameter.
"option"	VARCHAR(128)	The name of the message link parameter.

Constraints on underlying system table

```

PRIMARY KEY (option_id)
FOREIGN KEY (type_id) references SYS.ISYSREMOTETYPE (type_id)

```

SYSRE MOTETYPE system view

The SYSRE MOTETYPE system view contains information about SQL Remote. The underlying system table for this view is ISYSRE MOTETYPE.

Column name	Data type	Description
type_id	SMALLINT	Identifies which of the message systems supported by SQL Remote is to be used to send messages to the user.
object_id	UNSIGNED BIGINT	The internal ID for the remote type, uniquely identifying it in the database.
type_name	CHAR(128)	The name of the message system supported by SQL Remote.
publisher_address	LONG VARCHAR	The address of the remote database publisher.
remarks	LONG VARCHAR	Remarks about the remote type. This value is stored in the ISYSRE MARK system table.

Constraints on underlying system table

```
PRIMARY KEY (type_id)
```

```
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH  
UNIQUE FULL
```

```
UNIQUE Index (type_name)
```

SYSRE MOTETYPES consolidated view

Each row of the SYSRE MOTETYPES view describes one of the message types, including the publisher address.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSRE MOTETYPES"  
as select rt.type_id,rt.type_name,rt.publisher_address,rm.remarks  
from SYS.ISYSRE MOTETYPE as rt  
left outer join SYS.ISYSRE MARK as rm on(rt.object_id =  
rm.object_id)
```

SYSREMOTUSER system view

Each row in the SYSREMOTUSER system view describes a user ID with REMOTE permissions (a subscriber), together with the status of messages that were sent to and from that user. The underlying system table for this view is ISYSREMOTUSER.

Column name	Data type	Description
user_id	UNSIGNED INT	The user number of the user with REMOTE permissions.
consolidate	CHAR(1)	Indicates whether the user was granted CONSOLIDATE permissions (Y) or REMOTE permissions (N).
type_id	SMALLINT	Identifies which of the message systems is used to send messages to the user.
address	LONG VARCHAR	The address to which messages are to be sent. The address must be appropriate for the address_type.
frequency	CHAR(1)	How frequently messages are sent.
send_time	TIME	The next time messages are to be sent to this user.
log_send	UNSIGNED BIGINT	Messages are sent only to subscribers for whom log_send is greater than log_sent.
time_sent	TIMESTAMP	The time the most recent message was sent to this subscriber.
log_sent	UNSIGNED BIGINT	The log offset for the most recently sent operation.
confirm_sent	UNSIGNED BIGINT	The log offset for the most recently confirmed operation from this subscriber.
send_count	INTEGER	How many messages have been sent.
resend_count	INTEGER	Counter to ensure that messages are applied only once at the subscriber database.

Column name	Data type	Description
time_received	TIMESTAMP	The time when the most recent message was received from this subscriber.
log_received	UNSIGNED BIGINT	The log offset in the subscriber's database for the operation that was most recently received at the current database.
confirm_received	UNSIGNED BIGINT	The log offset in the subscriber's database for the most recent operation for which a confirmation message has been sent.
receive_count	INTEGER	How many messages have been received.
rereceive_count	INTEGER	Counter to ensure that messages are applied only once at the current database.

Constraints on underlying system table

PRIMARY KEY (user_id)

FOREIGN KEY (user_id) references SYS.ISYSUSER (user_id)

FOREIGN KEY (type_id) references SYS.ISYSREMOTETYPE (type_id)

UNIQUE Index (type_id, address)

SYSREMOTEUSERS consolidated view

Each row of the SYSREMOTEUSERS view describes a user ID with REMOTE permissions (a subscriber), together with the status of messages that were sent to and from that user.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSREMOTEUSERS"
  as select
u.user_name,r consolidate,t.type_name,r.address,r.frequency,
  r.send_time,
  (if r.frequency = 'A' then null else if r.frequency = 'P' then
    if r.time_sent is null then current timestamp
    else(select min(minutes(a.time_sent,60*hour(a.send_time)
      +minute(seconds(a.send_time,59))))
      from SYS.ISYSREMOTEUSER as a where a.frequency = 'P'
      and a.send_time = r.send_time)
    endif
  else if current date+r.send_time
```

System Tables and Views

```
> coalesce(r.time_sent,current timestamp) then
    current date+r.send_time else current date+r.send_time+1
endif
endif endif) as next_send,
r.log_send,r.time_sent,r.log_sent,r.confirm_sent,r.send_count,
r.resend_count,r.time_received,r.log_received,
r.confirm_received,r.receive_count,r.rereceive_count
from SYS.ISYSREMOTEUSER as r
    join SYS.ISYSUSER as u on(u.user_id = r.user_id)
    join SYS.ISYSREMOUETYPE as t on(t.type_id = r.type_id)
```

SYSSCHEDULE System View

Each row in the SYSSCHEDULE system view describes a time at which an event is to fire, as specified by the SCHEDULE clause of CREATE EVENT.

The underlying system table for this view is ISYSSCHEDULE.

The SYSSCHEDULE view is a SQL Anywhere system view. See *SQL Anywhere Server - SQL Reference > Views > System views > SYSSCHEDULE system view*.

Note: This reference points to SQL Anywhere documentation.

SYSSERVER system view

Each row in the SYSSERVER system view describes a remote server. The underlying system table for this view is ISYSSERVER.

Note: Previous versions of the catalog contained a SYSSERVERS system table. That table has been renamed to be ISYSSERVER (without an 'S'), and is the underlying table for this view.

Column name	Data type	Description
srvid	UNSIGNED INT	An identifier for the remote server.
srvname	VARCHAR(128)	The name of the remote server.
srvclass	LONG VARCHAR	The server class, as specified in the CREATE SERVER statement.
srvinfo	LONG VARCHAR	Server information.
srvreadonly	CHAR(1)	Whether the server is read-only.

Constraints on underlying system table

```
PRIMARY KEY (srvid)
```


SYSSOURCE system view

Each row in the SYSSOURCE system view contains the source code, if applicable, for an object listed in the SYSOBJECT system view. The underlying system table for this view is ISYSSOURCE.

Column name	Data type	Description
object_id	UNSIGNED BIGINT	The internal ID for the object whose source code is being defined.
source	LONG VARCHAR	This column contains the original source code for the object if the preserve_source_format database option is On when the object was created.

Constraints on underlying system table

```
PRIMARY KEY (object_id)
```

```
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH  
UNIQUE FULL
```

SYSSQLSERVERTYPE system view

The SYSSQLSERVERTYPE system view contains information relating to compatibility with Adaptive Server Enterprise. The underlying system table for this view is ISYSSQLSERVERTYPE.

Column name	Data type	Description
ss_user_type	SMALLINT	The Adaptive Server Enterprise user type.
ss_domain_id	SMALLINT	The Adaptive Server Enterprise domain ID.
ss_type_name	VARCHAR (30)	The Adaptive Server Enterprise type name.
primary_sa_domain_id	SMALLINT	The corresponding SQL Anywhere primary domain ID.
primary_sa_user_type	SMALLINT	The corresponding SQL Anywhere primary user type.

Constraints on underlying system table

```
PRIMARY KEY (ss_user_type)
```

SYSSUBPARTITIONKEY System View

This view is reserved for future use. Sybase IQ 15.4 does not support subpartitioning.

SYSSUBSCRIPTION system view

Each row in the SYSSUBSCRIPTION system view describes a subscription from one user ID (which must have REMOTE permissions) to one publication. The underlying system table for this view is ISYSSUBSCRIPTION.

Column name	Data type	Description
publication_id	UNSIGNED INT	The identifier for the publication to which the user ID is subscribed.
user_id	UNSIGNED INT	The ID of the user who is subscribed to the publication.
subscribe_by	CHAR(128)	The value of the SUBSCRIBE BY expression, if any, for the subscription.
created	UNSIGNED BIGINT	The offset in the transaction log at which the subscription was created.
started	UNSIGNED BIGINT	The offset in the transaction log at which the subscription was started.

Constraints on underlying system table

```
PRIMARY KEY (publication_id, user_id, subscribe_by)
```

```
FOREIGN KEY (publication_id) references SYS.ISYSPUBLICATION (publication_id)
```

```
FOREIGN KEY (user_id) references SYS.ISYSUSER (user_id)
```

SYSSUBSCRIPTIONS consolidated view

Each row describes a subscription from one user ID (which must have REMOTE permissions) to one publication.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSSUBSCRIPTIONS"
as select p.publication_name,u.user_name,s.subscribe_by,s.created,
s.started
from SYS.ISYSSUBSCRIPTION as s
```

```

join SYS.ISYSPUBLICATION as p on(p.publication_id =
s.publication_id)
join SYS.ISYSUSER as u on u.user_id = s.user_id

```

SYSSYNC system view

The SYSSYNC system view contains information relating to synchronization. Some columns in this view contain potentially sensitive data. For that reason, access to this view is restricted to users with DBA authority. The SYSSYNC2 view provides public access to the data in this view except for the potentially sensitive columns. The underlying system table for this view is ISYSSYNC.

Column name	Data type	Description
sync_id	UNSIGNED INT	A number that uniquely identifies the row.
type	CHAR(1)	This value is always D.
publication_id	UNSIGNED INT	A publication_id found in the SYSPUBLICATION system view.
progress	UNSIGNED BIGINT	The log offset of the last successful upload.
site_name	CHAR(128)	A user name.
"option"	LONG VARCHAR	Synchronization options.
server_connect	LONG VARCHAR	The address or URL of the server.
server_conn_type	LONG VARCHAR	The communication protocol, such as TCP/IP, to use when synchronizing.
last_download_time	TIMESTAMP	Indicates the last time a download stream was received from the server.
last_upload_time	TIMESTAMP	Indicates the last time (measured at the server) that information was successfully uploaded. The default is jan-1-1900.
created	UNSIGNED BIGINT	The log offset at which the subscription was created.

Column name	Data type	Description
log_sent	UNSIGNED BIGINT	The log progress up to which information has been uploaded. It is not necessary that an acknowledgement of the upload be received for the entry in this column to be updated.
generation_number	INTEGER	For file-base downloads, the last generation number received for this subscription. The default is 0.
extended_state	VARCHAR(1024)	For internal use only.
script_version	CHAR(128)	Indicates the script version used by the CREATE and ALTER SYNCHRONIZATION SUBSCRIPTION statements and the START SYNCHRONIZATION SCHEMA CHANGE statement.
subscription_name	CHAR (128)	The name of the subscription.

Constraints on underlying system table

```
PRIMARY KEY (sync_id)
FOREIGN KEY (publication_id) references SYS.ISYSPUBLICATION (publication_id)
UNIQUE Index (publication_id, site_name)
UNIQUE Index (subscription_name)
```

SYSSYNC2 consolidated view

The SYSSYNC2 view provides public access to the data found in the SYSSYNC system view—information relating to synchronization—without exposing potentially sensitive data.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSSYNC2"
as select ISYSSYNC.sync_id,
        ISYSSYNC.type,
        ISYSSYNC.publication_id,
        ISYSSYNC.progress,
        ISYSSYNC.site_name,
        SYS.HIDE_FROM_NON_DBA(ISYSSYNC."option") as "option",
        SYS.HIDE_FROM_NON_DBA(ISYSSYNC.server_connect) as
```

```
server_connect,
  ISYSSYNC.server_conn_type,
  ISYSSYNC.last_download_time,
  ISYSSYNC.last_upload_time,
  ISYSSYNC.created,
  ISYSSYNC.log_sent,
  ISYSSYNC.generation_number,
  ISYSSYNC.extended_state,
  ISYSSYNC.script_version,
  ISYSSYNC.subscription_name
from SYS.ISYSSYNC
```

SYSSYNCPUBLICATIONDEFAULTS consolidated view

The SYSSYNCPUBLICATIONDEFAULTS view provides the default synchronization settings associated with publications involved in synchronization.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSSYNCPUBLICATIONDEFAULTS"
  as select s.sync_id,
  p.publication_name,
  SYS.HIDE_FROM_NON_DBA(s."option") as "option",
  SYS.HIDE_FROM_NON_DBA(s.server_connect) as server_connect,
  s.server_conn_type
  from SYS.ISYSSYNC as s join SYS.ISYSPUBLICATION as p
on(p.publication_id = s.publication_id) where
  s.site_name is null
```

SYSSYNCS consolidated view

The SYSSYNCS view contains information relating to synchronization. Some columns in this view contain potentially sensitive data. For that reason, access to this view is restricted to users with DBA authority.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSSYNCS"
  as select p.publication_name,s.progress,s.site_name,
  SYS.HIDE_FROM_NON_DBA(s."option") as "option",
  SYS.HIDE_FROM_NON_DBA(s.server_connect) as server_connect,
  s.server_conn_type,s.last_download_time,
  s.last_upload_time,s.created,s.log_sent,s.generation_number,
  s.extended_state
  from SYS.ISYSSYNC as s
  left outer join SYS.ISYSPUBLICATION as p
  on p.publication_id = s.publication_id
```

SYSSYNCSCRIPT system view

Each row in the SYSSYNCSCRIPT system view identifies a stored procedure for scripted upload. This view is almost identical to the SYSSYNCSCRIPTS view, except that the values in this view are in their raw format.

The underlying system table for this view is ISYSSYNCSCRIPT.

Column name	Data type	Description
pub_object_id	UNSIGNED BIGINT	The object ID of the publication to which the script belongs.
table_object_id	UNSIGNED BIGINT	The object ID of the table to which the script applies.
type	UNSIGNED INT	The type of upload procedure.
proc_object_id	UNSIGNED BIGINT	The object ID of the stored procedure to use for the publication.

Constraints on underlying system table

```
PRIMARY KEY (pub_object_id, table_object_id, type)
```

```
FOREIGN KEY (pub_object_id) references SYS.ISYSOBJECT (object_id)
```

```
FOREIGN KEY (table_object_id) references SYS.ISYSOBJECT (object_id)
```

```
FOREIGN KEY (proc_object_id) references SYS.ISYSOBJECT (object_id)
```

SYSSYNCSCRIPTS consolidated view

Each row in the SYSSYNCSCRIPTS view identifies a stored procedure for scripted upload. This view is almost identical to the SYSSYNCSCRIPT system view, except that the values are in human-readable format, as opposed to raw data.

```
ALTER VIEW "SYS"."SYSSYNCSCRIPTS"
as select p.publication_name,
       t.table_name,
       case s.type
       when 0 then 'upload insert'
       when 1 then 'upload delete'
       when 2 then 'upload update'
       else 'unknown'
       end as type,
       c.proc_name
from SYS.ISYSSYNCSCRIPT as s
     join SYS.ISYSPUBLICATION as p on p.object_id = s.pub_object_id
     join SYS.ISYSTAB as t on t.object_id = s.table_object_id
     join SYS.ISYSPROCEDURE as c on c.object_id = s.proc_object_id
```

SYSSYNC SUBSCRIPTIONS consolidated view

The SYSSYNC SUBSCRIPTIONS view contains the synchronization settings associated with synchronization subscriptions.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSSYNC SUBSCRIPTIONS"
as select s.sync_id,
p.publication_name,
s.progress,
s.site_name,
SYS.HIDE_FROM_NON_DBA(s."option") as "option",
SYS.HIDE_FROM_NON_DBA(s.server_connect) as server_connect,
s.server_conn_type,
s.last_download_time,
s.last_upload_time,
s.created,
s.log_sent,
s.generation_number,
s.extended_state
from SYS.ISYSSYNC as s join SYS.ISYSPUBLICATION as p
on(p.publication_id = s.publication_id)
where s.publication_id is not null and
s.site_name is not null and exists
(select 1 from SYS.SYSSYNCUSERS as u
where s.site_name = u.site_name)
```

SYSSYNCUSERS consolidated view

A view of synchronization settings associated with synchronization users.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSSYNCUSERS"
as select ISYSSYNC.sync_id,
ISYSSYNC.site_name,
SYS.HIDE_FROM_NON_DBA(ISYSSYNC."option") as "option",
SYS.HIDE_FROM_NON_DBA(ISYSSYNC.server_connect) as
server_connect,
ISYSSYNC.server_conn_type
from SYS.ISYSSYNC where
ISYSSYNC.publication_id is null
```

SYSTAB system view

Each row of the SYSTAB system view describes one table or view in the database. Additional information for views can be found in the SYSVIEW system view. The underlying system table for this view is ISYSTAB.

Column name	Data type	Description
table_id	UNSIGNED INT	Each table is assigned a unique number (the table number).
dbspace_id	SMALLINT	A value indicating which dbspace contains the table.
count	UNSIGNED BIGINT	The number of rows in the table or materialized view. This value is updated during each successful checkpoint. This number is used to optimize database access. The count is always 0 for a non-materialized view or remote table.
creator	UNSIGNED INT	The user number of the owner of the table or view.
table_page_count	INTEGER	The total number of main pages used by the underlying table.
ext_page_count	INTEGER	The total number of extension pages used by the underlying table.
commit_action	INTEGER	For global temporary tables, 0 indicates that the ON COMMIT PRESERVE ROWS clause was specified when the table was created, 1 indicates that the ON COMMIT DELETE ROWS clause was specified when the table was created (the default behavior for temporary tables), and 3 indicates that the NOT TRANSACTIONAL clause was specified when the table was created. For non-temporary tables, commit_action is always 0.

Column name	Data type	Description
share_type	INTEGER	For global temporary tables, 4 indicates that the SHARE BY ALL clause was specified when the table was created, and 5 indicates that the SHARE BY ALL clause was <i>not</i> specified when the table was created. For non-temporary tables, share_type is always 5 because the SHARE BY ALL clause cannot be specified when creating non-temporary tables.
object_id	UNSIGNED BIGINT	The object ID of the table.
last_modified_at	TIMESTAMP	The time at which the data in the table was last modified. This column is only updated at checkpoint time.
last_modified_tsn	UNSIGNED BIGINT	A sequence number assigned to the transaction that modified the table.
file_id	SMALLINT	DEPRECATED. This column is present in SYSVIEW, but not in the underlying system table ISYSTAB. The contents of this column is the same as dbspace_id and is provided for compatibility. Use dbspace_id instead.
table_name	CHAR(128)	The name of the table or view. One creator cannot have two tables or views with the same name.

System Tables and Views

Column name	Data type	Description
table_type	TINYINT	The type of table or view. Values include: <ul style="list-style-type: none"> • 1 – Base table • 2 – Materialized view • 3 – Global temporary table • 4 – Local temporary table • 5 – Text index base table • 6 – Text index global temporary table • 21 – View
replicate	CHAR(1)	This value is for internal use only.
server_type	TINYINT	The location of the data for the underlying table. Values include: <ul style="list-style-type: none"> • 1 – Local server • 3 – Remote server
tab_page_list	LONG VARBIT	For internal use only. The set of pages that contain information for the table, expressed as a bitmap.
ext_page_list	LONG VARBIT	For internal use only. The set of pages that contain row extensions and large object (LOB) pages for the table, expressed as a bitmap.
pct_free	UNSIGNED INT	The PCT_FREE specification for the table, if one has been specified; otherwise, NULL.
clustered_index_id	UNSIGNED INT	The ID of the clustered index for the table. If none of the indexes are clustered, then this field is NULL.
encrypted	CHAR(1)	Whether the table or materialized view is encrypted.

Column name	Data type	Description
table_type_str	CHAR(9)	Readable value for table_type. Values include: <ul style="list-style-type: none"> • BASE – Base table • MAT VIEW – Materialized view • GBL TEMP – Global temporary table • VIEW – View

Constraints on underlying system table

FOREIGN KEY (dbspace_id) references SYS.ISYSDBSPACE (dbspace_id)

FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id)

PRIMARY KEY (table_id)

FOREIGN KEY (creator) references SYS.ISYSUSER (user_id)

UNIQUE Index (table_name, creator)

SYSTABAUTH consolidated view

The SYSTABAUTH view contains information from the SYSTABLEPERM system view, but in a more readable format.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSTABAUTH" ( grantor,
  grantee, screator, stname, tcreator, tname,
  selectauth, insertauth, deleteauth,
  updateauth, updatecols, alterauth, referenceauth )
as select u1.user_name, u2.user_name, u3.user_name, tab1.table_name,
  u4.user_name, tab2.table_name, tp.selectauth, tp.insertauth,
  tp.deleteauth, tp.updateauth, tp.updatecols, tp.alterauth,
  tp.referenceauth
  from SYS.ISYSTABLEPERM as tp
    join SYS.ISYSUSER as u1 on u1.user_id = tp.grantor
    join SYS.ISYSUSER as u2 on u2.user_id = tp.grantee
    join SYS.ISYSTAB as tab1 on tab1.table_id = tp.stable_id
    join SYS.ISYSUSER as u3 on u3.user_id = tab1.creator
    join SYS.ISYSTAB as tab2 on tab2.table_id = tp.stable_id
    join SYS.ISYSUSER as u4 on u4.user_id = tab2.creator
```

SYSTABCOL system view

The SYSTABCOL system view contains one row for each column of each table and view in the database. The underlying system table for this view is ISYSTABCOL.

Column name	Data type	Description
table_id	UNSIGNED INT	The object ID of the table or view to which the column belongs.
column_id	UNSIGNED INT	The ID of the column. For each table, column numbering starts at 1. The column_id value determines the order of columns in the result set when SELECT * is used. It also determines the column order for an INSERT statement when a list of column names is not provided.
domain_id	SMALLINT	The data type for the column, indicated by a data type number listed in the SYSDOMAIN system view.
nulls	CHAR(1)	Indicates whether NULL values are allowed in the column.
width	UNSIGNED INT	The length of a string column, the precision of numeric columns, or the number of bytes of storage for any other data type.
scale	SMALLINT	The number of digits after the decimal point for NUMERIC or DECIMAL data type columns. For string columns, a value of 1 indicates character-length semantics and 0 indicates byte-length semantics.
object_id	UNSIGNED BIGINT	The object ID of the table column.
max_identity	BIGINT	The largest value of the column, if it is an AUTOINCREMENT, IDENTITY, or GLOBAL AUTOINCREMENT column.

Column name	Data type	Description
column_name	CHAR(128)	The name of the column.
"default"	LONG VARCHAR	The default value for the column. This value, if specified, is only used when an INSERT statement does not specify a value for the column.
user_type	SMALLINT	The data type, if the column is defined using a user-defined data type.
column_type	CHAR(1)	The type of column (C=computed column, and R=other columns).
compressed	TINYINT	Whether this column is stored in a compressed format.
collect_stats	TINYINT	Whether the system automatically collects and updates statistics on this column.
inline_max	SMALLINT	The maximum number of bytes of a BLOB to store in a row. A NULL value indicates that either the default value has been applied, or that the column is not a character or binary type. A non-NULL inline_max value corresponds to the INLINE value specified for the column using the CREATE TABLE or ALTER TABLE statement.
inline_long	SMALLINT	The number of duplicate bytes of a BLOB to store in a row if the BLOB size exceeds the inline_max value. A NULL value indicates that either the default value has been applied, or that the column is not a character or binary type. A non-NULL inline_long value corresponds to the PREFIX value specified for the column using the CREATE TABLE or ALTER TABLE statement.

Column name	Data type	Description
lob_index	TINYINT	Whether to build indexes on BLOB values in the column that exceed an internal threshold size (approximately eight database pages). A NULL value indicates either that the default is applied, or that the column is not BLOB type. A value of 1 indicates that indexes will be built. A value of 0 indicates that no indexes will be built. A non-NULL lob_index value corresponds to whether INDEX or NO INDEX was specified for the column using the CREATE TABLE or ALTER TABLE statement.
base_type_str	VARCHAR(32,767)	The annotated type string representing the physical type of the column.

Constraints on underlying system table

```
PRIMARY KEY (table_id, column_id)
FOREIGN KEY (table_id) references SYS.ISYSTAB (table_id)
FOREIGN KEY (domain_id) references SYS.ISYSDOMAIN (domain_id)
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH
UNIQUE FULL
FOREIGN KEY (user_type) references SYS.ISYSUSERTYPE (type_id)
```

SYSTABLE compatibility view (deprecated)

The SYSTABLE view is provided for compatibility with older versions of the software that offered a SYSTABLE system table. However, the SYSTABLE system table has been replaced by the ISYSTAB system table, and its corresponding SYSTAB system view, which you should use instead.

Each row of SYSTABLE view describes one table in the database.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSTABLE"
as select b.table_id,
b.file_id,
```

```

b.count,
0 as first_page,
b.commit_action as last_page,
COALESCE(ph.root,0) as primary_root,
b.creator,
0 as first_ext_page,
0 as last_ext_page,
b.table_page_count,
b.ext_page_count,
b.object_id,
b.table_name,
b.table_type_str as table_type,
v.view_def,
r.remarks,
b.replicate,
p.existing_obj,
p.remote_location,
'T' as remote_objtype,
p.srvid,
case b.server_type
when 1 then 'SA'
when 2 then 'IQ'
when 3 then 'OMNI'
else 'INVALID'
end as server_type,
10 as primary_hash_limit,
0 as page_map_start,
s.source,
b."encrypted"
from SYS.SYSTAB as b
  left outer join SYS.ISYSREMARK as r on(b.object_id =
r.object_id)
  left outer join SYS.ISYSSOURCE as s on(b.object_id =
s.object_id)
  left outer join SYS.ISYSVIEW as v on(b.object_id =
v.view_object_id)
  left outer join SYS.ISYSPROXYTAB as p on(b.object_id =
p.table_object_id)
  left outer join(SYS.ISYSIDX as i left outer join SYS.ISYSPHYSIDX
as ph on(i.table_id = ph.table_id and i.phys_index_id =
ph.phys_index_id)
    on(b.table_id = i.table_id and i.index_category = 1 and
i.index_id = 0)

```

SYSTABLEPERM system view

Permissions given by the GRANT statement are stored in the SYSTABLEPERM system view. Each row in this view corresponds to one table, one user ID granting the permission (grantor)

System Tables and Views

and one user ID granted the permission (grantee). The underlying system table for this view is ISYSTABLEPERM.

Column name	Data type	Description
stable_id	UNSIGNED INT	The table number of the table or view to which the permissions apply.
grantee	UNSIGNED INT	The user number of the user ID receiving the permission.
grantor	UNSIGNED INT	The user number of the user ID granting the permission.
selectauth	CHAR(1)	Indicates whether SELECT permission has been granted. Possible values are Y, N, or G. See the Remarks area below for more information about what these values mean.
insertauth	CHAR(1)	Indicates whether INSERT permission has been granted. Possible values are Y, N, or G. See the Remarks area below for more information about what these values mean.
deleteauth	CHAR(1)	Indicates whether DELETE permission has been granted. Possible values are Y, N, or G. See the Remarks area below for more information about what these values mean.
updateauth	CHAR(1)	Indicates whether UPDATE permission has been granted for all columns in the table. Possible values are Y, N, or G. See the Remarks area below for more information about what these values mean.

Column name	Data type	Description
updatecols	CHAR(1)	Indicates whether UPDATE permission has only been granted for some of the columns in the underlying table. If updatecols has the value Y, there will be one or more rows in the SYSCOLPERM system view granting update permission for the columns.
alterauth	CHAR(1)	Indicates whether ALTER permission has been granted. Possible values are Y, N, or G. See the Remarks area below for more information about what these values mean.
referenceauth	CHAR(1)	Indicates whether REFERENCE permission has been granted. Possible values are Y, N, or G. See the Remarks area below for more information about what these values mean.

Remarks

There are several types of permission that can be granted. Each permission can have one of the following three values.

- **N** – No, the grantee has not been granted this permission by the grantor.
- **Y** – Yes, the grantee has been given this permission by the grantor.
- **G** – The grantee has been given this permission and can grant the same permission to another user.

Note: The grantee might have been given permission for the same table by another grantor. If so, this information would be found in a different row of the SYSTABLEPERM system view.

Constraints on underlying system table

PRIMARY KEY (stable_id, grantee, grantor)

FOREIGN KEY (stable_id) references SYS.ISYSTAB (table_id)

FOREIGN KEY (grantor) references SYS.ISYSUSER (user_id)

FOREIGN KEY (grantee) references SYS.ISYSUSER (user_id)

SYSTEXTCONFIG system view

Each row in the SYSTEXTCONFIG system view describes one text configuration object, for use with the full text search feature. The underlying system table for this view is ISYSTEXTCONFIG.

Column name	Data type	Description
object_id	UNSIGNED BIGINT	The object ID for the text configuration object.
creator	UNSIGNED INT	The creator of the text configuration object.
term_breaker	TINYINT	The algorithm used to separate a string into terms or words. Values are 0 for GENERIC and 1 for NGRAM. With GENERIC, any string of one or more alphanumeric characters separated by non-alphanumerics are treated as a term. NGRAM is for approximate matching or for documents that do not use a whitespace to separate terms.
stemmer	TINYINT	For internal use only.
min_term_length	TINYINT	The minimum length, in characters, allowed for a term. Terms that are shorter than min_term_length are ignored. The MINIMUM TERM LENGTH setting is only meaningful for the GENERIC term breaker. For NGRAM text indexes, the setting is ignored.
max_term_length	TINYINT	For GENERIC text indexes, the maximum length, in characters, allowed for a term. Terms that are longer than max_term_length are ignored. For NGRAM text indexes, this is the length of the n-grams into which terms are broken.
collation	CHAR(128)	For internal use only.

Column name	Data type	Description
text_config_name	CHAR(128)	The name of the text configuration object.
prefilter	LONG VARCHAR	The function and library name for an external prefilter library.
postfilter	LONG VARCHAR	For internal use only.
char_stoplist	LONG VARCHAR	Terms to ignore when performing a full text search on CHAR columns. These terms are also omitted from text indexes. This column is used when the text configuration object is created from default_char.
nchar_stoplist	LONG NVARCHAR	Terms to ignore when performing a full text search on NCHAR columns. These terms are also omitted from text indexes. This column is used when the text configuration object is created from default_nchar.
external_term_breaker	LONG VARCHAR	The function and library name for an external term breaker library.

Constraints on underlying system table

PRIMARY KEY (object_id)

FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH UNIQUE FULL

FOREIGN KEY (creator) references SYS.ISYSUSER (user_id)

UNIQUE Index (creator, text_config_name)

SYSTEXTIDX system view

Each row in the SYSTEXTIDX system view describes one text index. The underlying system table for this view is ISYSTEXTIDX.

Column name	Data type	Description
index_id	UNSIGNED BIGINT	The object ID of the text index in SYSIDX.
sequence	UNSIGNED INT	For internal use only.

Column name	Data type	Description
status	UNSIGNED INT	For internal use only.
text_config	UNSIGNED BIGINT	The object ID of the text configuration object in SYSTEXTCONFIG.
next_handle	UNSIGNED INT	For internal use only.
last_handle	UNSIGNED INT	For internal use only.
deleted_length	UNSIGNED BIGINT	The total size of deleted indexed values in the text index.
pending_length	UNSIGNED BIGINT	The total size of indexed values that will be added to the text index at the next refresh.
refresh_type	TINYINT	The type of refresh. One of: <ul style="list-style-type: none"> • 1 – MANUAL • 2 – AUTO • 3 – IMMEDIATE
refresh_interval	UNSIGNED INT	The AUTO REFRESH interval, in minutes.
last_refresh	TIMESTAMP	The time of the last refresh.

Constraints on underlying system table

PRIMARY KEY (index_id, sequence)

FOREIGN KEY (index_id) references SYS.ISYSOBJECT (object_id)

FOREIGN KEY (text_config) references SYS.ISYSTEXTCONFIG (object_id)

SYSTEXTIDXTAB system view

Each row in the SYSTEXTIDXTAB system view describes a generated table that is part of a text index. The underlying system table for this view is ISYSTEXTIDXTAB.

Column name	Data type	Description
index_id	UNSIGNED BIGINT	For internal use only.
sequence	UNSIGNED INT	For internal use only.
table_type	UNSIGNED INT	For internal use only.
table_id	UNSIGNED INT	For internal use only.

Constraints on underlying system table

PRIMARY KEY (index_id, sequence, table_type)

FOREIGN KEY (index_id, sequence) references SYS.ISYSTEXTIDX
(index_id, sequence)

FOREIGN KEY (table_id) references SYS.ISYSTAB (table_id)

SYSTRIGGER system view

Each row in the SYSTRIGGER system view describes one trigger in the database. This view also contains triggers that are automatically created for foreign key definitions which have a referential triggered action (such as ON DELETE CASCADE). The underlying system table for this view is ISYSTRIGGER.

Column name	Data type	Description
trigger_id	UNSIGNED INT	A unique number for the trigger in the SYSTRIGGER view.
table_id	UNSIGNED INT	The table ID of the table to which this trigger belongs.
object_id	UNSIGNED BIGINT	The object ID for the trigger in the database.
event	CHAR(1)	The operation that causes the trigger to fire. <ul style="list-style-type: none"> • A – INSERT, DELETE • B – INSERT, UPDATE • C – UPDATE COLUMNS • D – DELETE • E – DELETE, UPDATE • I – INSERT • M – INSERT, DELETE, UPDATE • U – UPDATE

Column name	Data type	Description
trigger_time	CHAR(1)	<p>The time when the trigger fires relative to the event.</p> <ul style="list-style-type: none"> • A – AFTER (row-level trigger) • B – BEFORE (row-level trigger) • I – INSTEAD OF (row-level trigger) • K – INSTEAD OF (statement-level trigger) • R – RESOLVE • S – AFTER (statement-level trigger)
trigger_order	SMALLINT	<p>The order in which are fired when there are multiple triggers of the same type (insert, update, or delete) set to fire at the same time (applies to BEFORE or AFTER triggers only, only).</p>
foreign_table_id	UNSIGNED INT	<p>The ID of the table containing a foreign key definition that has a referential triggered action (such as ON DELETE CASCADE). The foreign_table_id value reflects the value of ISY-SIDX.table_id.</p>
foreign_key_id	UNSIGNED INT	<p>The ID of the foreign key for the table referenced by foreign_table_id. The foreign_key_id value reflects the value of ISY-SIDX.index_id.</p>
referential_action	CHAR(1)	<p>The action defined by a foreign key. This single-character value corresponds to the action that was specified when the foreign key was created.</p> <ul style="list-style-type: none"> • C – CASCADE • D – SET DEFAULT • N – SET NULL • R – RESTRICT

Column name	Data type	Description
trigger_name	CHAR(128)	The name of the trigger. One table cannot have two triggers with the same name.
trigger_defn	LONG VARCHAR	The command that was used to create the trigger.
remarks	LONG VARCHAR	Remarks about the trigger. This value is stored in the ISYSREMARK system table.
source	LONG VARCHAR	The SQL source for the trigger. This value is stored in the ISYSOURCE system table.

Constraints on underlying system table

```
PRIMARY KEY (trigger_id)
```

```
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH  
UNIQUE FULL
```

```
FOREIGN KEY (table_id) references SYS.ISYSTAB (table_id)
```

```
FOREIGN KEY fkey_index (foreign_table_id, foreign_key_id) references  
SYS.ISYSIDX (table_id, index_id)
```

```
UNIQUE Index (table_id, event, trigger_time, trigger_order)
```

```
UNIQUE Index (trigger_name, table_id)
```

```
UNIQUE Index (table_id, foreign_table_id, foreign_key_id, event)
```

SYSTRIGGERS consolidated view

Each row in the SYSTRIGGERS view describes one trigger in the database. This view also contains triggers that are automatically created for foreign key definitions which have a referential triggered action (such as ON DELETE CASCADE).

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSTRIGGERS"( owner,  
    trigname,tname,event,trigtime,trigdefn )  
as select u.user_name,trig.trigger_name,tab.table_name,  
    if trig.event = 'I' then 'INSERT'  
    else if trig.event = 'U' then 'UPDATE'  
    else if trig.event = 'C' then 'UPDATE'  
    else if trig.event = 'D' then 'DELETE'  
    else if trig.event = 'A' then 'INSERT,DELETE'  
    else if trig.event = 'B' then 'INSERT,UPDATE'  
    else if trig.event = 'E' then 'DELETE,UPDATE'
```

```

        else 'INSERT,DELETE,UPDATE'
        endif
    endif
endif
endif
endif
endif
endif
endif,if trig.trigger_time = 'B' or trig.trigger_time = 'P' then
'BEFORE'
    else if trig.trigger_time = 'A' or trig.trigger_time = 'S' then
'AFTER'
        else if trig.trigger_time = 'R' then 'RESOLVE'
        else 'INSTEAD OF'
        endif
    endif
endif, trig.trigger_defn
from SYS.ISYSTRIGGER as trig
    join SYS.ISYSTAB as tab on(tab.table_id = trig.table_id)
    join SYS.ISYSUSER as u on u.user_id = tab.creator where
trig.foreign_table_id is null

```

SYSTYPEMAP system view

The SYSTYPEMAP system view contains the compatibility mapping values for entries in the SYSSQLSERVERTYPE system view. The underlying system table for this view is ISYSTYPEMAP.

Column name	Data type	Description
ss_user_type	SMALLINT	Contains the Adaptive Server Enterprise user type.
sa_domain_id	SMALLINT	Contains the corresponding SQL Anywhere domain_id.
sa_user_type	SMALLINT	Contains the corresponding SQL Anywhere user type.
nullable	CHAR(1)	Whether the type allows NULL values.

Constraints on underlying system table

```
FOREIGN KEY (sa_domain_id) references SYS.ISYSDOMAIN (domain_id)
```

SYSTYPES ASE Compatibility View

systypes contains one row for each system-supplied and user-defined datatype. Domains (defined by rules) and defaults are given, if they exist.

This view is owned by user DBO. You cannot alter the rows that describe system-supplied datatypes.

See also

- *Tables in Each Adaptive Server Enterprise Database* on page 679
- *SYSCOLUMNS ASE Compatibility View* on page 586
- *SYSCOMMENTS ASE Compatibility View* on page 587
- *SYSINDEXES ASE Compatibility View* on page 607
- *SYSIQOBJECTS ASE Compatibility View* on page 619
- *SYSIQVINDEXT ASE Compatibility View* on page 623
- *SYSOBJECTS ASE Compatibility View* on page 629
- *SYSUSERS ASE Compatibility View* on page 676

SYSUSER system view

Each row in the SYSUSER system view describes a user in the system. The underlying system table for this view is ISYSUSER.

Column name	Data type	Description
user_id	UNSIGNED INT	A unique identifier for the user assigned to the login policy.
object_id	UNSIGNED BIGINT	A unique identifier for the user in the database.
user_name	CHAR(128)	The login name for the user.
password	BINARY(128)	The password for the user.
login_policy_id	UNSIGNED BIGINT	A unique identifier for the login policy.
expired_password_on_login	TINYINT	Indicates if the user's password expires at the next login.
password_creation_time	TIMESTAMP	The time the user's password was created.
failed_login_attempts	UNSIGNED INT	The number of times a user can fail to log in before the account is locked.
last_login_time	TIMESTAMP	The time the user last logged in.

Constraints on underlying system table

```
PRIMARY KEY (user_id)
```

```
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH  
UNIQUE FULL
```

```
FOREIGN KEY (login_policy_id) references SYS.ISYSLOGINPOLICY  
(login_policy_id)
```

```
UNIQUE Index (user_name)
```

SYSUSERAUTH compatibility view (deprecated)

The SYSUSERAUTH view is provided for compatibility with older versions of the software. Use the SYSUSERAUTHORITY system view instead.

Each row of the SYSUSERAUTH view describes a user, without exposing their user_id. Instead, each user is identified by their user name. Because this view displays passwords, this view does not have PUBLIC select permission.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSUSERAUTH" ( name,
    password, resourceauth, dbaauth, scheduleauth, user_group )
as select
SYSUSERPERM.user_name, SYSUSERPERM.password, SYSUSERPERM.resourceauth
, SYSUSERPERM.dbaauth, SYSUSERPERM.scheduleauth, SYSUSERPERM.user_grou
P
from SYS.SYSUSERPERM
```

SYSUSERAUTHORITY system view

Each row of SYSUSERAUTHORITY system view describes an authority granted to one user ID. The underlying system table for this view is ISYSUSERAUTHORITY.

Column name	Data type	Description
user_id	UNSIGNED INT	The ID of the user to whom the authority belongs.
auth	VARCHAR(20)	The authority granted to the user.

Constraints on underlying system table

```
PRIMARY KEY (user_id, auth)
```

```
FOREIGN KEY (user_id) references SYS.ISYSUSER (user_id)
```

SYSUSERLIST compatibility view (deprecated)

The SYSUSERAUTH view is provided for compatibility with older versions of the software.

Each row of the SYSUSERLIST view describes a user, without exposing their user_id and password. Each user is identified by their user name.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSUSERLIST" ( name,
    resourceauth, dbaauth, scheduleauth, user_group )
```

```

as select
SYSUSERPERM.user_name, SYSUSERPERM.resourceauth, SYSUSERPERM.dbauth,
SYSUSERPERM.scheduleauth, SYSUSERPERM.user_group
from SYS.SYSUSERPERM

```

SYSUSERMESSAGE system view

Each row in the SYSUSERMESSAGE system view holds a user-defined message for an error condition. The underlying system table for this view is ISYSUSERMESSAGE.

Note: Previous versions of the catalog contained a SYSUSERMESSAGES system table. That table has been renamed to be ISYSUSERMESSAGE (without an 'S'), and is the underlying table for this view.

Column name	Data type	Description
error	INTEGER	A unique identifying number for the error condition.
uid	UNSIGNED INT	The user number that defined the message.
description	VARCHAR(255)	The message corresponding to the error condition.
langid	SMALLINT	Reserved.

Constraints on underlying system table

```
FOREIGN KEY (uid) references SYS.ISYSUSER (user_id)
```

```
UNIQUE Constraint (error, langid)
```

SYSUSEROPTIONS consolidated view

The SYSUSEROPTIONS view contains the option settings that are in effect for each user. If a user has no setting for an option, this view displays the public setting for the option.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```

ALTER VIEW "SYS"."SYSUSEROPTIONS"( user_name,
"option",setting )
as select u.user_name,
o."option",
isnull((select s.setting
from SYS.ISYSOPTION as s
where s.user_id = u.user_id
and s."option" = o."option"),
o.setting)
from SYS.SYSOPTION as o,SYS.ISYSUSER as u
where o.user_name = 'PUBLIC'

```

SYSUSERPERM compatibility view (deprecated)

This view is deprecated because it only shows the authorities and permissions available in previous versions. You should change your application to use the SYSUSERAUTHORITY system view instead.

Each row of the SYSUSERPERM view describes one user ID.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSUSERPERM"
  as select b.user_id,
    b.object_id,
    b.user_name,
    b.password,
    if exists(select * from SYS.ISYSUSERAUTHORITY
      where ISYSUSERAUTHORITY.user_id = b.user_id and
      ISYSUSERAUTHORITY.auth = 'RESOURCE') then
      'Y' else 'N' endif as resourceauth,
    if exists(select * from SYS.ISYSUSERAUTHORITY
      where ISYSUSERAUTHORITY.user_id = b.user_id and
      ISYSUSERAUTHORITY.auth = 'DBA') then
      'Y' else 'N' endif as dbaauth,
    'N' as scheduleauth,
    if exists(select * from SYS.ISYSUSERAUTHORITY
      where ISYSUSERAUTHORITY.user_id = b.user_id and
      ISYSUSERAUTHORITY.auth = 'PUBLISH') then
      'Y' else 'N' endif as publishauth,
    if exists(select * from SYS.ISYSUSERAUTHORITY
      where ISYSUSERAUTHORITY.user_id = b.user_id and
      ISYSUSERAUTHORITY.auth = 'REMOTE DBA') then
      'Y' else 'N' endif as remotedbauth,
    if exists(select * from SYS.ISYSUSERAUTHORITY
      where ISYSUSERAUTHORITY.user_id = b.user_id and
      ISYSUSERAUTHORITY.auth = 'GROUP') then
      'Y' else 'N' endif as user_group,
    r.remarks
  from SYS.ISYSUSER as b
  left outer join SYS.ISYSREMARK as r on(b.object_id =
  r.object_id)
```

SYSUSERPERMS compatibility view (deprecated)

This view is deprecated because it only shows the authorities and permissions available in previous versions. You should change your application to use the SYSUSERAUTHORITY and SYSUSER system views instead.

Similar to the SYSUSERPERM view, each row of the SYSUSERPERMS view describes one user ID. However, password information is not included. All users are allowed to read from this view.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSUSERPERMS"
  as select
SYSUSERPERM.user_id,SYSUSERPERM.user_name,SYSUSERPERM.resourceauth,
SYSUSERPERM.dbauth,

SYSUSERPERM.scheduleauth,SYSUSERPERM.user_group,SYSUSERPERM.publish
auth,SYSUSERPERM.remotedbaauth,SYSUSERPERM.remarks
  from SYS.SYSUSERPERM
```

SYSUSERTYPE system view

Each row in the SYSUSERTYPE system view holds a description of a user-defined data type. The underlying system table for this view is ISYSUSERTYPE.

Column name	Data type	Description
type_id	SMALLINT	A unique identifying number for the user-defined data type.
creator	UNSIGNED INT	The user number of the owner of the data type.
domain_id	SMALLINT	The data type on which this user defined data type is based, indicated by a data type number listed in the SYSDOMAIN system view.
nulls	CHAR(1)	Whether the user-defined data type allows nulls. Possible values are Y, N, or U. A value of U indicates that nullability is unspecified.
width	BIGINT	The length of a string column, the precision of a numeric column, or the number of bytes of storage for any other data type.
scale	SMALLINT	The number of digits after the decimal point for numeric data type columns, and zero for all other data types.
type_name	CHAR(128)	The name for the data type.
"default"	LONG VARCHAR	The default value for the data type.

Column name	Data type	Description
"check"	LONG VARCHAR	The CHECK condition for the data type.
base_type_str	VARCHAR(32767)	The annotated type string representing the physical type of the user type.

Constraints on underlying system table

```
PRIMARY KEY (type_id)
FOREIGN KEY (creator) references SYS.ISYSUSER (user_id)
FOREIGN KEY (domain_id) references SYS.ISYSDOMAIN (domain_id)
UNIQUE Constraint (type_name)
```

SYSUSERS ASE Compatibility View

sysusers contains one row for each user allowed in the database, and one row for each group or roles.

This view is owned by user DBO.

See also

- *Tables in Each Adaptive Server Enterprise Database* on page 679
- *SYSCOLUMNS ASE Compatibility View* on page 586
- *SYSCOMMENTS ASE Compatibility View* on page 587
- *SYSINDEXES ASE Compatibility View* on page 607
- *SYSIQOBJECTS ASE Compatibility View* on page 619
- *SYSIQVINDEXT ASE Compatibility View* on page 623
- *SYSOBJECTS ASE Compatibility View* on page 629
- *SYSTYPES ASE Compatibility View* on page 670

SYSVIEW System View

Each row in the SYSVIEW system view describes a view in the database.

You can find additional information about views in the SYSTAB system view. The underlying system table for this view is ISYSVIEW.

You can also use the **sa_materialized_view_info** system procedure for a readable format of the information for materialized views. Materialized views are only supported for SQL Anywhere tables in the IQ catalog store.

The SYSVIEW view is a SQL Anywhere system view. See *SQL Anywhere Server - SQL Reference > Views > System views > SYSVIEW system view*.

Note: This reference points to SQL Anywhere documentation.

SYSVIEWS consolidated view

Each row of the `SYSVIEWS` view describes one view, including its view definition.

The tables and columns that make up this view are provided in the SQL statement below. To learn more about a particular table or column, use the links provided beneath the view definition.

```
ALTER VIEW "SYS"."SYSVIEWS"( vcreator,
    viewname,viewtext )
as select u.user_name,t.table_name,v.view_def
    from SYS.ISYSTAB as t
        join SYS.ISYSVIEW as v on(t.object_id = v.view_object_id)
        join SYS.ISYSUSER as u on(u.user_id = t.creator)
```

SYSWEBSERVICE system view

Each row in the `SYSWEBSERVICE` system view holds a description of a web service. The underlying system table for this view is `ISYSWEBSERVICE`.

Column name	Data type	Description
service_id	UNSIGNED INT	A unique identifying number for the web service.
object_id	UNSIGNED BIGINT	The ID of the webservice.
service_name	CHAR(128)	The name assigned to the web service.
service_type	VARCHAR(40)	The type of the service; for example, RAW, HTTP, XML, SOAP, or DISH.
auth_required	CHAR(1)	Whether all requests must contain a valid user name and password.
secure_required	CHAR(1)	Whether insecure connections, such as HTTP, are to be accepted, or only secure connections, such as HTTPS.
url_path	CHAR(1)	Controls the interpretation of URLs.

Column name	Data type	Description
user_id	UNSIGNED INT	If authentication is enabled, identifies the user, or group of users, that have permission to use the service. If authentication is disabled, specifies the account to use when processing requests.
parameter	LONG VARCHAR	A prefix that identifies the SOAP services to be included in a DISH service.
statement	LONG VARCHAR	A SQL statement that is always executed in response to a request. If NULL, arbitrary statements contained in each request are executed instead. Ignored for services of type DISH.
remarks	LONG VARCHAR	Remarks about the webservice. This value is stored in the ISYS-REMARK system table.
enabled	CHAR(1)	Indicates whether the web service is currently enabled or disabled (see CREATE SERVICE).

Constraints on underlying system table

```
PRIMARY KEY (service_id)
FOREIGN KEY (object_id) references SYS.ISYSOBJECT (object_id) MATCH
UNIQUE FULL
UNIQUE Constraint (service_name)
```

Transact-SQL Compatibility Views

Adaptive Server Enterprise and Sybase IQ have different system catalogs, reflecting the different uses for the two products.

In Adaptive Server Enterprise, there is a single master database containing a set of system tables holding information that applies to all databases on the server. Many databases may exist within the master database, and each has additional system tables associated with it.

In Sybase IQ, each database exists independently, and contains its own system tables. There is no master database that contains system information on a collection of databases. Each server may run several databases at a time, dynamically loading and unloading each database as needed.

The Adaptive Server Enterprise and Sybase IQ system catalogs are different. The Adaptive Server Enterprise system tables and views are owned by the special user `dbo`, and exist partly in the master database, partly in the **sybsecurity** database, and partly in each individual database; the Sybase IQ system tables and views are owned by the special user `SYS` and exist separately in each database.

To assist in preparing compatible applications, Sybase IQ provides a set of views owned by the special user `dbo`, which correspond to the Adaptive Server Enterprise system tables and views. Where architectural differences make the contents of a particular Adaptive Server Enterprise table or view meaningless in a Sybase IQ context, the view is empty, containing only the column names and data types.

These topics list the Adaptive Server Enterprise system tables and their implementation in the Sybase IQ system catalog. The owner of all tables is `dbo` in each DBMS.

See also

- *ASE T-SQL Compatibility Views* on page 578

Tables in Each Adaptive Server Enterprise Database

Not all Adaptive Server Enterprise system tables are implemented in the Sybase IQ system catalog.

Table 212. Tables in each ASE database

Table name	Description	Data?	Supported by IQ?
<code>sysalternates</code>	One row for each user mapped to a database user	No	No
<code>syscolumns</code>	One row for each column in a table or view, and for each parameter in a procedure. In Sybase IQ, use the owner name <code>dbo</code> when querying, i.e. <code>dbo.syscolumns</code> .	Yes	Yes
<code>syscomments</code>	One or more rows for each view, rule, default, and procedure, giving SQL definition statement.	Yes	Yes
<code>sysconstraints</code>	One row for each referential and check constraint associated with a table or column.	No	No
<code>sysdepends</code>	One row for each procedure, view, or table that is referenced by a procedure, view.	No	No

Table name	Description	Data?	Supported by IQ?
sysindexes	One row for each clustered or nonclustered index, and one row for each table with no indexes, and an additional row for each table containing text or image data. In Sybase IQ, use the owner name dbo when querying, i.e. dbo.sysindexes.	Yes	Yes
sysiqobjects	One row for each system table, user table, view, procedure, trigger, event, join index, constraint, domain (sysdomain), domain (sysusertype), column, and index.	Yes	Yes
sysiqvindex	One row for each non-fp iq index.	Yes	Yes
syskeys	One row for each primary, foreign, or common key; set by user (not maintained by Adaptive Server Enterprise).	No	No
syslogs	Transaction log.	No	No
sysobjects	One row for each table, view, procedure, rule, default, log, and (in tempdb only) temporary object.	Contains compatible data only	Yes
sysprocedures	One row for each view, rule, default, and procedure, giving internal definition.	No	No
sysprotects	User permissions information.	No	No
sysreferences	One row for each referential integrity constraint declared on a table or column.	No	No
sysroles	Maps server-wide roles to local database groups.	No	No
syssegments	One row for each segment (named collection of disk pieces).	No	No
systhresholds	One row for each threshold defined for the database.	No	No
systypes	One row for each system-supplied and user-defined data type.	Yes	Yes
sysusers	One row for each user allowed in the database.	Yes	Yes

See also

- *SYSCOLUMNS ASE Compatibility View* on page 586
- *SYSCOMMENTS ASE Compatibility View* on page 587

- *SYSINDEXES ASE Compatibility View* on page 607
- *SYSIQOBJECTS ASE Compatibility View* on page 619
- *SYSIQVINDEXT ASE Compatibility View* on page 623
- *SYSOBJECTS ASE Compatibility View* on page 629
- *SYSTYPES ASE Compatibility View* on page 670
- *SYSUSERS ASE Compatibility View* on page 676

Tables in the Adaptive Server Enterprise Master Database

Not all Adaptive Server Enterprise master database tables are implemented in the Sybase IQ system catalog.

Table 213. ASE master database tables

Table name	Description	Data?	Supported by IQ?
syscharsets	One row for each character set or sort order	No	No
sysconfigures	One row for each configuration parameter that can be set by a user	No	No
syscurconfigs	Information about configuration parameters currently being used by the server	No	No
sysdatabases	One row for each database on the server	No	No
sysdevices	One row for each tape dump device, disk dump device, disk for databases, and disk partition for databases	No	No
sysengines	One row for each server currently online	No	No
syslanguages	One row for each language (except U.S. English) known to the server	No	No
syslocks	Information about active locks	No	No
sysloginroles	One row for each server login that possesses a system-defined role	No	No
syslogins	One row for each valid user account	Yes	Yes
sysmessages	One row for each system error or warning	No	No
sysprocesses	Information about server processes	No	No
sysremotelogins	One row for each remote user	No	No
sysrvroles	One row for each server-wide role	No	No
syssservers	One row for each remote server	No	No

System Tables and Views

Table name	Description	Data?	Supported by IQ?
sysusages	One row for each disk piece allocated to a database	No	No

Tables in the Adaptive Server Enterprise Sybsecurity Database

No Adaptive Server Enterprise sybsecurity database tables are implemented in the Sybase IQ system catalog.

Table 214. ASE sybsecurity database tables

Table name	Description	Data?	Supported by IQ?
sysaudits	One row for each audit record	No	No
sysauditoptions	One row for each global audit option	No	No

Compatibility with Other Sybase Databases

Use the topics in this section to simplify migration to Sybase IQ from other Sybase databases, and to serve as a guide for creating Sybase IQ applications that are compatible with Adaptive Server Enterprise or SQL Anywhere.

Compatibility features are addressed in each new version of Sybase IQ. This appendix compares Sybase IQ with Adaptive Server Enterprise, and SQL Anywhere.

About SQL Anywhere

Sybase IQ is an extension of SQL Anywhere.

In most cases, SQL syntax, functions, options, utilities, procedures, and other features are common to both products. There are, however, important differences. Do not assume that features described in SQL Anywhere documentation are supported for Sybase IQ.

The Sybase IQ documentation set calls many, but not all differences. *Sybase IQ documentation always supersedes the SQL Anywhere documentation.*

An Overview of Transact-SQL Support

Sybase IQ, like SQL Anywhere, supports a large subset of *Transact-SQL*, which is the dialect of SQL supported by Sybase Adaptive Server Enterprise.

The goal of Transact-SQL support in Sybase IQ is to provide application portability. Many applications, stored procedures, and batch files can be written for use with both Adaptive Server Enterprise and Sybase IQ databases.

The aim is to write applications to work with both Adaptive Server Enterprise and Sybase IQ. Existing Adaptive Server Enterprise applications generally require some changes to run on SQL Anywhere or Sybase IQ databases.

Transact-SQL support in Sybase IQ takes the following form:

- Most SQL statements are compatible between Sybase IQ and Adaptive Server Enterprise.
- For some statements, particularly in the procedure language used in procedures and batches, a separate Transact-SQL statement is supported along with the syntax supported in earlier versions of Sybase IQ. For these statements, SQL Anywhere and Sybase IQ support two dialects of SQL. In this appendix, we name those dialects Transact-SQL and Watcom-SQL.

Compatibility with Other Sybase Databases

- A procedure or batch is executed in either the Transact-SQL or Watcom-SQL dialect. You must use control statements from one dialect only throughout the batch or procedure. For example, each dialect has different flow control statements.

Sybase IQ supports a high percentage of Transact-SQL language elements, functions, and statements for working with existing data.

Further, Sybase IQ supports a very high percentage of the Transact-SQL stored procedure language (**CREATE PROCEDURE** syntax, control statements, and so on), and many, but not all, aspects of Transact-SQL data definition language statements.

There are design differences in the architectural and configuration facilities supported by each product. Device management, user management, and maintenance tasks such as backups tend to be system-specific. Even here, however, Sybase IQ provides Transact-SQL system tables as views, where the tables that are not meaningful in Sybase IQ have no rows. Also, Sybase IQ provides a set of system procedures for some of the more common administrative tasks.

Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ Architectures

Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ are complementary products, with architectures designed to suit their distinct purposes.

Sybase IQ is a high-performance decision-support server designed specifically for data warehousing and analytic processing. SQL Anywhere works well as a workgroup or departmental server requiring little administration, and as a personal database. Adaptive Server Enterprise works well as an enterprise-level server for large databases, with a focus on transaction processing.

This section describes architectural differences among the three products. It also describes the Adaptive Server Enterprise-like tools that Sybase IQ and SQL Anywhere include for compatible database management.

Servers and Databases

The relationship between servers and databases is different in Adaptive Server Enterprise from Sybase IQ and SQL Anywhere.

In Adaptive Server Enterprise, each database exists inside a server, and each server can contain several databases. Users can have login rights to the server, and can connect to the server. They can then connect to any of the databases on that server, provided that they have permissions. System-wide system tables, held in a master database, contain information common to all databases on the server.

In Sybase IQ, there is nothing equivalent to the Adaptive Server Enterprise master database. Instead, each database is an independent entity, containing all of its system tables. Users can have connection rights to a database, rather than to the server. When a user connects, he or she

connects to an individual database. There is no system-wide set of system tables maintained at a master database level. Each Sybase IQ database server can dynamically start and stop a database, to which users can maintain independent connections. Sybase strongly recommends that you run only one Sybase IQ database per server.

SQL Anywhere and Sybase IQ provide tools in their Transact-SQL support and Open Server support to allow some tasks to be carried out in a manner similar to Adaptive Server Enterprise. There are differences, however, in exactly how these tools are implemented.

Space Allocation and Device Management

Adaptive Server Enterprise, SQL Anywhere and Sybase IQ use different models for managing devices and allocating disk space initially and later, reflecting the different uses for the products.

For example:

- In Adaptive Server Enterprise, you allocate space in database devices initially using **DISK INIT** and then create a database on one or more database devices. You can add more space using **ALTER DATABASE** or automatically, using thresholds.
- In Sybase IQ, you initially allocate space by listing raw devices in the **CREATE DATABASE** statement. You can add more space manually using **CREATE DBSPACE**. Although you cannot add space automatically, you can create events to warn the DBA before space is actually needed. Sybase IQ can also use file system space. Sybase IQ does not support Transact-SQL **DISK** statements, such as **DISK INIT**, **DISK MIRROR**, **DISK REFIT**, **DISK REINIT**, **DISK REMIRROR**, and **DISK UNMIRROR**.
- SQL Anywhere is similar to Sybase IQ, except that the initial **CREATE DATABASE** statement takes a single file system file instead of a list of raw devices. SQL Anywhere lets you initialize its databases using a command utility named **dbinit**. Sybase IQ provides an expanded version of this utility called **iqinit** for initializing IQ databases.

For information on disk management, see the *System Administration Guide: Volume 1*.

System Tables, Catalog Store, and IQ Store

An IQ database is a joint data store.

The joint store consists of:

- The catalog store includes system tables and stored procedures, and resides in a set of tables that are compatible with SQL Anywhere.
- The permanent IQ store is the set of Sybase IQ tables. Table data is stored in indexes.
- The temporary store consists of a set of temporary tables which the database server uses for sorting and other temporary processing.

Catalog distinctions and compatibility features include:

Compatibility with Other Sybase Databases

- SQL Anywhere and Sybase IQ use a different schema from Adaptive Server Enterprise for the catalog (tables, columns, and so on).
- SQL Anywhere and Sybase IQ provide compatibility views that mimic relevant parts of the Adaptive Server Enterprise system tables, although there are performance implications when using them.
- In Adaptive Server Enterprise, the database owner (user ID `dbo`) owns the catalog objects.
- In SQL Anywhere and Sybase IQ, the system owner (user ID `SYS`) owns the catalog objects.

Note: A `dbo` user ID owns the Adaptive Server Enterprise-compatible system views provided by Sybase IQ.

Administrative Roles

Adaptive Server Enterprise has a more elaborate set of administrative roles than either SQL Anywhere or Sybase IQ.

In Adaptive Server Enterprise, there is a set of distinct roles, although more than one login account on an Adaptive Server Enterprise can be granted any role, and one account can possess more than one role.

Adaptive Server Enterprise Role	Description
System administrator	Responsible for general administrative tasks unrelated to specific applications; can access any database object.
System security officer	Responsible for security-sensitive tasks in Adaptive Server Enterprise, but has no special permissions on database objects.
Database owner	Has full permissions on objects inside the database he or she owns, can add users to a database and grant other users the permission to create objects and execute commands within the database.
Data definition statements	Permissions can be granted to users for specific data definition statements, such as CREATE TABLE or CREATE VIEW , enabling the user to create database objects.
Object owner	Each database object has an owner who may grant permissions to other users to access the object. The owner of an object automatically has all permissions on the object.

SQL Anywhere/Sybase IQ Role	Description
Database Administrator (DBA authority)	Has, like the Adaptive Server Enterprise Database Owner, full permissions on all objects inside the database (other than objects owned by SYS) and can grant other users the permission to create objects and execute commands within the database. The default database administrator is user ID DBA.
RESOURCE permission	Allows a user to create any kind of object within a database. This is in place of the Adaptive Server Enterprise scheme of granting permissions on individual CREATE statements.
Object owner	Sybase IQ has object owners in the same way Adaptive Server Enterprise does. The owner of an object automatically has all permissions on the object, including the right to grant permissions.

For seamless access to data held in both Adaptive Server Enterprise and Sybase IQ, create user IDs with appropriate permissions in the database (RESOURCE in Sybase IQ, or permission on individual **CREATE** statements in Adaptive Server Enterprise) and create objects from that user ID. If you use the same user ID in each environment, object names and qualifiers can be identical in the two databases, providing compatible access.

Data Types

Adaptive Server Enterprise, SQL Anywhere and Sybase IQ treat data types differently.

This section discusses compatibility information for data types.

Note: Data types that are not included in this section are currently supported by all three products.

Bit Data Type

Adaptive Server Enterprise, SQL Anywhere and Sybase IQ support the BIT data type, with differences.

The differences are:

- SQL Anywhere permits only 0 or 1.
- Adaptive Server Enterprise and Sybase IQ implicitly convert integral data types to BIT. Nonzero values are stored as 1 (TRUE).

Character Data Types

Adaptive Server Enterprise, SQL Anywhere and Sybase IQ permit CHAR and VARCHAR data, but each product treats these types differently.

- SQL Anywhere treats all strings as VARCHAR, even in a blank-padded database.
- Adaptive Server Enterprise and Sybase IQ differentiate between CHAR (fixed-length) and VARCHAR (variable-length) data.

Adaptive Server Enterprise trims trailing blank spaces from VARCHAR values. Sybase IQ trims trailing blanks from VARCHAR values depending on the form of the data and the operation.

When inserting into CHAR or VARCHAR:

- SQL Anywhere permits inserting integral data types into CHAR or VARCHAR (implicit conversion).
- Adaptive Server Enterprise and Sybase IQ require explicit conversion.

The maximum size of a column is determined as follows:

- Adaptive Server Enterprise CHAR and VARCHAR depend on the logical page size, which can be 2K, 4K, 8K, and 16K. For example:
 - 2K page size allows a column as large as a single row, about 1962 bytes.
 - 4K page size allows a column as large as about 4010 bytes.
- SQL Anywhere supports up to 32K-1 with CHAR and VARCHAR, and up to 2GB with LONG VARCHAR.
- SQL Anywhere supports the name LONG VARCHAR and its synonym TEXT, while Adaptive Server Enterprise supports only the name TEXT, not the name LONG VARCHAR.
- Sybase IQ supports CHAR and VARCHAR up to 32K-1 bytes.

Sybase IQ also supports up to 512TB (with an IQ page size of 128KB) and 2PB (with an IQ page size of 512KB) with LONG VARCHAR. For information on the LONG VARCHAR data type in Sybase IQ, see *Unstructured Data Analytics in Sybase IQ*.
- Adaptive Server Enterprise supports NCHAR, NVARCHAR, UNICHAR, UNIVARCHAR data types. N is for multibyte character sets; UNI is for single-byte character sets.
- SQL Anywhere and Sybase IQ support Unicode in the CHAR and VARCHAR data types, rather than as a separate data type.
- For compatibility between Sybase IQ and Adaptive Server Enterprise, always specify a length for character data types.

See also

- *Character Data Types* on page 69

Binary Data Types

Binary data type support differs between Adaptive Server Enterprise, SQL Anywhere and Sybase IQ.

Table 215. Binary data type supported sizes

Data type	Adaptive Server Enterprise	SQL Anywhere	Sybase IQ
BINARY	< page size	32KB - 1	255
VARBINARY	< page size	32KB - 1	32KB - 1
LONG BINARY*	not supported	2GB - 1	512TB (IQ page size 128KB) 2PB (IQ page size 512KB)
IMAGE	2GB	2GB - 1	use LONG BINARY*

*For information on the LONG BINARY data type in Sybase IQ, see *Unstructured Data Analytics in Sybase IQ*. This feature requires a separate license.

Adaptive Server Enterprise and SQL Anywhere display binary data differently when projected:

- Sybase IQ supports both Adaptive Server Enterprise and SQL Anywhere display formats.
- If '123' is entered in a BINARY field the SQL Anywhere display format is by bytes, as '123'; the Adaptive Server Enterprise display format is by nibbles, as '0x616263'.

See also

- *Binary Data Types* on page 78
- *NEWID Function [Miscellaneous]* on page 239
- *STRTOUUID Function [String]* on page 311
- *UUIDTOSTR Function [String]* on page 334
- *Character Data Types* on page 69

Date, Time, Datetime, and Timestamp Data Types

Although Adaptive Server Enterprise, SQL Anywhere and Sybase IQ all support some form of date and time data, there are some differences.

- SQL Anywhere and Sybase IQ support the 4-byte date and time data types.
- Adaptive Server Enterprise supports an 8-byte datetime type, and timestamp as a user-defined data type (domain) implemented as binary (8).
- SQL Anywhere and Sybase IQ support an 8-byte timestamp type, and an 8-byte datetime domain implemented as timestamp. The millisecond precision of the SQL

Compatibility with Other Sybase Databases

Anywhere/Sybase IQ `datetime` data type differs from that of Adaptive Server Enterprise.

Display formats for dates have different defaults:

- Adaptive Server Enterprise defaults to displaying dates in the format “MMM-DD-YYYY” but can be changed by setting an option.
- SQL Anywhere and Sybase IQ default to the ISO “YYYY-MM-DD” format but can be changed by setting an option.

Time conversions are as follows:

- Adaptive Server Enterprise varies the way it converts time stored in a string to an internal time, depending on whether the fraction part of the second was delimited by a colon or a period.
- SQL Anywhere and Sybase IQ convert times in the same way, regardless of the delimiter.

When you insert a time into a `DATETIME` column:

- Adaptive Server Enterprise and Sybase IQ default to supplying 1st January 1900.
- SQL Anywhere defaults to supplying the current date.

`TIME` and `DATETIME` values retrieved from an Adaptive Server Enterprise database change when inserted into a Sybase IQ table with a `DATETIME` column using **`INSERT...LOCATION`**. The **`INSERT...LOCATION`** statement uses Open Client, which has a `DATETIME` precision of 1/300 of a second.

For example, assume that the following value is stored in a table column in an Adaptive Server Enterprise database:

```
2004-11-08 10:37:22.823
```

When you retrieve and store it in a Sybase IQ table using **`INSERT...LOCATION`**, the value becomes:

```
2004-11-08 10:37:22.823333
```

Compatibility of Datetime and Time Values from ASE

A `DATETIME` or `TIME` value retrieved from an Adaptive Server Enterprise database using **`INSERT...LOCATION`** can have a different value due to the datetime precision of Open Client.

For example, the `DATETIME` value in the Adaptive Server Enterprise database is ‘2004-11-08 10:37:22.823’ as retrieved using **`INSERT...LOCATION`** is ‘2004-11-08 10:37:22.823333’.

BIGTIME and BIGDATETIME Support

Sybase IQ supports the Adaptive Server Enterprise (ASE) data types `BIGTIME` and `BIGDATETIME` for Component Integration Services (CIS) and **`INSERT...LOCATION`**.

- Component Integration Services with Adaptive Server Enterprise—aseodbc server class proxy tables mapped to Adaptive Server Enterprise tables that contain columns of data type `BIGTIME` and `BIGDATETIME`.

When you create a proxy table mapped to an Adaptive Server Enterprise table, a `BIGDATETIME` column is mapped to a `TIMESTAMP` column by default, if no mapping is specified. A `BIGTIME` column is mapped to a `TIME` column by default.

The asejdbc server class does not support the `BIGTIME` and `BIGDATETIME` data types.

- **INSERT...LOCATION**—the **INSERT...LOCATION** command to load data into Sybase IQ tables from Adaptive Server Enterprise tables that contain columns of data type `BIGTIME` and `BIGDATETIME`.

Sybase IQ inserts the Adaptive Server Enterprise data type `BIGTIME` into the Sybase IQ data type `TIME`.

Sybase IQ inserts the Adaptive Server Enterprise data type `BIGDATETIME` into the Sybase IQ data types `DATETIME`, `DATE`, `TIME`, and `TIMESTAMP`.

Numeric Data Types

Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ have different default precision and scale.

- In Adaptive Server Enterprise, the default is precision 18 scale 0.
- In SQL Anywhere, the default is precision 30 scale 6.
- In Sybase IQ, the default is precision 126 scale 38. Because these defaults are too large for TDS and for some client tools, always specify a precision and scale for Sybase IQ exact numeric types.

Text Data Type

Support for `TEXT` data differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Adaptive Server Enterprise supports up to 2GB with `LONG VARBINARY` (`LONG BINARY` in SQL Anywhere) and `TEXT`. SQL Anywhere does not support `LONG VARBINARY` as a column type, but uses `LONG BINARY` for the same purpose. SQL Anywhere supports up to 2GB with `LONG BINARY` and `TEXT`.
- Sybase IQ supports up to 32KB - 1 with `VARCHAR`. Sybase IQ also supports up to 512TB (with an IQ page size of 128KB) and 2PB (with an IQ page size of 512KB) with `LONG VARCHAR`. For information on the `LONG VARCHAR` data type in Sybase IQ, see *Unstructured Data Analytics in Sybase IQ*.

Image Data Type

Support for `IMAGE` data differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Adaptive Server Enterprise and SQL Anywhere support up to 2GB with `IMAGE`.
- Sybase IQ supports up to 512TB (with an IQ page size of 128KB) and 2PB (with an IQ page size of 512KB) with `LONG BINARY`. For information on the `LONG BINARY` data type in Sybase IQ, see *Unstructured Data Analytics in Sybase IQ*.

Java Data Types

Adaptive Server Enterprise allows Java data types in the database. SQL Anywhere and Sybase IQ do not.

Data Definition Language

Differences exist between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ in how you create databases and database objects.

Creating a Transact-SQL Compatible Database from Sybase Central

Create a Transact-SQL compatible database with case sensitivity settings that emulate Adaptive Server Enterprise.

1. In the Create Database wizard, access the **Specify a Collation Sequence** page.
2. Set **Case Sensitivity for String comparisons** to ON. This emulates Adaptive Server Enterprise.
3. Set **Case Sensitivity for Passwords** to ON.

These case sensitivity settings emulate Adaptive Server Enterprise.

Creating a Transact-SQL Compatible Database Using the CREATE DATABASE statement

Use Interactive SQL to create a Transact-SQL compatible database. Type the following statement, for example, in Interactive SQL:

```
CREATE DATABASE 'db-name.db' CASE RESPECT BLANK PADDING ON
```

Case-sensitivity

Case-sensitivity in databases refers to the case-sensitivity of data, identifiers, and passwords.

Case-sensitivity of Data

The case-sensitivity of the data is reflected in indexes, in the results of queries, and so on.

You decide the case-sensitivity of Sybase IQ data in comparisons when you create the database. By default, Sybase IQ databases are case-sensitive in comparisons, although data is always held in the case in which you enter it.

Adaptive Server Enterprise sensitivity to case depends on the sort order installed on the Adaptive Server Enterprise system. You can change case-sensitivity for single-byte character sets by reconfiguring the Adaptive Server Enterprise sort order.

Case-sensitivity of Identifiers

Identifiers include table names, column names, user IDs, and so on.

Sybase IQ does not support case-sensitive identifiers. In Adaptive Server Enterprise, the case-sensitivity of identifiers follows the case-sensitivity of the data.

In Adaptive Server Enterprise, user-defined data type names are case-sensitive. In Sybase IQ, they are case-insensitive.

Case-sensitivity of User IDs and Passwords

Case-sensitivity of passwords is treated differently from other identifiers.

In Sybase IQ and SQL Anywhere, all passwords in newly-created databases are case-sensitive, regardless of the case-sensitivity of the database. The default user ID is DBA and the password for this user is lowercase *sql*.

When you rebuild an existing database, Sybase IQ and SQL Anywhere determine the case-sensitivity of the password as follows:

- If the database was originally entered in a case-insensitive database, the password remains case-insensitive.
- If the password was originally entered in a case-sensitive database, uppercase and mixed-case passwords remain case-sensitive. If the password was entered in all lowercase, then the password becomes case-insensitive.
- Changes to both existing passwords and new passwords are case-sensitive.

In Adaptive Server Enterprise, the case-sensitivity of user IDs and passwords follows the case-sensitivity of the server.

Ensuring Compatible Object Names

Each database object must have a unique name within a certain name space.

Outside this name space, duplicate names are allowed. Some database objects occupy different name spaces in Adaptive Server Enterprise as compared to SQL Anywhere and Sybase IQ.

Table Name Uniqueness

Table name uniqueness requirements apply within a database:

- For Sybase IQ and SQL Anywhere, table names must be unique within a database for a given owner. For example, both `user1` and `user2` can create a table called `employee`; uniqueness is provided by the fully qualified names, `user1.employee` and `user2.employee`.
- For Adaptive Server Enterprise, table names must be unique within the database and to the owner.

Index Name Uniqueness

Index name uniqueness requirements apply within a table. In all three products, indexes are owned by the owner of the table on which they are created. Index names must be unique on a given table, but any two tables can have an index of the same name, even for the same owner. For example, in all three products, tables `t1` and `t2` can have indexes of the same name, whether they are owned by the same or different users.

Renaming Indexes and Foreign Keys

Sybase IQ allows you to rename explicitly created indexes, foreign key role names of indexes, and foreign keys, using the **ALTER INDEX** statement. SQL Anywhere allows you to rename indexes, foreign key role names, and foreign keys, using the **ALTER INDEX** statement. Adaptive Server Enterprise does not allow you to rename these objects.

Considerations When Using the CREATE TABLE Statement

When creating tables for compatibility, be aware of the following compatibility considerations for NULL treatment, check constraints, referential integrity, default values, identify columns, computed columns, temporary tables, and table location.

NULL in Columns

For compatible treatment of NULL:

- SQL Anywhere and Sybase IQ assume that columns can be null unless **NOT NULL** is stated in the column definition. You can change this behavior by setting the database option **ALLOW_NULLS_BY_DEFAULT** to the Transact-SQL compatible setting of **OFF**.
- SQL Anywhere assumes that **BIT** columns only cannot be NULL.
- Adaptive Server Enterprise assumes that columns cannot be null unless **NULL** is stated.

Check Constraints

Sybase IQ enforces check constraints on base, global temporary, and local temporary tables, and on user-defined data types. Users can log check integrity constraint violations and specify the number of violations that can occur before a **LOAD** statement rolls back.

Sybase IQ does not allow the creation of a check constraint that it cannot evaluate, such as those composed of user-defined functions, proxy tables, or non-Sybase IQ tables. Constraints

that cannot be evaluated are detected the first time the table on which the check constraint is defined is used in a **LOAD**, **INSERT**, or **UPDATE** statement. Sybase IQ does not allow check constraints containing:

- Subqueries
- Expressions specifying a host language parameter, a SQL parameter, or a column as the target for a data value
- Set functions
- Invocations of nondeterministic functions or functions that modify data

Adaptive Server Enterprise and SQL Anywhere enforce **CHECK** constraints. SQL Anywhere allows subqueries in check constraints.

Sybase IQ supports user-defined data types that allow constraints to be encapsulated in the data type definition.

Referential Integrity Constraints

Sybase IQ enforces referential integrity as described in the *System Administration Guide: Volume 1*.

Actions for enforcing integrity are supported as follows:

- SQL Anywhere supports all ANSI actions: SET NULL, CASCADE, DEFAULT, RESTRICT.
- Adaptive Server Enterprise supports two of these actions: SET NULL, DEFAULT.

Note: You can achieve CASCADE in Adaptive Server Enterprise by using triggers instead of referential integrity.

- Sybase IQ supports the RESTRICT action only.
- Sybase IQ does not support NOT NULL FOREIGN KEY.
- Sybase IQ has the restriction that a column cannot be both a candidate key and a foreign key at the same time.

Default Values in a Column

Default value support differs as follows:

- Adaptive Server Enterprise and SQL Anywhere support specifying a default value for a column.
- Only SQL Anywhere supports DEFAULT UTC TIMESTAMP.
- Sybase IQ supports specifying a default value for a column, except for the special values DEFAULT UTC TIMESTAMP and DEFAULT CURRENT UTC TIMESTAMP. Sybase IQ also ignores settings for the DEFAULT_TIMESTAMP_INCREMENT database option.

Identity Columns

Identity column support differs as follows:

Compatibility with Other Sybase Databases

- Sybase IQ supports `IDENTITY` or `DEFAULT AUTOINCREMENT` as a default value. Sybase IQ supports identity columns of any numeric type with any precision and scale 0, and the column can be `NULL`. Sybase IQ identity columns must be positive and are limited by the range of the data type. Sybase IQ supports a single identity column per table, and requires database option `IDENTITY_INSERT` set to a table name for explicit inserts and updates. To drop a table with an `IDENTITY` column, you cannot have `IDENTITY_INSERT` set to that table. The table can contain data when adding an identity column. Tables derived using `SELECT INTO` do not have Identity/Autoincrement columns. Sybase IQ views cannot contain `IDENTITY/DEFAULT AUTOINCREMENT` columns.
- SQL Anywhere supports the `AUTOINCREMENT` default value. SQL Anywhere supports identity columns of any numeric type with any allowable scale and precision. The identity column value can be positive, negative, or zero, limited by the range of the data type. SQL Anywhere supports any number of identity columns per table, and does not require `identity_insert` for explicit inserts, drops, and updates. The table must be empty when adding identity columns. SQL Anywhere identity columns can be altered to be nonidentity columns, and vice versa. You can add or drop `AUTOINCREMENT` columns from SQL Anywhere views.
- Adaptive Server Enterprise supports a single identity column per table. Adaptive Server Enterprise identity columns are restricted to only numeric data type scale 0, maximum precision 38. They must be positive, are limited by the range of the data type, and cannot be null. Adaptive Server Enterprise requires `identity_insert` for explicit inserts and drops, but not for updates to the identity column. The table can contain data when you add an identity column. Adaptive Server Enterprise users cannot explicitly set the next value chosen for an identity column. Adaptive Server Enterprise views cannot contain `IDENTITY/AUTOINCREMENT` columns. When using `SELECT INTO` under certain conditions, Adaptive Server Enterprise allows Identity/Autoincrement columns in the result table if they were in the table being selected from.

Computed Columns

Computed column support differs as follows:

- SQL Anywhere supports computed columns that can be indexed.
- Adaptive Server Enterprise and Sybase IQ do not.

Temporary Tables

You can create a temporary table by placing a pound sign (#) without an owner specification in front of the table name in a `CREATE TABLE` statement. These temporary tables are Sybase IQ-declared temporary tables and are available only in the current connection.

Locating Tables

Physical placement of a table is carried out differently in Adaptive Server Enterprise and Sybase IQ. Sybase IQ supports the `ON segment-name` clause, but `segment-name` refers to a Sybase IQ dbspace.

Considerations When Using the CREATE DEFAULT, CREATE RULE, and CREATE DOMAIN Statements

Sybase IQ provides an alternative means of incorporating rules.

- Adaptive Server Enterprise supports the Create Default and Create Rule statements to create named defaults.
- SQL Anywhere and Sybase IQ support the **CREATE DOMAIN** statement to achieve the same objective.

Considerations When Using the CREATE TRIGGER Statement

Support for triggers differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- SQL Anywhere supports both row-level and statement-level triggers.
- Adaptive Server Enterprise supports only statement-level triggers.
- Sybase IQ does not support triggers.

Note: A trigger is effectively a stored procedure that is run automatically either immediately before or immediately after an **INSERT**, **UPDATE**, or **DELETE** as part of the same transaction, that can be used to cause a dependent change (for example, to automatically update the name of an employee's manager when the employee is moved to a different department). It can also be used to write an audit trail to identify which modifications made which changes to the database, and at what time.

Considerations When Using the CREATE INDEX Statement

CREATE INDEX syntax differs slightly between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Adaptive Server Enterprise and SQL Anywhere support clustered or nonclustered indexes, using the following syntax:

```
CREATE [UNIQUE] [CLUSTERED] INDEX name
ON table (column,...)
ON dbspace
```

Adaptive Server Enterprise also allows the **NONCLUSTERED** keyword, but for both products the default is **NONCLUSTERED**.

- Adaptive Server Enterprise **CREATE INDEX** statements work in SQL Anywhere because SQL Anywhere allows, but ignores, the keywords **FILLFACTOR**, **IGNORE_DUP_KEY**, **SORTED_DATA**, **IGNORE_DUP_ROW**, and **ALLOW_DUP_ROW**.
- SQL Anywhere **CREATE INDEX** syntax supports the **VIRTUAL** keyword for use by its Index Consultant, but not for actual query executions.
- Sybase IQ supports seven specialized index types: **LF**, **HG**, **HNG**, **DATE**, **TIME**, **DTTM**, and **WD**. Sybase IQ also supports a **CMP** index on the relationship between two columns of

Compatibility with Other Sybase Databases

identical data type, precision, and scale. Sybase IQ defaults to creating an **HG** index unless the index type is specified in the **CREATE INDEX** statement:

```
CREATE [UNIQUE] [type] INDEX name
ON table (column,...)
```

Note: Sybase IQ also supports **CREATE JOIN INDEX**, which lets you create a prejoined index on a certain set of columns that are joined consistently and frequently in queries.

Users, Groups, and Permissions

There are some differences between the Adaptive Server Enterprise and SQL Anywhere and Sybase IQ models of users and groups.

In Adaptive Server Enterprise, users connect to a server, and each user requires a login ID and password to the server as well as a user ID for each database they want to access on that server.

SQL Anywhere and Sybase IQ users do not require a server login ID. All SQL Anywhere and Sybase IQ users receive a user ID and password for a database.

User Groups

All three products support user groups, so you can grant permissions to many users at one time. However, there are differences in the specifics of groups:

- Adaptive Server Enterprise allows each user to be a member of only one group.
- SQL Anywhere and Sybase IQ allow users to be members of multiple groups, and group hierarchies are allowed.

All three products have a `public` group, for defining default permissions. Every user automatically becomes a member of the `public` group.

Database Object Permissions

GRANT and **REVOKE** statements for granting permissions on individual database objects are very similar in all three products.

- All three products allow **SELECT**, **INSERT**, **DELETE**, **UPDATE**, and **REFERENCES** permissions on database tables and views, and **UPDATE** permissions on selected columns of database tables.

For example, the following statement is valid in all three products:

```
GRANT INSERT, DELETE
ON TITLES
TO MARY, SALES
```

This statement grants permission to use the **INSERT** and **DELETE** statements on the `TITLES` table to user `MARY` and to the `SALES` group.

- All three products allow **EXECUTE** permissions to be granted on stored procedures.
- Adaptive Server Enterprise also supports **GRANT** and **REVOKE** on additional items:

- Objects: columns within tables, columns within views, and stored procedures
- User abilities: **CREATE DATABASE**, **CREATE DEFAULT**, **CREATE PROCEDURE**, **CREATE RULE**, **CREATE TABLE**, **CREATE VIEW**
- SQL Anywhere and Sybase IQ require a user to have RESOURCE authority to create database objects. (A closely corresponding Adaptive Server Enterprise permission is GRANT ALL, used by a Database Owner.)
- All three products support the **WITH GRANT OPTION** clause, allowing the recipient of permissions to grant them in turn, although Sybase IQ does not permit **WITH GRANT OPTION** to be used on a **GRANT EXECUTE** statement.

Database-wide Permissions

Adaptive Server Enterprise uses a different model for database-wide user permissions.

- SQL Anywhere and Sybase IQ employ DBA permissions to allow a user full authority within a database.
- The System Administrator in Adaptive Server Enterprise enjoys this permission for all databases on a server. However, DBA authority on a Sybase IQ database is different from the permissions of an Adaptive Server Enterprise Database Owner, who must use the Adaptive Server Enterprise **SETUSER** statement to gain permissions on objects owned by other users.

Adding Users

Adaptive Server Enterprise requires a two-step process to add a user: **sp_addlogin** followed by **sp_add_user**.

SQL Anywhere and Sybase IQ add users in a single step.

Sybase IQ Login Management stored procedures, although not required to add or drop users, allow DBAs to add or drop Sybase IQ user accounts. When Sybase IQ User Administration is enabled, these Sybase IQ user accounts let DBAs control user connections and password expirations.

Although SQL Anywhere and Sybase IQ allow Adaptive Server Enterprise system procedures for managing users and groups, the exact syntax and function of these procedures differs in some cases.

See also

- *Adaptive Server Enterprise System Procedures* on page 568

Load Formats

Load format support differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Sybase IQ handles ASCII, BINARY, and BCP load formats.

Compatibility with Other Sybase Databases

- SQL Anywhere, in addition to ASCII and BINARY, also lets you import dBase, Excel, FoxPro, and Lotus file formats.
- Adaptive Server Enterprise handles ASCII and BINARY load formats through BCP.

Note: The syntax of the Sybase IQ and SQL Anywhere **LOAD** statement is based on BCP and designed to offer exactly the same functionality.

Options for Transact-SQL Compatibility

Set Sybase IQ database options using the **SET OPTION** statement.

See the Transact-SQL compatibility options in *Reference: Statements and Options*.

Data Manipulation Language

Query requirements differ between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

General Guidelines for Writing Portable SQL

Even if more than one server supports a given SQL statement, it might be a mistake to assume that default behavior is the same on each system.

General guidelines applicable to writing compatible SQL include:

- When writing SQL for use on more than one database management system, make your SQL statements as explicit as possible.
- Spell out all of the available options, rather than using default behavior.
- Use parentheses to make the order of execution within statements explicit, rather than assuming identical default order of precedence for operators.
- Use the Transact-SQL convention of an @ sign preceding variable names for Adaptive Server Enterprise portability.
- Declare variables and cursors in procedures and batches immediately following a **BEGIN** statement. Sybase IQ requires this, although Adaptive Server Enterprise allows declarations to be made anywhere in a procedure or batch.
- Do not use reserved words from either Adaptive Server Enterprise or Sybase IQ as identifiers in your databases.

Criteria for Writing Compatible Queries

There are two criteria for writing a query that runs on both Sybase IQ and Adaptive Server Enterprise databases.

- The data types, expressions, and search conditions in the query must be compatible.
- The syntax of the **SELECT** statement itself must be compatible.

Sybase IQ supports the following subset of the Transact-SQL **SELECT** statement.

Syntax

```

SELECT [ ALL | DISTINCT ] select-list
...[ INTO #temporary-table-name ]
...[ FROM table-spec,
...     table-spec, ... ]
...[ WHERE search-condition ]
...[ GROUP BY column-name, ... ]
...[ HAVING search-condition ]
...| [ ORDER BY expression [ ASC | DESC ], ... ] |
...| [ ORDER BY integer [ ASC | DESC ], ... ] |

```

Parameters

```

select-list:
{ table-name.* }...
{ * }...
{ expression }...
{ alias-name = expression }...
{ expression as identifier }...
{ expression as T_string }...

```

```

table-spec:
[ owner. ]table-name
... [ [ AS ] correlation-name ]
...

```

```

alias-name:
identifier | 'string' | "string"

```

The sections that follow provide details on several items to be aware of when writing compatible queries.

See also

- *Variables in Transact-SQL Procedures* on page 711

Subquery Support

Sybase IQ currently provides support for subqueries that is somewhat different from that provided by Adaptive Server Enterprise and SQL Anywhere.

Adaptive Server Enterprise and SQL Anywhere support subqueries in the **ON** clause; Sybase IQ does not currently support this.

UNION in subqueries is supported as follows:

- SQL Anywhere supports **UNION** in both correlated and uncorrelated subqueries.
- Sybase IQ supports **UNION** only in uncorrelated subqueries.
- Adaptive Server Enterprise does not support **UNION** in any subqueries.

SQL Anywhere supports subqueries in many additional places that a scalar value might appear in the grammar. Adaptive Server Enterprise and Sybase IQ follow the ANSI standard as to where subqueries can be specified.

GROUP BY Clause Support

GROUP BY ALL support differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Adaptive Server Enterprise supports **GROUP BY ALL**, which returns all possible groups including those eliminated by the **WHERE** clause and **HAVING** clause. These have the NULL value for all aggregates.
- SQL Anywhere does not support the **GROUP BY ALL** Transact-SQL extension.

ROLLUP and **CUBE** in the **GROUP BY** clause are supported as follows:

- Sybase IQ and SQL Anywhere support **ROLLUP** and **CUBE** in the **GROUP BY** clause.
- Adaptive Server Enterprise does not currently support **ROLLUP** and **CUBE**.

Adaptive Server Enterprise supports projecting nongrouped columns in the **SELECT** clause. This is known as extended group by semantics and returns a set of values. Sybase IQ supports and SQL Anywhere do not support extended group by semantics. Only SQL Anywhere supports the List() aggregate to return a list of values.

COMPUTE Clause Support

COMPUTE support differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Adaptive Server Enterprise supports the Transact-SQL **COMPUTE** clause.
- SQL Anywhere and Sybase IQ do not support the Transact-SQL **COMPUTE** clause since it is not in the ANSI standard and this functionality is provided by most third-party front-end tools.

WHERE Clause Support

The **WHERE** clause differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ in support for the Contains() predicate, and treatment of trailing white space in the **Like()** predicate.

- Sybase IQ supports the **Contains()** predicate for word searches in character data (similar to Contains in MS SQL Server and Verity). Sybase IQ uses **WORD** indexes and **TEXT** indexes to optimize these, if possible.
- Adaptive Server Enterprise does not support **Contains()**.

Transact-SQL Outer Joins Support

Supported syntax for outer joins differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Adaptive Server Enterprise fully supports *= and =* Transact-SQL syntax for outer joins.

- SQL Anywhere and Sybase IQ support Transact-SQL outer joins, but reject some complex Transact-SQL outer joins that are potentially ambiguous.
- Sybase IQ does not support chained (nested) Transact-SQL outer joins. Use ANSI syntax for this type of multiple outer join.

Note: Transact-SQL outer join syntax is deprecated in SQL Anywhere and Sybase IQ.

For detailed information on Transact-SQL outer joins, including ANSI syntax alternatives, see the white paper “Semantics and Compatibility of Transact-SQL Outer Joins,” from *MySybase*. Although written for SQL Anywhere, the information in the document also applies to Sybase IQ.

ANSI Joins Support

Support for ANSI join syntax differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Sybase IQ does not currently support subqueries in the **ON** clause.
- Adaptive Server Enterprise and SQL Anywhere support subqueries in the **ON** clause.
- A **CONTAINS** condition in the **FROM** clause in queries using ANSI join syntax is supported but may have suboptimal performance. Using Outer Joins for **CONTAINS** in the **FROM** clause should only be used if the “score” column from each of the **CONTAINS** clauses is required. Otherwise **CONTAINS** should be moved to an **ON** condition or **WHERE** clause.

Full outer join support is as follows:

- SQL Anywhere and Sybase IQ support **FULL OUTER JOIN**.
- Adaptive Server Enterprise does not support **FULL OUTER JOIN**.

Null Comparisons Support

Adaptive Server Enterprise has Transact-SQL extensions that permit predicates to compare the null value.

For example, `{col} = Null` means `{col} Is Null`.

SQL Anywhere and Sybase IQ use ANSI semantics for null comparisons unless the `ANSINULL` option is set to `OFF`, in which case such comparisons are Adaptive Server Enterprise-compatible.

Note: SQL Anywhere 8.0 and later adds support for the `TDS_EMPTY_STRING_AS_NULL` to offer Adaptive Server Enterprise compatibility in mapping empty strings to the null value.

Zero-length Strings Support

Zero-length strings are treated differently in Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Adaptive Server Enterprise treats zero-length strings as the null value.

Compatibility with Other Sybase Databases

Adaptive Server Enterprise users store a single space for blank strings.

- SQL Anywhere and Sybase IQ follow ANSI semantics for zero-length strings, that is, a zero-length string is a real value; it is not null.

HOLDLOCK, SHARED, and FOR BROWSE Support

HOLDLOCK, SHARED, and FOR BROWSE syntax differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

- Adaptive Server Enterprise supports HOLDLOCK, SHARED, and FOR BROWSE syntax.
- SQL Anywhere supports HOLDLOCK but does not support SHARED or FOR BROWSE.
- Sybase IQ does not support these keywords.

SQL Function Support

Sybase IQ supports most of the same functions as SQL Anywhere and Adaptive Server Enterprise, with some differences.

- Adaptive Server Enterprise supports the **USING CHARACTERS | USING BYTES** syntax in **PatIndex()**; SQL Anywhere and Sybase IQ do not.
- Adaptive Server Enterprise supports the **Reverse()** function; SQL Anywhere and Sybase IQ do not.
- Adaptive Server Enterprise supports **Len()** as alternative syntax for **Length()**; SQL Anywhere does not support this alternative.
- Adaptive Server Enterprise supports the **Square()** and **Str_Replace()** Microsoft compatibility functions; SQL Anywhere does not.
- Sybase IQ supports **Str_Replace()**.
- Adaptive Server Enterprise and SQL Anywhere support **TSEQUAL()** to compare two timestamps for modification time; Sybase IQ does not support **TSEQUAL()**. (**TSEQUAL** is not relevant in the Sybase IQ table-level versioning model.)
- Sybase IQ supports **ROWID()**; Adaptive Server Enterprise and SQL Anywhere do not.
- SQL Anywhere and Sybase IQ support **Cast()** in addition to Adaptive Server Enterprise's **Convert()** for data type conversions.

Note: **Cast()** is the ANSI-compliant name.

- SQL Anywhere and Sybase IQ support **Lcase()** and **Ucase()** as synonyms of **Lower()** and **Upper()**; Adaptive Server Enterprise does not.
- SQL Anywhere and Sybase IQ support the **Locate()** string function; Adaptive Server Enterprise does not.
- SQL Anywhere supports the **IsDate()** and **IsNumeric()** function to test the ability to convert a string to the respective data type; Adaptive Server Enterprise does not. Sybase IQ supports **IsDate()**. You can use **IsNumeric** in Sybase IQ, but CIS functional compensation performance considerations apply.

- SQL Anywhere supports the **NEWID**, **STRTOUID**, and **UIDTOSTR** functions; Adaptive Server Enterprise does not. These are native functions in Sybase IQ, so CIS functional compensation performance considerations do not apply.

Note: Some SQL functions, including **SOUNDEX** and **DIFFERENCE** string functions, and some date functions operate differently in Sybase IQ and SQL Anywhere. The Sybase IQ database option `ASE_FUNCTION_BEHAVIOR` specifies that output of some of the Sybase IQ data type conversion functions, including **HEXTOINT** and **INTTOHEX**, is consistent with the output of Adaptive Server Enterprise functions.

OLAP Function Support

Currently, Adaptive Server Enterprise does not support OLAP functions. Sybase IQ and SQL Anywhere do.

Sybase IQ currently supports these OLAP functions:

- **Corr()**
- **Covar_Pop()**
- **Covar_Samp()**
- **Cume_Dist**
- **Dense_Rank()**
- **Exp_Weighted_Avg**
- **First_Value**
- **Last_Value**
- **Median**
- **Ntile()**
- **Percent_Rank()**
- **Percentile_Cont()**
- **Percentile_Disc()**
- **Rank()**
- **Regr_Avgx()**
- **Regr_Avgy()**
- **Regr_Intercept()**
- **Regr_R2**
- **Regr_Slope()**
- **Regr_Sxx()**
- **Regr_Sxy()**
- **Regr_Syy()**
- **StdDev()**
- **Stddev_Pop**
- **Stddev_Samp**
- **Var_Pop**

Compatibility with Other Sybase Databases

- **Var_Samp**
- **Variance()**
- **Weighted_Avg**

SQL Anywhere supports all of the Sybase IQ OLAP functions.

Currently, Adaptive Server Enterprise does not support OLAP functions.

CIS functional compensation does not support OLAP functions.

Note: Support for OLAP functions is a rapidly evolving area of Sybase product development.

System Function Support

SQL Anywhere and SQL AnywhereSybase IQ do not support certain Adaptive Server Enterprise system functions.

These Adaptive Server Enterprise system functions are not supported by SQL Anywhere and Sybase IQ:

- **curunreservedpgs()** – number of pages free on a dbspace.
- **data_pgs()** – number of pages used by a table or index.
- **host_id()** – UNIX pid of the server process.
- **hos_name()** – name of the machine on which the server is running.
- **lct_admin()** – manages the “last chance threshold” for the Transaction manager.
- **reserved_pgs()** – number of pages allocated to a table or index.
- **rowcnt()** – number of rows in the specified table.
- **valid_name()** – whether a name would be a valid name if used, for example, for a table.
- **valid_user()** – returns TRUE if that user has connect permissions.
- **ptn_data_pgs()** – number of data pages in a partition.
- **index_colorder()** – returns the column order in an index.

User Defined Function Support

User-defined function (UDF) support differs between Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ:

- SQL Anywhere supports UDFs in SQL, Java, and C.
- Adaptive Server Enterprise supports UDFs written only in Java .
- Sybase IQ offers support for UDFs via CIS query decomposition, but there are performance implications.

Differences Interpreting Arithmetic Expressions on Dates

SQL Anywhere and Sybase IQ interpret arithmetic expressions on dates as shorthand notation for various date functions. Adaptive Server Enterprise does not.

- Date +/- integer is equivalent to **Dateadd()**.

- Date – date is equivalent to **Datediff()**.
- Date + time creates a timestamp from the two.

SELECT INTO Statement Support

There are differences in the types of tables permitted in SELECT INTO statements in Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

Consider this example statement:

```
select into table1 from table2
```

- Adaptive Server Enterprise permits *table1* to be permanent, temporary or a proxy table. Adaptive Server Enterprise also supports **SELECT INTO EXISTING TABLE**.
- SQL Anywhere and Sybase IQ permit *table1* to be a permanent or a temporary table. A permanent table is created only when you select into *table* and specify more than one column. **SELECT INTO #table**, without an owner specification, always creates a temporary table, regardless of the number of columns specified. **SELECT INTO** table with just one column selects into a host variable.

Updatable Views Support

Adaptive Server Enterprise and SQL Anywhere are more liberal than ANSI permits on the view definitions that are updatable when the WITH CHECK option is not requested.

SQL Anywhere offers the **ANSI_UPDATE_CONSTRAINTS** option to control whether updates are restricted to those supported by SQL92, or a more liberal set of rules.

Sybase IQ permits **UPDATE** only on single-table views that can be flattened. Sybase IQ does not support **WITH CHECK**.

Support for FROM Clause in UPDATE and DELETE

Sybase IQ Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ all support the FROM clause with multiple tables in UPDATE and DELETE.

Transact-SQL Procedure Language Overview

The stored procedure language is the part of SQL used in stored procedures and batches.

SQL Anywhere and Sybase IQ support a large part of the Transact-SQL stored procedure language in addition to the Watcom-SQL dialect based on SQL92.

Transact-SQL Stored Procedure Overview

Because it is based on the ISO/ANSI draft standard, the SQL Anywhere and Sybase IQ stored procedure language differs from the Transact-SQL dialect in many ways.

Many of the concepts and features are similar, but the syntax is different. SQL Anywhere and Sybase IQ support for Transact-SQL takes advantage of the similar concepts by providing automatic translation between dialects. However, you must write a procedure exclusively in one of the two dialects, not in a mixture of the two.

There are a variety of aspects to SQL Anywhere and Sybase IQ support for Transact-SQL stored procedures, including:

- Passing parameters
- Returning result sets
- Returning status information
- Providing default values for parameters
- Control statements
- Error handling

Transact-SQL Batch Overview

In Transact-SQL, a batch is a set of SQL statements submitted together and executed as a group, one after the other.

Batches can be stored in command files. The ISQL utility in SQL Anywhere and Sybase IQ and the `isql` utility in Adaptive Server Enterprise provide similar capabilities for executing batches interactively.

The control statements used in procedures can also be used in batches. SQL Anywhere and Sybase IQ support the use of control statements in batches and the Transact-SQL-like use of nondelimited groups of statements terminated with a **GO** statement to signify the end of a batch.

For batches stored in command files, SQL Anywhere and Sybase IQ support the use of parameters in command files. Adaptive Server Enterprise does not support parameters.

SQL Statements in Procedures and Batches

Some SQL statements supported by Sybase IQ are part of one dialect, but not the other.

You cannot mix the two dialects within a procedure or batch. This means that:

- You can include Transact-SQL-only statements with statements that are part of both dialects in a batch or procedure.
- You can include statements not supported by Adaptive Server Enterprise with statements that are supported by both servers in a batch or procedure.

- You cannot include Transact-SQL—only statements with Sybase IQ—only statements in a batch or procedure.

SQL statements not separated by semicolons are part of a Transact-SQL procedure or batch. See *Reference: Statements and Options* for details of individual statements.

Transact-SQL compatibility has improved; incorrect SQL syntax that was previously accepted now fails with an error.

Expression Subqueries in IF Statements

Adaptive Server Enterprise and SQL Anywhere support comparisons between a variable and a scalar value returned by an expression subquery.

For example:

```
create procedure testIf ()
begin
  declare var4 int;
set var4 = 10;
  if var4 = (select MIN (a_i1) from a) then set
    var4 = 100;
end if;
end;
```

CASE Statement Support

Permitted usage of the CASE statement differs in Sybase IQ and SQL Anywhere.

The **CASE** statement is not supported in Adaptive Server Enterprise, which supports case expressions only.

See also

- *Expressions* on page 22

Row-level Cursor Operations Support

Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ support the use of cursors with UPDATE and DELETE.

Consider this example:

```
UPDATE WHERE CURRENT OF {cursor}
```

```
DELETE WHERE CURRENT OF {cursor}
```

In Sybase IQ, updatable cursors are asensitive only, for one table only, and chained only. Updatable hold cursors are not permitted. Updatable cursors in Sybase IQ get a table lock.

PRINT Command Support

Support for PRINT differs in Adaptive Server Enterprise, SQL Anywhere, and Sybase IQ.

The effect of **PRINT** depends on the client:

Compatibility with Other Sybase Databases

- Adaptive Server Enterprise **PRINT** always sends a message to the client.
- In SQL Anywhere and Sybase IQ, **PRINT** sends a message to the client for Open Client and JDBC connections.
- Adaptive Server Enterprise stored procedures that rely on **PRINT** work in Sybase IQ using Interactive SQL.

Note: Sybase IQ users might prefer Interactive SQL with JDBC, rather than the iAdaptive Server Anywhere JDBC driver (formerly called the JDBC-ODBC bridge).

Automatic Translation of Stored Procedures

In addition to supporting Transact-SQL alternative syntax, SQL Anywhere and Sybase IQ provide aids for translating statements between the Watcom-SQL and Transact-SQL dialects.

Functions returning information about SQL statements and enabling automatic translation of SQL statements include:

Table 216. Functions enabling automatic translation

Function	Description
SQLDialect(statement)	Returns Watcom-SQL or Transact-SQL.
WatcomSQL(statement)	Returns the Watcom-SQL syntax for the statement.
TransactSQL(statement)	Returns the Transact-SQL syntax for the statement.

These are functions and thus can be accessed using a **SELECT** statement from ISQL. For example, the following statement returns the value Watcom-SQL:

```
SELECT SqlDialect('select * from Employees')
```

Result Sets from Transact-SQL Procedures

SQL Anywhere/Sybase IQ procedures and Transact-SQL procedures return result sets differently.

SQL Anywhere and Sybase IQ use a **RESULT** clause to specify returned result sets.

In Transact-SQL procedures, column names or alias names of the first query are returned to the calling environment.

The following Transact-SQL procedure illustrates how Transact-SQL stored procedures return result sets:

```
CREATE PROCEDURE showdept (@deptname varchar(30))
AS
```



```

SELECT Employees.Surname, Employees.GivenName
FROM Departments, Employees
WHERE Departments.DepartmentName = @deptname
AND Departments.DepartmentID = Employees.DepartmentID

```

The following is the corresponding SQL Anywhere or Sybase IQ procedure:

```

CREATE PROCEDURE showdept(in deptname varchar(30))
RESULT ( lastname char(20), firstname char(20))
BEGIN
    SELECT Employees.Surname, Employees.GivenName
    FROM Departments, Employees
    WHERE Departments.DepartmentName = deptname
    AND Departments.DepartmentID = Employee.DepartmentID
END

```

There are minor differences in the way the three Sybase client tools present multiple results to the client:

- **isql** displays all results in a single stream.
- Interactive SQL presents each result set on a separate tab. You must enable this functionality in the Option menu. Make it a permanent change, then restart or reconnect to Interactive SQL.
- Interactive SQL Classic (deprecated) provides **RESUME** to display each successive result set.

Variables in Transact-SQL Procedures

SQL Anywhere and Sybase IQ assign values to variables in procedures differently than Transact-SQL.

SQL Anywhere and Sybase IQ use the SET statement to assign values to variables in a procedure.

In Transact-SQL, values are assigned using the **SELECT** statement with an empty table list. The following simple procedure illustrates how the Transact-SQL syntax works:

```

CREATE PROCEDURE multiply
    @mult1 int,
    @mult2 int,
    @result int output
AS
SELECT @result = @mult1 * @mult2

```

This procedure can be called as follows:

```

CREATE VARIABLE @product int
go
EXECUTE multiply 5, 6, @product OUTPUT
go

```

The variable *@product* has a value of 30 after the procedure executes.

There are some differences in order and persistence of variable declarations:

- In Adaptive Server Enterprise, you can declare variables anywhere in the body of a stored procedure. Variables persist for the duration of the procedure.
- In SQL Anywhere and Sybase IQ, you must declare variables at the beginning of a compound statement (that is, immediately after **BEGIN** in a **BEGIN...END** pair). Variables persist only for the duration of the compound statement.

See also

- *Criteria for Writing Compatible Queries* on page 700

Error Handling in Transact-SQL Procedures

Default procedure error handling is different in the Watcom-SQL and Transact-SQL dialects.

By default, Watcom-SQL dialect procedures exit when they encounter an error, returning `SQLSTATE` and `SQLCODE` values to the calling environment.

You can build explicit error handling into Watcom-SQL stored procedures using the **EXCEPTION** statement, or you can instruct the procedure to continue execution at the next statement when it encounters an error, using the **ON EXCEPTION RESUME** statement.

When a Transact-SQL dialect procedure encounters an error, execution continues at the following statement. The global variable `@@error` holds the error status of the most recently executed statement. You can check this variable following a statement to force return from a procedure. For example, the following statement causes an exit if an error occurs:

```
IF @@error != 0 RETURN
```

When the procedure completes execution, a return value indicates the success or failure of the procedure. This return status is an integer, and can be accessed as follows:

```
DECLARE @status INT
EXECUTE @status = proc_sample
IF @status = 0
    PRINT 'procedure succeeded'
ELSE
    PRINT 'procedure failed'
```

This table describes the built-in procedure return values and their meanings:

Table 217. Built-in procedure return values

Value	Meaning
0	Procedure executed without error
-1	Missing object
-2	Data type error

Value	Meaning
-3	Process was chosen as deadlock victim
-4	Permission error
-5	Syntax error
-6	Miscellaneous user error
-7	Resource error, such as out of space
-8	Nonfatal internal problem
-9	System limit was reached
-10	Fatal internal inconsistency
-11	Fatal internal inconsistency
-12	Table or index is corrupt
-13	Database is corrupt
-14	Hardware error

The **RETURN** statement can be used to return other integers, with their own user-defined meanings.

Using the RAISERROR Statement in Procedures

The **RAISERROR** statement is a Transact-SQL statement for generating user-defined errors. It has a similar function to the **SIGNAL** statement.

By itself, **RAISERROR** does not cause an exit from the procedure, but it can be combined with a **RETURN** statement or a test of the **@@error** global variable to control execution following a user-defined error.

If you set the **ON_TSQL_ERROR** database option to **CONTINUE**, **RAISERROR** no longer signals an execution-ending error. Instead, the procedure completes and stores the **RAISERROR** status code and message, and returns the most recent **RAISERROR**. If the procedure causing the **RAISERROR** was called from another procedure, **RAISERROR** returns after the outermost calling procedure terminates.

You lose intermediate **RAISERROR** statuses and codes when the procedure terminates. If, at return time, an error occurs along with **RAISERROR**, the error information is returned and you lose the **RAISERROR** information. The application can query intermediate **RAISERROR** statuses by examining **@@error** global variable at different execution points.

Transact-SQL-like Error Handling in the Watcom-SQL Dialect

You can make a Watcom-SQL dialect procedure handle errors in a Transact-SQL-like manner.

Supply the **ON EXCEPTION RESUME** clause to the **CREATE PROCEDURE** statement:

```
CREATE PROCEDURE sample_proc()  
ON EXCEPTION RESUME  
BEGIN  
    ...  
END
```

The presence of an **ON EXCEPTION RESUME** clause prevents explicit exception handling code from being executed, so avoid using these two clauses together.

SQL Anywhere and Sybase IQ Differences and Shared Functionality

Sybase IQ and SQL Anywhere have differences in starting and managing databases and servers, database option support, DDL support, and DML support.

For additional information, always refer to the Sybase IQ documentation set when using the product. Refer to the SQL Anywhere documentation set when using SQL Anywhere, or when the Sybase IQ documentation refers to SQL Anywhere documentation for specific functionality only.

SQL Anywhere Server and Database Startup and Administration

Starting and managing databases and servers differs between Sybase IQ and SQL Anywhere.

- Sybase IQ uses the server startup command **start_iq**, instead of the SQL Anywhere network server startup command.
- Sybase IQ does not support personal servers.
- Sybase IQ supports many SQL Anywhere server command line options, but not all. Other server options are supported for Sybase IQ but not for SQL Anywhere.
- Sybase IQ provides the **stop_iq** utility (UNIX) to shut down servers.
- Clauses permitted in the **BACKUP** and **RESTORE** statements differ in Sybase IQ and SQL Anywhere.
- SQL Remote is supported in Sybase IQ only for multiplex operations.

Sybase IQ supports many SQL Anywhere database administration utilities, but not all:

- The following SQL Anywhere utilities are *not* supported by Sybase IQ: **backup**, **compression**, **console**, **initialization**, **license**, **log transfer**, **log translation**, **rebuild**, **spawn**, some **transaction log** options (**-g**, **-il**, **-ir**, **-n**, **-x**, **-z**), **uncompression**, **unload**, **upgrade**, and **write file**.
- Sybase IQ supports the SQL Anywhere **validation** utility only on the catalog store. To validate the IQ store, use **sp_iqcheckdb**.

SQL Anywhere Database Options

Some SQL Anywhere database options are not supported by Sybase IQ, including `DEFAULT_TIMESTAMP_INCREMENT`.

Some database options apply only to the catalog store, including:

`FOR_XML_NULL_TREATMENT`, `ISOLATION_LEVEL`, `PREFETCH`, `PRECISION`, and `SCALE`.

Options with differences in behavior, default, or allowed values include `DELAYED_COMMITS`, `TIME_FORMAT`, `TIMESTAMP_FORMAT`.

Sybase IQ also includes many options that SQL Anywhere does not support.

SQL Anywhere Data Definition Language (DDL) Differences

In addition to the DDL differences discussed previously:

- In a **DELETE/DROP** or **PRIMARY KEY** clause of an **ALTER TABLE** statement, Sybase IQ takes the **RESTRICT** action (reports an error if there are associated foreign keys). SQL Anywhere always takes the **CASCADE** action.
- Similarly, **DROP TABLE** statement reports an error in Sybase IQ if there are associated foreign-key constraints.
- Sybase IQ does not support these DDL statements: **CREATE COMPRESSED DATABASE**, **CREATE TRIGGER**, **SETUSER**.
- Sybase IQ supports referential integrity at the statement level, rather than the transaction-level integrity that SQL Anywhere supports with the **CHECK ON COMMIT** clause of the **CREATE TABLE** statement.
- A Sybase IQ table cannot have a foreign key that references a SQL Anywhere (or catalog) table, and a SQL Anywhere table cannot have a foreign key that references a Sybase IQ table.
- In a Sybase IQ database, publications can only be created on SQL Anywhere tables.
- In **CREATE DATABASE**, the defaults for case-sensitivity and collation differ. The defaults for Sybase IQ are `CASE RESPECT` and the `ISO_BINENG` collation; for SQL Anywhere, the defaults are `CASE IGNORE`, and collation inferred from the language and character set of the operating system.
- Sybase IQ does not support the **CREATE ENCRYPTED DATABASE** and **CREATE DECRYPTED DATABASE** commands supported by SQL Anywhere. See *Advanced Security in Sybase IQ*.

SQL Anywhere Data Manipulation Language (DML) Differences

Not all SQL Anywhere DML objects and syntax are supported by Sybase IQ.

- Sybase IQ does not support these DML and procedural statements: **EXPLAIN**, **GET DATA**, **INPUT**, **PREPARE TO COMMIT**, **PUT**, **READTEXT**, **ROLLBACK TRIGGER**, **SYSTEM**, **UNLOAD TABLE**, **VALIDATE TABLE**.

Note: A set of extraction options perform a role similar to **UNLOAD TABLE**; for details, see the data extraction options in *System Administration Guide: Volume 1*.

- Sybase IQ supports the **INSERT...LOCATION** syntax; SQL Anywhere does not.
- **LOAD TABLE** options differ in Sybase IQ and SQL Anywhere.
- **OPEN** statement in Sybase IQ does not support **BLOCK** and **ISOLATION LEVEL** clauses.
- Sybase IQ does not support triggers.
- Use of transactions, isolation levels, checkpoints, and automatically generated COMMITs, as well as cursor support, is different in Sybase IQ and SQL Anywhere.
- When you **SELECT** from a stored procedure in Sybase IQ, CIS functional compensation performance considerations apply.
- Sybase IQ ignores the database name qualifier in fully qualified names in Adaptive Server Enterprise **SELECT** statements, such as a **FROM** clause with `<database name>.<owner>.<table name>`. For example, Sybase IQ interprets the query `SELECT * FROM XXX..TEST` as `SELECT * FROM TEST`.

Adaptive Server Enterprise and Sybase IQ Differences and Shared Functionality

Sybase IQ and Adaptive Server Enterprise have differences in stored procedure support and views support.

For additional information, always refer to the Sybase IQ documentation set when using the product. Refer to the Adaptive Server Enterprise documentation set when using Adaptive Server Enterprise, or when the Sybase IQ documentation refers to Adaptive Server Enterprise documentation for specific functionality only.

Adaptive Server Enterprise Stored Procedures

Certain stored procedures are not supported by Sybase IQ.

Sybase IQ no longer supports these Adaptive Server Enterprise stored procedures:

- **sp_addserver**
- **sp_configure**
- **sp_estspace**
- **sp_help**
- **sp_helpuser**
- **sp_who**

Sybase IQ no longer supports the following catalog procedures:

- **sp_column_privileges**
- **sp_databases**
- **sp_datatype_info**

- `sp_server_info`

Adaptive Server Enterprise System Views

Certain views are not supported by Sybase IQ.

Sybase IQ no longer supports these Adaptive Server Enterprise views:

- `sysalternates`
- `sysaudits`
- `sysauditoptions`
- `sysconstraints`
- `syscharsets`
- `sysconfigures`
- `syscurconfigs`
- `sysdatabases`
- `sysdepends`
- `sysdevices`
- `sysengines`
- `syskeys`
- `syslanguages`
- `syslocks`
- `syslogs`
- `sysloginroles`
- `sysmessages`
- `sysprocedures`
- `sysprocesses`
- `sysprotects`
- `sysreferences`
- `sysremotelogins`
- `sysroles`
- `syssegments`
- `syssservers`
- `sysssrvroles`
- `systhresholds`
- `sysusages`

Column Name Differences

The column name used in the Adaptive Server Enterprise view SYSTYPES is “allownulls”.

The column name used in the Sybase IQ view SYSTYPES is “allowsnulls”.

Index

A

ABS function 123
 absolute value 123
 ACOS function 124
 Adaptive Server Enterprise
 compatibility 683, 690
 advice
 clearing 437
 displaying 437
 storing 437
 aggregate functions 99
 AVG 129
 COUNT 153
 LIST 224
 MAX 230
 MIN 232
 STDDEV 303
 STDDEV_POP 305
 STDDEV_SAMP 306
 SUM 315
 VAR_POP 335
 VAR_SAMP 336
 VARIANCE 338
 ALL
 conditions 38
 allocation map
 resetting 375
 alphabetic characters
 defined 20
 analytic functions
 DENSE_RANK 180
 NTILE 244
 PERCENT_RANK 252
 PERCENTILE_CONT 254
 PERCENTILE_DISC 256
 RANK 265
 analytical functions 101
 AND conditions 45
 ANY
 conditions 38
 apostrophe
 in strings 21
 arc-cosine 124
 arc-sine 126
 arc-tangent 127

arc-tangent ratio 128
 architectures
 Adaptive Server Enterprise 684
 SQL Anywhere 684
 ARGN function 125
 argument selection 125
 arithmetic expressions 25
 on dates 706
 ASCII function 125
 ASCII value 125, 136
 ASIN function 126
 ATAN function 127
 ATAN2 function 128
 automatic joins
 and foreign keys 349
 average 129
 AVG function 129

B

backslashes
 not allowed in SQL identifiers 20
 backup history file
 location 10
 backup operations
 summary 366
 backups
 during checkpoint 524
 during low activity 524
 batches
 Transact-SQL overview 708
 writing 708
 BETWEEN conditions 39
 BFILE function
 description 130
 BIGDATETIME data type
 compatibility 690
 BIGINTTOHEX function 130
 BIGTIME data type
 compatibility 690
 binary data
 compatibility 689
 BINARY data type 79
 BIT data type
 compatibility 687
 Transact-SQL 83

Index

- bit length 131
- BIT_LENGTH function 131
- bitwise operators 26
- blanks
 - trimming trailing 71, 688
- BLOB data type
 - LIKE conditions 39
- brackets
 - database objects 20
 - SQL identifiers 20
- breaking
 - terms 359, 360
- buffer cache
 - monitoring with sp_iqsysmon 494
- byte length 250
- BYTE_LENGTH function 132
- BYTE_LENGTH64 function
 - description 133
- BYTE_SUBSTR function
 - description 133
- BYTE_SUBSTR64 function
 - description 133
- C**
- CASE expression 28
 - NULLIF function 246
- case sensitivity
 - and pattern matching 42
 - comparison conditions 33
 - data 693
 - databases 692
 - identifiers 693
 - passwords 693
 - Transact-SQL compatibility 692
 - user IDs 693
 - user-defined data types 692, 693
- CAST function 93, 133
- catalog
 - Adaptive Server Enterprise compatibility 685
 - system tables 573
- catalog store
 - IQ 685
 - monitoring with 494
- CEIL function 135
- CEILING function 136
- chained outer joins 703
- char
 - breaking into terms¶ 359
- CHAR data type
 - about 69
- CHAR function 136
- CHAR_LENGTH function 137
- CHAR_LENGTH64 function
 - description 138
- character data
 - compatibility 688
- CHARACTER data type
 - about 69
- character sets
 - specifying 7
- CHARACTER VARYING data type
 - about 69
 - trimming trailing blanks 71, 688
- CHARINDEX function 139
- CHECK conditions
 - Transact-SQL 694
- check constraints 694
 - enforced 694
 - Transact-SQL compatibility 694
- CHECKPOINT statement
 - backup during checkpoint 524
- clauses
 - ON EXCEPTION RESUME 713
- CLOB data type 39
- COALESCE function 140
- code pages
 - and data storage 71
- COL_LENGTH function 141
- COL_NAME function 141
- collation
 - SORTKEY function 294
- column default
 - not supported 695
- column length 141
- column name 141
- columns
 - and user-defined data types 89
 - naming 24
- comments
 - comment indicators 65
- comparing dates and times 87
- comparisons
 - about 32
- compatibility
 - Adaptive Server Enterprise 683
 - Adaptive Server Enterprise data types 690
 - referential integrity constraints 695
- COMPUTE clause
 - Transact-SQL 702

- computed columns
 - not supported 696
 - concatenating strings 25
 - condition hint strings 49
 - conditions
 - user-supplied for queries 48
 - connection information
 - sp_iqcontext 391
 - connection property value 142
 - CONNECTION_PROPERTY function 142
 - connection-level variables
 - about 59
 - connections
 - determining ID number 240
 - displaying information about 517
 - properties 119
 - consistency checking
 - partitions 373
 - consistent state
 - restoring 481
 - constants
 - in expressions 23
 - Transact-SQL 30
 - CONTAINS conditions
 - with TEXT index 44
 - with WD index 44
 - CONVERT function 93, 144
 - date to integer conversion 147
 - date to string conversion 147
 - integer to date conversion 147
 - string to date conversion 147
 - CORR function 148
 - COS function 149
 - cosine 149
 - COT function 150
 - cotangent 150
 - COUNT function 153
 - COVAR_POP function 151
 - COVAR_SAMP function 152
 - CPU utilization
 - database consistency checker 373
 - CREATE DECRYPTED DATABASE statement
 - 715
 - CREATE DEFAULT statement
 - unsupported 697
 - CREATE DOMAIN statement
 - syntax 90
 - Transact-SQL compatibility 697
 - using 89
 - CREATE ENCRYPTED DATABASE statement
 - 715
 - CREATE INDEX statement
 - IQ 697
 - Transact-SQL 697
 - CREATE RULE statement
 - unsupported 697
 - CREATE TABLE statement
 - Transact-SQL 694
 - CREATE TRIGGER
 - not supported 697
 - creating
 - data types 89, 90
 - CUBE operation
 - GROUPING function 196
 - CUME_DIST function 154
 - CURRENT DATABASE
 - special value 54
 - CURRENT DATE
 - default 54
 - special value 54
 - CURRENT PUBLISHER
 - default 55
 - special value 55
 - CURRENT TIME
 - default 55
 - special value 55
 - CURRENT TIMESTAMP
 - default 55
 - special value 55
 - current user
 - environment settings 15
 - CURRENT USER
 - default 56
 - special value 56
 - cursors
 - displaying information about 395
 - row-level in IQ 709
 - Transact-SQL 709
- ## D
- data
 - case sensitivity 693
 - data type compatibility
 - binary data 689
 - bit data 687
 - character data 688
 - date and time data 689
 - datetime and time data 690

Index

- IMAGE data 692
- Java data 692
- numeric data 691
- TEXT data 691
- data type conversion
 - about 93
 - BIT to BINARY 94
 - BIT to CHAR 96
 - BIT to VARBINARY 94
 - BIT to VARCHAR 96
 - CHAR to BIT 96
 - functions 106
 - VARCHAR to BIT 96
- data type conversion functions 106
 - BIGINTTOHEX 130
 - CAST 133
 - CONVERT 144
 - HEXTOBIGINT 198
 - HEXTOINT 199
 - INTTOHEX 209
- data types
 - Adaptive Server Enterprise 687
 - and compatibility 93
 - binary 78
 - character 69
 - creating 90
 - date and time 84
 - displaying information about 398, 426
 - IMAGE 689, 692
 - IQ 687
 - LONG BINARY 689, 692
 - numeric 73
 - SQL Anywhere 687
 - TEXT 70, 688, 691
 - UNIQUEIDENTIFIERSTR 69
 - user-defined 89
- database administrator
 - roles 686
- database object
 - determining ID 248
 - determining name 249
 - identifying 20
- database options
 - DATE_ORDER 88
 - QUOTED_IDENTIFIER 31
- databases
 - case sensitivity 692
 - determining ID number 176, 242
 - determining name 177
 - properties 120
 - property value 178
 - system procedures 357
 - system tables 573
- DATALENGTH function 155
- date and time data types
 - compatibility 689
- date and time functions 107
 - DATE 156
 - DATEADD 157
 - DATECEILING 158
 - DATEDIFF 161
 - DATEFLOOR 164
 - DATEFORMAT 166
 - DATENAME 168
 - DATEPART 169
 - DATEROUND 170
 - DATETIME 173
 - DAY 174
 - DAYNAME 174
 - DAYS 175
 - DOW 182
 - GETDATE 193
 - getting consistent results 107
 - HOUR 201
 - HOURS 201
 - IQ features 349
 - MINUTE 233
 - MINUTES 234
 - MONTH 236
 - MONTHNAME 236
 - MONTHS 237
 - NOW 244
 - QUARTER 262
 - SECOND 290
 - SECONDS 291
 - TODAY 319
 - WEEKS 340
 - YEAR 345
 - YEARS 346
 - YMD 347
- DATE data type 84
- DATE function 156
- DATE_ORDER option 88
- DATEADD function 157
- DATECEILING function 158
- DATEDIFF function 161
- DATEFLOOR function 164
- DATEFORMAT function 166

- DATENAME function 168
 - DATEPART function 169
 - DATEROUND function 170
 - dates
 - arithmetic expressions 706
 - determining current 244, 319
 - interpreting strings as dates 88
 - queries 87
 - datetime and time data types
 - compatibility 690
 - DATETIME function 173
 - DAY function 174
 - day of the week (DOW) 182
 - DAYNAME function 174
 - DAYS function 175
 - DB_ID function 176
 - DB_NAME function 177
 - DB_PROPERTY function 178
 - dbcc
 - thread usage 373
 - DBCC
 - database verification 371
 - output 377
 - performance 377
 - time to run 377
 - DBCC_LOG_PROGRESS option 377
 - dbinit utility 685
 - dbspaces
 - managing 685
 - preventing read-write operations 375
 - DDL
 - SQL Anywhere 715
 - DECIMAL data type 73
 - default values
 - CURRENT DATABASE 54
 - CURRENT PUBLISHER 55
 - CURRENT USER 56
 - LAST USER 56
 - not supported 695
 - TIMESTAMP 57
 - USER 58
 - defaults
 - CURRENT DATE 54
 - CURRENT PUBLISHER 55
 - CURRENT TIME 55
 - CURRENT TIMESTAMP 55
 - CURRENT USER 56
 - Transact-SQL 697
 - defining a window 103
 - DEGREES function 179
 - delimiting SQL strings 20
 - DENSE_RANK function 180
 - devices
 - managing 685
 - DIFFERENCE function 181
 - directory structure 3
 - disjunction of subquery predicates 36
 - DISK statements
 - unsupported 685
 - distribution functions 102
 - DML
 - SQL Anywhere 715
 - domains 90
 - about 89
 - double quotes
 - database objects 20
 - not allowed in SQL identifiers 20
 - DOW function 182
 - dropleaks mode 375
 - dummy IQ table 99
 - getting consistent results 107
 - DUMMY table 576
- ## E
- ELSE
 - IF expression 28
 - ENDIF
 - IF expression 28
 - environment variables
 - about 6
 - IQLOGDIR 15 10
 - IQPORT 8
 - IQTIMEOUT 10
 - PATH 12
 - SQLCONNECT 13
 - SYBASE 13
 - SYBASE_JRE 14
 - SYBASE_OCS 14
 - error messages
 - ERRORMSG function 183
 - getting text 183
 - ERRORMSG function
 - SQL syntax 183
 - errors
 - Transact-SQL 712, 713
 - estimates
 - optimizer 48
 - EVENT_CONDITION function 184

Index

- EVENT_CONDITION_NAME function 186
- EVENT_PARAMETER function 186
- events
 - displaying information about 420, 426
 - EVENT_CONDITION function 184
 - EVENT_CONDITION_NAME function 186
 - EVENT_PARAMETER function 186
- execution phase hints 52
- EXISTS conditions 44
- EXP function 188
- EXP_WEIGHTED_AVG function 188
- exponential function 188
- expression
 - converting to timestamp 173
 - length in bytes 155
- expression subqueries
 - in IF statements 709
- expressions 22
 - CASE 28
 - Transact-SQL 29
- external libraries
 - listing 360
 - unloading 360
- F**
- files
 - location 4
- FIRST_VALUE function 190
- FLOAT data type 75
- FLOOR function 192
- FOR BROWSE syntax
 - Transact-SQL 704
- FP indexes
 - verifying 375
- FROM clause 112
 - UPDATE and DELETE 707
- functions
 - ABS function 123
 - ACOS function 124
 - Adaptive Server Enterprise system functions 118
 - aggregate 99
 - alphabetical list 123
 - analytical 101
 - ARGN function 125
 - ASCII function 125
 - ASIN function 126
 - ATAN function 127
 - ATAN2 function 128
 - AVG function 129
 - BFILE function 130
 - BIGINTTOHEX function 130
 - BIT_LENGTH function 131
 - BYTE_LENGTH function 132
 - BYTE_LENGTH64 function 133
 - BYTE_SUBSTR function 133
 - BYTE_SUBSTR64 function 133
 - CAST function 133
 - CEIL function 135
 - CEILING function 136
 - CHAR function 136
 - CHAR_LENGTH function 137
 - CHAR_LENGTH64 function 138
 - CHARINDEX function 139
 - COALESCE function 140
 - COL_LENGTH function 141
 - COL_NAME function 141
 - CONNECTION_PROPERTY function 142
 - consistent results 112
 - CONVERT function 144
 - CORR function 148
 - COS function 149
 - COT function 150
 - COUNT function 153
 - COVAR_POP function 151
 - COVAR_SAMP function 152
 - CUME_DIST function 154
 - data type conversion 106
 - DATALENGTH function 155
 - date and time 107
 - DATE function 156
 - DATEADD function 157
 - DATECEILING function 158
 - DATEDIFF function 161
 - DATEFLOOR function 164
 - DATEFORMAT function 166
 - DATENAME function 168
 - DATEPART function 169
 - DATEROUND function 170
 - DATETIME function 173
 - DAY function 174
 - DAYNAME function 174
 - DAYS function 175
 - DB_ID function 176
 - DB_NAME function 177
 - DB_PROPERTY function 178
 - DEGREES function 179
 - DENSE_RANK function 180

- DIFFERENCE function 181
- distribution 102
- DOW function 182
- ERRORMSG function SQL syntax 183
- EVENT_CONDITION function 184
- EVENT_CONDITION_NAME function 186
- EVENT_PARAMETER function 186
- EXP function 188
- EXP_WEIGHTED_AVG function 188
- FIRST_VALUE function 190
- FLOOR function 192
- GETDATE function 193
- GRAPHICAL_PLAN 193
- GROUP_MEMBER function SQL syntax 197
- GROUPING function SQL syntax 196
- HEXTOBIGINT function 198
- HEXTOINT function 199
- HOUR function 201
- HOURS function 201
- HTML_DECODE function 203
- HTML_ENCODE function 203
- HTML_PLAN function 204
- HTTP 112
- HTTP_DECODE function 205
- HTTP_ENCODE function 206
- HTTP_HEADER function 206
- HTTP_VARIABLE function 206
- IFNULL function 207
- INDEX_COL function 208
- INSERTSTR function 208
- INTTOHEX function 209
- IQ extensions 351
- ISDATE function SQL syntax 211
- ISNULL function 212
- ISNUMERIC function SQL syntax 213
- LAST_VALUE 215
- LCASE function 217
- LEFT function 220
- LEN function SQL syntax 222
- LENGTH function 223, 224
- LIST function 224
- LOCATE function 225
- LOG function 226
- LOG10 function 227
- LOWER function 228
- LTRIM function 229
- MAX function 230
- MEDIAN function 231
- MIN function 232
- MINUTE function 233
- MINUTES function 234
- miscellaneous 122
- MOD function 235
- MONTH function 236
- MONTHNAME function 236
- MONTHS function 237
- NEWID function SQL syntax 239
- NEXT_CONNECTION function 240
- NEXT_DATABASE function 242
- NEXT_HTTP_HEADER function 243
- NEXT_HTTP_VARIABLE function 243
- NOW function 244
- NTILE function 244
- NULLIF function 246
- NUMBER function 247
- numeric 102, 112
- OBJECT_ID function 248
- OBJECT_NAME function 249
- OCTET_LENGTH function 250
- PATINDEX function 250
- PERCENT_RANK function 252
- PERCENTILE_CONT function 254
- PERCENTILE_DISC function 256
- PI function 258
- POWER function 258
- PROPERTY function 259
- PROPERTY_DESCRIPTION function 260
- PROPERTY_NAME function 261
- PROPERTY_NUMBER function 261
- QUARTER function 262
- RADIANS function 263
- RAND function 264
- RANK function 265
- ranking 102
- REGR_AVGX function 266
- REGR_AVGY function 267
- REGR_COUNT function 269
- REGR_INTERCEPT function 270
- REGR_R2 function 271
- REGR_SLOPE function 272
- REGR_SXX function 274
- REGR_SXY function 275
- REGR_SYY function 276
- REMAINDER function 278
- REPEAT function 278
- REPLACE function 279
- REPLICATE function 282

- REVERSE function SQL syntax 283
- RIGHT function 284
- ROUND function 285
- ROWID function 287
- RTRIM function 289
- SECOND function 290
- SECONDS function 291
- SIGN function 292
- SIMILAR function 292
- SIN function 293
- SORTKEY function 294
- SOUNDEX function 299
- SPACE function 300
- SQLFLAGGER function 301
- SQRT function 302
- SQUARE function 303
- statistical 102
- STDDEV function 303
- STDDEV_POP function 305
- STDDEV_SAMP function 306
- STR function 307
- STR_REPLACE function SQL syntax 308
- string 114
- STRING function 310
- STRTOUUID function SQL syntax 311
- STUFF function 312
- SUBSTR function 313
- SUBSTRING function 313
- SUBSTRING64 function 315
- SUM function 315
- SUSER_ID function 316
- SUSER_NAME function 317
- TAN function 318
- time series 122
- today 576
- TODAY function 319
- Transact-SQL 704
- TRIM function 319
- TRUNCNUM function 320
- TS_ARMA_AR function 321
- TS_ARMA_CONST function 321
- TS_ARMA_MA function 321
- TS_AUTO_ARIMA function 322
- TS_AUTO_ARIMA_OUTLIER function 322
- TS_AUTO_ARIMA_RESULT_AIC function 323
- TS_AUTO_ARIMA_RESULT_AICC function 323
- TS_AUTO_ARIMA_RESULT_BIC function 323
- TS_AUTO_ARIMA_RESULT_FORECAST_ERROR function 324
- TS_AUTO_ARIMA_RESULT_FORECAST_VALUE function 323
- TS_AUTO_ARIMA_RESULT_MODEL_D function 324
- TS_AUTO_ARIMA_RESULT_MODEL_P function 324
- TS_AUTO_ARIMA_RESULT_MODEL_S function 325
- TS_AUTO_ARIMA_RESULT_RESIDUAL_SIGMA function 325
- TS_AUTO_UNI_AR function 325
- TS_AUTOCORRELATION function 322
- TS_BOX_COX_XFORM function 325
- TS_DIFFERENCE function 326
- TS_DOUBLE_ARRAY function 326
- TS_ESTIMATE_MISSING function 326
- TS_GARCH function 326
- TS_GARCH_RESULT_A function 327
- TS_GARCH_RESULT_AIC function 327
- TS_GARCH_RESULT_USER function 327
- TS_INT_ARRAY function 327
- TS_LACK OF FIT function 328
- TS_LACK OF FIT_P function 328
- TS_MAX_ARMA_AR function 328
- TS_MAX_ARMA_CONST function 329
- TS_MAX_ARMA_LIKELIHOOD function 329
- TS_MAX_ARMA_MA function 329
- TS_OUTLIER_IDENTIFICATION function 330
- TS_PARTIAL_AUTOCORRELATION function 330
- TS_VWAP function 330
- UCASE function 331
- UPPER function 332
- USER_ID function 333
- USER_NAME function 333
- user-defined 121
- UIDTOSTR function SQL syntax 334
- VAR_POP function 335
- VAR_SAMP function 336
- VARIANCE function 338
- WEEKS function 340
- WEIGHTED_AVG function 341
- WIDTH_BUCKET function 343

- windowing aggregate 102
- YEAR function 345
- YEARS function 346
- YMD function 347
- functions, aggregate
 - GROUPING 196
- functions, data type conversion
 - ISDATE 211
- functions, miscellaneous
 - ERRORMSG 183
 - ISNUMERIC 213
 - NEWID 239
- functions, string 204, 222, 283, 308
 - STRTOUUID 311
 - UIDTOSTR 334

G

- GETDATE function 193
- global variables
 - about 58, 60
 - compatibility 62
 - list of 61
- globally unique identifiers
 - SQL syntax for NEWID function 239
- GRAPHICAL_PLAN function 193
- GROUP BY
 - compatibility 702
- GROUP_MEMBER function
 - SQL syntax 197
- GROUPING function 196
- groups
 - Adaptive Server Enterprise 698
- GUIDs
 - SQL syntax for NEWID function 239
 - SQL syntax for STRTOUUID function 311
 - SQL syntax for UIDTOSTR function 334

H

- HEXTOBIGINT function 198
- HEXTOINT function 199
 - ASE_FUNCTION_BEHAVIOR option 200
- hints
 - execution phase 52, 53
 - index preference 50
- HOLDLOCK syntax
 - Transact-SQL 704
- HOUR function 201

- HOURS function 201
- HTML_DECODE function 203
- HTML_ENCODE function 203
- HTML_PLAN 204
- HTML_PLAN function 204
- HTTP functions 112
 - HTML_DECODE 203
 - HTML_ENCODE 203
 - HTTP_DECODE 205
 - HTTP_ENCODE 206
 - HTTP_HEADER 206
 - HTTP_VARIABLE 206
 - NEXT_HTTP_HEADER 243
 - NEXT_HTTP_VARIABLE 243
- HTTP_DECODE function 205
- HTTP_ENCODE function 206
- HTTP_HEADER function 206
- HTTP_VARIABLE function 206

I

- identifiers
 - about 20
 - case sensitivity 693
 - maximum length in SQL Anywhere 20
 - SQL syntax 20
 - uniqueness 693
- identify
 - external libraries 360
- identity columns
 - compatibility 695
 - supported as default value 695
- IF expression 28
- IFNULL function 207
- IMAGE data type 689
 - compatibility 692
- immediate refresh 361
- IN conditions 43
- index preference hints 50
- INDEX_COL function 208
- INDEX_PREFERENCE option 50
- indexes
 - Adaptive Server Enterprise 697
 - IQ 697
 - SQL Anywhere 697
 - Transact-SQL 693
- inserts
 - SQL Anywhere 715
- INSERTSTR function 208

Index

- installation directory
 - about 3
- instances
 - external libraries¶ 360
- INTEGER data type 74
- INTTOHEX function 209
 - ASE_FUNCTION_BEHAVIOR option 209
- IQ Agent
 - port 8
 - wait time 10
- IQ store 685
- iq_dummy table 99, 576
- IQCHARSET environment variable 7
- IQDIR15 environment variable 7
- iqinit utility 685
- IQLANG 9
- IQLANG environment variable 8
- IQLOGDIR15 environment variable 10
- IQPORT environment variable 8
- IQTIMEOUT environment variable
 - specifying IQ Agent wait time 10
- IQTMP15 environment variable 11
- IS NULL conditions 45
- ISDATE function
 - SQL syntax 211
- ISNULL function 212
- ISNUMERIC function
 - SQL syntax 213

J

- Java
 - user-defined functions 122
- Java data types
 - compatibility 692
- Java Runtime Environment
 - setting 14
- JAVA_HOME environment variable 12
- join equality conditions 53
- join indexes
 - displaying information about 447
- join operators
 - ANSI 703
 - Transact-SQL 702
- joins
 - automatic 350
 - outer operators 27
 - Transact-SQL 702

K

- keys
 - displaying information about 467
 - verifying 375
- keywords
 - listing 17
 - SQL 17

L

- language_code 9
- languages
 - specifying 8
- large object data
 - LIKE conditions 39
- LAST USER
 - special value 56
- LAST_VALUE function 215
- LCASE function 217
- LD_LIBRARY_PATH environment variable 12
- LEFT function 220
- LEN function
 - SQL syntax 222
- LENGTH function 223, 224
- LIBPATH environment variable 12
- LIKE conditions 39
 - large object data 39
 - LONG BINARY data 39
 - LONG VARCHAR data 39
- list
 - external libraries 360
- LIST function 224
- literal strings 21, 23
- literals
 - maximum length 353
- load formats
 - Transact-SQL and SQL Anywhere 699
- LOB data type
 - LIKE conditions 39
- local machine
 - environment settings 15
- local variables
 - about 58
- locales
 - setting 9
- LOCATE function 225
- locks
 - displaying 454
- LOG function 226

LOG10 function 227
 logarithm (base 10) 227
 logarithm of a number 226
 login management

- sp_expireallpasswords 362
- sp_iqaddlogin 363
- sp_iqcopyloginpolicy 394, 457

 login policies

- assigning user to 458
- copying 394, 457

 LONG BINARY data type 689, 692

- LIKE conditions 39

 LONG VARCHAR data type 39
 LOWER function 228
 LTRIM function 229
 LVC cells 376

M

master database

- unsupported 684

 mathematical expressions 25
 MAX function 230
 MEDIAN function 231
 memory

- monitoring with sp_iqsysmon 494

 MIN function 232
 MINUTE function 233
 MINUTES function 234
 miscellaneous functions 122

- ARGN 125
- COALESCE 140
- IFNULL 207
- ISNULL 212
- NULLIF 246
- NUMBER 247
- ROWID 287
- SQLFLAGGER 301

 MOD function 235
 MONEY data type 77
 monitor

- sp_iqsysmon procedure 494

 MONTH function 236
 MONTHNAME function 236
 MONTHS function 237
 MPXServerName column 387
 multiplex

- synchronizing query servers 10
- system procedures 386

N

name spaces

- indexes 693

 nchar

- breaking into terms¶ 360

 nested outer joins 703
 NEWID function

- SQL syntax 239

 NEXT_CONNECTION function 240
 NEXT_DATABASE function 242
 NEXT_HTTP_HEADER function 243
 NEXT_HTTP_VARIABLE function 243
 NOT conditions 46
 NOW function 244
 NTILE function 244
 NULL

- Transact-SQL compatibility 694

 null comparisons

- Transact-SQL 703

 NULL value

- about 66

 NULLIF function 29, 246
 NUMBER function 247
 numbers 23
 numeric data types

- compatibility 691

 numeric functions 102, 112

- ABS 123
- ACOS 124
- ASIN 126
- ATAN 127
- ATAN2 128
- CEIL 135
- CEILING 136
- consistent results 112
- COS 149
- COT 150
- DEGREES 179
- EXP 188
- FLOOR 192
- LOG 226
- LOG10 227
- MOD 235
- PI 258
- POWER 258
- RADIANS 263
- RAND 264
- REMAINDER 278
- ROUND 285

Index

- SIGN 292
 - SIN 293
 - SQRT 302
 - SQUARE 303
 - TAN 318
 - TRUNCNUM 320
 - WIDTH_BUCKET 343
- O**
- object
 - determining ID 248
 - determining name 249
 - displaying information about 426
 - renaming 477
 - OBJECT_ID function 248
 - OBJECT_NAME function 249
 - OCTET_LENGTH function 250
 - OLAP
 - DENSE_RANK function 180
 - distribution functions 102
 - GROUPING function 196
 - NTILE function 244
 - numeric functions 102
 - PERCENT_RANK function 252
 - PERCENTILE_CONT function 254
 - PERCENTILE_DISC function 256
 - RANK function 265
 - ranking functions 102
 - row comparison functions 102
 - statistical functions 102
 - STDDEV function 303
 - window function type 103
 - window functions 103
 - window name 103
 - window specification 103
 - windows aggregate functions 102
 - OLAP functions
 - compatibility 705
 - OLAP OVER clause 103
 - ON EXCEPTION RESUME clause
 - Transact-SQL 713
 - Open Client setting 14
 - operators
 - comparison operators 33
 - precedence of 28
 - optimizer
 - estimates 48
 - user-defined selectivity 48
 - options
 - DBCC_LOG_PROGRESS 377
 - QUOTED_IDENTIFIER 31
 - SQL Anywhere 715
 - OR keyword 45
 - outer joins
 - and subqueries 24
 - chained 703
 - nested 703
 - operators 27
 - Transact-SQL 702
 - OVER clause 103
- P**
- partitioned tables
 - verifying 375
 - partitions
 - consistency checking 373
 - passwords
 - adding or modifying 465
 - case sensitivity 693
 - expiring 362
 - PATH environment variable 12
 - PATINDEX function 250
 - pattern matching
 - about 39
 - and collations 42
 - limits 40
 - PERCENT_RANK function 252
 - percentile
 - computing with NTILE function 244
 - PERCENTILE_CONT function 254
 - PERCENTILE_DISC function 256
 - performance
 - monitoring 494
 - sp_iqshowpsex connection information 484
 - sp_iqsysmon procedure 494
 - permissions
 - Adaptive Server Enterprise 698
 - PI function 258
 - population variance function 335
 - portable SQL 700
 - POWER function 258
 - precedence of operators 28
 - predicates
 - about 32
 - disjunction of 36
 - prefetching
 - monitoring with sp_iqsysmon 494
 - primary keys
 - displaying information about 467
-

- generating unique values 239
 - generating unique values using UUIDs 239
 - UUIDs and GUIDs 239
 - PRINT command
 - Transact-SQL 709
 - procedure language
 - overview 707
 - procedures
 - displaying information about 426, 469
 - displaying parameter information 472
 - error handling 712, 713
 - return values 712
 - Transact-SQL 710
 - Transact-SQL overview 708
 - translation 710
 - properties
 - connection 119
 - databases 120
 - description of ID 260
 - determining name 261
 - determining number 261
 - server 120
 - server level 259
 - PROPERTY function 259
 - PROPERTY_DESCRIPTION function 260
 - PROPERTY_NAME function 261
 - PROPERTY_NUMBER function 261
 - publisher
 - SQL Remote 55
- Q**
- QUARTER function 262
 - queries
 - Transact-SQL 700
 - query servers
 - synchronizing 10
 - quotation marks
 - database objects 20
 - SQL identifiers 20
 - QUOTED_IDENTIFIER option 31
 - quotes
 - strings 21
- R**
- RADIANS function 263
 - RAISERROR statement
 - ON EXCEPTION RESUME 713
 - Transact-SQL 713
 - RAND function 264
 - RANK function 265
 - ranking functions 102
 - referential integrity constraints
 - CASCADE not supported 695
 - compatibility 695
 - refresh
 - immediate¶ 361
 - registry entries
 - about 14
 - REGR_AVGX function 266
 - REGR_AVGY function 267
 - REGR_COUNT function 269
 - REGR_INTERCEPT function 270
 - REGR_R2 function 271
 - REGR_SLOPE function 272
 - REGR_SXX function 274
 - REGR_SXY function 275
 - REGR_SYY function 276
 - REMAINDER function 278
 - remote tables
 - columns 560
 - rename objects
 - sp_iqrename procedure 477
 - REPEAT function 278
 - REPLACE function 279
 - in SELECT INTO statement 25, 147, 208, 217, 220, 228, 229, 278, 279, 282, 284, 289, 319, 331, 332
 - REPLICATE function 282
 - reserved words 17
 - listing 17
 - resetclocks
 - sp_iqcheckdb option 374
 - restore operations
 - consistent state 481
 - result sets
 - Transact-SQL 710
 - return values
 - procedures 712
 - REVERSE function
 - SQL syntax 283
 - RIGHT function 284
 - roles
 - Adaptive Server Enterprise 686
 - ROLLUP operation
 - GROUPING function 196
 - ROUND function 285

Index

ROWID function 287

rows

counting 153

RTRIM function 289

rules

Transact-SQL 697

S

sa_char_terms stored procedure 359

sa_checkpoint_execute system procedure 524

sa_conn_list system procedure 531

sa_conn_properties system procedure 531

sa_external_library_unload stored procedure 360

sa_get_table_definition system procedure 571

sa_list_external_library stored procedure 360

sa_nchar_terms stored procedure 360

sa_text_index_vocab stored procedure 361

sample variance function 336

search conditions

about 32

ALL or ANY conditions 38

BETWEEN conditions 39

comparison conditions 33

CONTAINS conditions 44

EXISTS conditions 44

IN conditions 43

IS NULL conditions 45

LEADING SUBSTRING SEARCH

conditions 39

LIKE conditions 39

NOT conditions 46

subqueries 35

three-valued logic 46

truth value conditions 46

SECOND function 290

SECONDS function 291

SELECT INTO

Transact-SQL 707

using REPLACE function 25, 147, 208, 217,
220, 228, 229, 278, 279, 282, 284,
289, 319, 331, 332

SELECT statement

Transact-SQL 700

selectivity

explicit 48

hints 49

user-supplied conditions 48

server

properties 120

server administration

SQL Anywhere and IQ 714

services

registry entries 15

SET OPTION statement

Transact-SQL 700

SHARED syntax

Transact-SQL 704

SIGN function 292

SIGNAL statement

Transact-SQL 713

SIMILAR function 292

SIN function 293

SMALLDATETIME data type 85

SMALLMONEY data type 77

SOME conditions 38

SORTKEY function 294

SOUNDEX function 299

sp_expireallpasswords system procedure 362

sp_iq_reset_identity system procedure 480

sp_iqaddlogin system procedure 363

sp_iqbackupdetails stored procedure 364

sp_iqbackupsummary stored procedure 366

sp_iqbrestoreaction stored procedure 481

sp_iqcardinality_analysis system procedure 368

sp_iqcheckdb

allocation mode 374

check mode 375

DBCC_LOG_PROGRESS option 377

dropleaks mode 375

output 377

performance 377

resetclocks option 374

sample output 377

syntax 371

time to run 377

verify mode 375

sp_iqcheckdb system procedure 370

sp_iqcheckoptions system procedure 379

sp_iqcolumn system procedure 382

sp_iqcolumnuse system procedure 385

sp_iqconnection system procedure 386

sp_iqcontext system procedure 391

sp_iqcopyloginpolicy system procedure 394, 457

sp_iqcursorinfo system procedure 395

sp_iqdatatype system procedure 398

sp_iqdbsize system procedure 401

sp_iqdbspace system procedure 402

sp_iqdbspaceinfo system procedure 405

- sp_iqdbspaceobjectinfo system procedure 409
- sp_iqdbstatistics system procedure 413
- sp_iqdroplogin system procedure 414
- sp_iqemptyfile system procedure 415
- sp_iquestdbspaces system procedure 417
- sp_iquestjoin system procedure 416
- sp_iquestspace system procedure 420
- sp_iqevent system procedure 420
- sp_iqfile system procedure 424
- sp_iqhelp system procedure 426
- sp_iqindex system procedure 433
- sp_iqindex_alt system procedure 433
- sp_iqindexadvice system procedure 437
- sp_iqindexfragmentation system procedure 438
- sp_iqindexinfo
 - displaying index information 441, 443
- sp_iqindexinfo system procedure 440
- sp_iqindexmetadata system procedure 442
- sp_iqindexsize system procedure 443
- sp_iqindexuse system procedure 446
- sp_iqjoinindex system procedure 447
- sp_iqjoinindexsize system procedure 451
- sp_iqlocks system procedure 454
- sp_iqmodifylogin 458
- sp_iqmodifylogin system procedure 458
- sp_iqmpxfilestatus system procedure 461
- sp_iqmpxinconnpoolinfo stored procedure 461
- sp_iqmpxinheartbeatinfo stored procedure 462
- sp_iqmpxinfo stored procedure 462
- sp_iqmpxvalidate stored procedure 462
- sp_iqobjectinfo system procedure 462
- sp_iqpassword system procedure 465
- sp_iqpkeys system procedure 467
- sp_iqprocedure system procedure 469
- sp_iqprocparm system procedure 472
- sp_iqrebuildindex system procedure 475, 482
- sp_iqrename system procedure 477
- sp_iqsetcompression system procedure 357
- sp_iqsharedtempdistrib procedure 483
- sp_iqshowcompression system procedure 357
- sp_iqshowpsex system procedure 484
- sp_iqspaceinfo system procedure 486
 - sample output 486
- sp_iqspaceused system procedure 487
- sp_iqstatistics system procedure 489
- sp_iqstatus system procedure 491
 - sample output 492
- sp_iqsysmon system procedure 494
- sp_iqtable system procedure 501
- sp_iqtablesize system procedure 504
- sp_iqtableuse system procedure 506
- sp_iqtransaction system procedure 506
- sp_iqunusedcolumn system procedure 510
- sp_iqunusedindex system procedure 511
- sp_iqunusedtable system procedure 513
- sp_iqversionuse system procedure 514
- sp_iqview system procedure 515
- sp_iqwho system procedure 517
- sp_iqworkmon system procedure 521
- sp_remote_columns system procedure 560
- sp_remote_primary_keys system procedure
 - syntax 564
- SPACE function 300
- special characters
 - in strings 21
- special values
 - CURRENT DATABASE 54
 - CURRENT DATE 54
 - CURRENT PUBLISHER 55
 - CURRENT TIME 55
 - CURRENT TIMESTAMP 55
 - CURRENT USER 56
 - LAST USER 56
 - SQLCODE 57
 - SQLSTATE 57
 - TIMESTAMP 57
 - USER 58
- SQL
 - IQ dialect differences 349
 - user-defined functions 121
- SQL Anywhere 683
 - administrative roles 686
 - referential integrity constraints 695
- SQL functions
 - compatibility 704
 - ERRORMSG function syntax 183
 - GRAPHICAL_PLAN function syntax 193
 - GROUP_MEMBER function syntax 197
 - GROUPING function syntax 196
 - HTML_PLAN function syntax 204
 - ISDATE function syntax 211
 - ISNUMERIC function syntax 213
 - LEN function syntax 222
 - NEWID function syntax 239
 - REVERSE function syntax 283
 - STR_REPLACE function syntax 308
 - STRTOUUID function syntax 311
 - UIDTOSTR function syntax 334

Index

- SQL statements
 - CREATE DECRYPTED DATABASE 715
 - CREATE ENCRYPTED DATABASE 715
- SQL syntax
 - CURRENT DATABASE special value 54
 - CURRENT PUBLISHER special value 55
 - CURRENT USER special value 56
 - identifiers 20
 - LAST USER special value 56
 - TIMESTAMP special value 57
 - USER special value 58
- SQL92 conformance 349
- SQLCODE
 - special value 57
- SQLCONNECT environment variable 13
- SQLFLAGGER function 301
- SQLSTATE
 - special value 57
- SQRT function 302
- square brackets
 - database objects 20
 - SQL identifiers 20
- SQUARE function 303
- square root function 302, 303
- standard deviation
 - function 303
 - of a popular function 305
 - of a sample function 306
- standards conformity function 301
- statements
 - CREATE DEFAULT 697
 - CREATE DOMAIN 697
 - CREATE RULE 697
 - CREATE TABLE 694
 - DISK INIT 685
 - DISK MIRROR 685
 - DISK REFIT 685
 - DISK REINIT 685
 - DISK REMIRROR 685
 - DISK UNMIRROR 685
 - RAISERROR 713
 - SELECT 700
 - SIGNAL 713
- statistical functions 102
- statistics
 - TEXT indexes 361
- STDDEV function 303
- STDDEV_POP function 305
- STDDEV_SAMP function 306
- stored procedure language
 - overview 707
- stored procedures
 - Adaptive Server Enterprise 716
 - sa_char_terms 359
 - sa_external_library_unload 360
 - sa_list_external_library 360
 - sa_nchar_terms 360
 - sa_text_index_vocab 361
 - sp_iqbackupdetails 364
 - sp_iqbackupsummary 366
 - sp_iqclient_lookup 380
 - sp_iqmpxinconnpoolinfo 461
 - sp_iqmpxincheartbeatinfo 462
 - sp_iqmpxinfo 462
 - sp_iqmpxvalidate 462
 - sp_iqrestoreaction 481
- STR function 307
- STR_REPLACE function
 - SQL syntax 308
- string
 - insert 208
 - length 132, 137
 - position 139
- STRING function 310
- string functions 114
 - ASCII 125
 - BFILE 130
 - BIT_LENGTH 131
 - BYTE_LENGTH 132
 - BYTE_LENGTH64 133
 - BYTE_SUBSTR 133
 - BYTE_SUBSTR64 133
 - CHAR 136
 - CHAR_LENGTH 137
 - CHAR_LENGTH64 138
 - CHARINDEX 139
 - DIFFERENCE 181
 - INSERTSTR 208
 - LCASE 217
 - LEFT 220
 - LENGTH 223, 224
 - LOCATE 225
 - LOWER 228
 - LTRIM 229
 - OCTET_LENGTH 250
 - PATINDEX 250
 - REPEAT 278
 - REPLACE 279

- REPLICATE 282
- RIGHT 284
- RTRIM 289
- SIMILAR 292
- SORTKEY 294
- SOUNDEX 299
- SPACE 300
- STR 307
- STRING 310
- STUFF 312
- SUBSTR 313
- SUBSTRING 313
- SUBSTRING64 315
- TRIM 319
- UCASE 331
- UPPER 332
- string literals
 - maximum length 353
- strings
 - about 21
 - concatenating 25, 282, 310
 - concatenation operators 25
 - constants 21, 23
 - converting to lowercase 217, 228
 - converting to uppercase 331
 - delimiter 31
 - determining length 223, 224
 - determining similarity 292
 - literal strings 21
 - removing blanks 319
 - removing leading blanks 229
 - removing trailing blanks 289
 - replacing substrings 279
 - returning a substring 313
 - SOUNDEX function 299
 - special characters 21
 - Transact-SQL 31
- STRTOUUID function
 - SQL syntax 311
- STUFF function 312
- subqueries
 - Adaptive Server Enterprise 701
 - disjunction of 36
 - in expressions 24
 - in search conditions 35
 - IQ 701
 - IQ implementation 351
 - SQL Anywhere 701
- SUBSTR function 313
- SUBSTRING function 313
- SUBSTRING64 function
 - description 315
- SUM function 315
- summary 364, 461, 462
- SUSER_ID function 316
- SUSER_NAME function 317
- SYBASE environment variable 13
- Sybase IQ User Administration
 - sp_iqdroplogin 414
- SYBASE_JRE environment variable 14
- SYBASE_OCS environment variable 14
- syntax
 - CURRENT DATABASE special value 54
 - CURRENT PUBLISHER special value 55
 - CURRENT USER special value 56
 - LAST USER special value 56
 - SQL identifiers 20
 - TIMESTAMP special value 57
 - USER special value 58
- SYSIQBACKUPHISTORY system view 608
- SYSIQBACKUPHISTORYDETAIL system view 609
- SYSIQDBFILE system view 611
- SYSIQDBSPACE system view 612
- SYSIQIDX system view 613
- SYSIQJOINIDX system view 615
- SYSIQJOINIXCOLUMN system view 616
- SYSIQJOINIXTABLE system view 617
- SYSIQLOGICALSERVER system view 618
- SYSIQLLOGINPOLICYLSINFO system view 618
- SYSIQLSLOGINPOLICYOPTION system view 618
- SYSIQLSMEMBER system view 618
- SYSIQLSPOLICY system view 619
- SYSIQPARTITIONCOLUMNS system view 620
- SYSIQTAB system view 620
- SYSPARTITION system view 631
- SYSPARTITIONKEY system view 632
- SYSPARTITIONSCHEME system view 632
- SYSSUBPARTITIONKEY system view 648
- system administrator
 - Adaptive Server Enterprise 686
- system catalog
 - Adaptive Server Enterprise compatibility 685
 - Transact-SQL 678
- system functions 116
 - COL_LENGTH 141
 - COL_NAME 141

Index

- CONNECTION_PROPERTY 142
- DATALENGTH 155
- DB_ID 176
- DB_NAME 177
- DB_PROPERTY 178
- EVENT_CONDITION 184
- EVENT_CONDITION_NAME 186
- EVENT_PARAMETER 186
- INDEX_COL 208
- NEXT_CONNECTION 240
- NEXT_DATABASE 242
- OBJECT_ID 248
- OBJECT_NAME 249
- PROPERTY 259
- PROPERTY_DESCRIPTION 260
- PROPERTY_NAME 261
- PROPERTY_NUMBER 261
- SUSER_ID 316
- SUSER_NAME 317
- Transact-SQL 706
- USER_ID 333
- USER_NAME 333
- system procedures
 - about 357
 - displaying information about 426
 - sa_checkpoint_execute 524
 - sa_conn_list 531
 - sa_conn_properties 531
 - sa_get_table_definition 571
 - sp_expireallpasswords 362
 - sp_iqaddlogin 363
 - sp_iqbackupdetails 364
 - sp_iqbackupssummary 366
 - sp_iqcardinality_analysis 368
 - sp_iqcheckdb 370
 - sp_iqcheckoptions 379
 - sp_iqclient_lookup 380
 - sp_iqcolumn 382
 - sp_iqcolumnuse 385
 - sp_iqconnection 386
 - sp_iqcontext 391
 - sp_iqcopyloginpolicy 394, 457
 - sp_iqcursinfo 395
 - sp_iqdatatype 398
 - sp_iqdbsize 401
 - sp_iqdbspaceobjectinfo 409
 - sp_iqdbstatistics 413
 - sp_iqdroplogin 414
 - sp_iqemptyfile 415
 - sp_iquestdbspaces 417
 - sp_iquestjoin 416
 - sp_iquestspace 420
 - sp_iqevent 420
 - sp_iqfile 424
 - sp_iqhelp 426
 - sp_iqindex 433
 - sp_iqindex_alt 433
 - sp_iqindexadvice 437
 - sp_iqindexsize 443
 - sp_iqindexuse 446
 - sp_iqjoinindex 447
 - sp_iqjoinindexsize 451
 - sp_iqmodifylogin 458
 - sp_iqmpxfilestatus 461
 - sp_iqmpxinconnpoolinfo 461
 - sp_iqmpxinheartbeatinfo 462
 - sp_iqmpxinfo 462
 - sp_iqmpxvalidate 462
 - sp_iqobjectinfo 462
 - sp_iqpassword 465
 - sp_iqpkeys 467
 - sp_iqprocedure 469
 - sp_iqprocparm 472
 - sp_iqrename 477
 - sp_iqrestoreaction 481
 - sp_iqsetcompression 357
 - sp_iqsharedtempdistrib 483
 - sp_iqshowcompression 357
 - sp_iqshowpsex 484
 - sp_iqspaceinfo 486
 - sp_iqspaceused 487
 - sp_iqstatistics 489
 - sp_iqstatus 491
 - sp_iqsysmon 494
 - sp_iqtable 501
 - sp_iqtablesize 504
 - sp_iqtableuse 506
 - sp_iqtransaction 506
 - sp_iqunusedcolumn 510
 - sp_iqunusedindex 511
 - sp_iqunusedtable 513
 - sp_iqversionuse 514
 - sp_iqview 515
 - sp_iqwho 517
 - sp_iqworkmon 521
 - sp_remote_columns 560
 - sp_remote_primary_keys 564

system security officer
 Adaptive Server Enterprise 686
 System tables
 about 573
 Adaptive Server Enterprise compatibility 685
 displaying information 426
 DUMMY 576
 system variables 60
 system views
 Adaptive Server Enterprise 717
 consolidated 577
 SYSCOLUMNS ASE compatibility view 586
 SYSDBSPACEPERM 589
 SYSEVENTTYPE 592
 SYSINDEXES ASE compatibility view 607
 SYSIQBACKUPHISTORY 608
 SYSIQBACKUPHISTORYDETAIL 609
 SYSIQCOLUMN 610
 SYSIQDBFILE 611
 SYSIQDBSPACE 612
 SYSIQFILE 612
 SYSIQIDX 613
 SYSIQINFO 614
 SYSIQITAB 620
 SYSIQJOINIDX 615
 SYSIQJOININDEX 616
 SYSIQJOINIXCOLUMN 616
 SYSIQJOINIXTABLE 617
 SYSIQLOGICALSERVER system view 618
 SYSIQLOGINPOLICYLSINFO 618
 SYSIQLSLOGINPOLICIES 619
 SYSIQLSLOGINPOLICYOPTION 618
 SYSIQLSMEMBER 618
 SYSIQLSMEMBERS 619
 SYSIQLSPOLICY 619
 SYSIQMPXSERVER 619
 SYSIQOBJECTS ASE compatibility view
 619
 SYSIQPARTITIONCOLUMN 620
 SYSIQTABCOL 621
 SYSIQTABLE system view 622
 SYSIQVINDEXT ASE compatibility view 623
 SYSLOGINS ASE compatibility view 627
 SYSOBJECTS ASE compatibility view 629
 SYSPARTITION 631
 SYSPARTITIONKEY 632
 SYSPARTITIONSCHEME 632
 SYSREMOPTOPTION 641
 SYSSCHEDULE 646

SYSSUBPARTITIONKEY 648
 SYSTYPES ASE compatibility view 670
 SYSUSERS ASE compatibility view 676
 SYSVIEW 676

T

tables
 displaying information about 426
 iq_dummy 99
 Transact-SQL 694
 TAN function 318
 tangent 318
 temporary tables
 Transact-SQL 696
 terms
 breaking 359, 360
 row position 359, 360
 TEXT data type 70, 688
 compatibility 688, 691
 TEXT indexes
 statistics 361
 THEN
 IF expression 28
 threads
 dbcc 373
 three-valued logic
 about 46
 NULL value 66
 TIME data type 85
 time series functions 122
 times
 queries 87
 TIMESTAMP
 converting an expression 173
 data type 85
 data type compatibility 689, 690
 special value 57
 TINYINT data type 74
 TODAY function 319, 576
 trailing blanks
 trimming 71, 688
 Transact-SQL
 about 683
 batches 708
 bitwise operators 26
 comparison conditions 34
 constants 30
 creating compatible databases 692
 expressions 29

Index

- joins 702
- local variables 59
- outer join operators 27
- overview 683
- procedure language overview 707
- procedures 708
- referential integrity constraints 695
- result sets 710
- strings 31
- system catalog 678
- user-defined data types 92
- variables 711
- writing portable SQL 700

Transact-SQL compatibility

- databases 692

transaction management

- monitoring with `sp_iqsysmon` 494

triggers

- not supported 697

TRIM function 319

trimming trailing blanks 71, 688

TRUNCNUM function 320

TS_ARMA_AR function 321

TS_ARMA_CONST function 321

TS_ARMA_MA function 321

TS_AUTO_ARIMA function 322

TS_AUTO_ARIMA_OUTLIER function 322

TS_AUTO_ARIMA_RESULT_AIC function 323

TS_AUTO_ARIMA_RESULT_AICC function 323

TS_AUTO_ARIMA_RESULT_BIC function 323

TS_AUTO_ARIMA_RESULT_FORECAST_ERR OR function 324

TS_AUTO_ARIMA_RESULT_FORECAST_VAL UE function 323

TS_AUTO_ARIMA_RESULT_MODEL_D function 324

TS_AUTO_ARIMA_RESULT_MODEL_P function 324

TS_AUTO_ARIMA_RESULT_MODEL_S function 325

TS_AUTO_ARIMA_RESULT_RESIDUAL_SIG MA function 325

TS_AUTO_UNI_AR function 325

TS_AUTOCORRELATION function 322

TS_BOX_COX_XFORM function 325

TS_DIFFERENCE function 326

TS_DOUBLE_ARRAY function 326

TS_ESTIMATE_MISSING function 326

TS_GARCH function 326

TS_GARCH_RESULT_A function 327

TS_GARCH_RESULT_AIC function 327

TS_GARCH_RESULT_USER function 327

TS_INT_ARRAY function 327

TS_LACK OF FIT function 328

TS_LACK OF FIT_P function 328

TS_MAX_ARMA_AR function 328

TS_MAX_ARMA_CONST function 329

TS_MAX_ARMA_LIKELIHOOD function 329

TS_MAX_ARMA_MA function 329

TS_OUTLIER_IDENTIFICATION function 330

TS_PARTIAL_AUTOCORRELATION function 330

TS_VWAP function 330

type conversions 93

U

UCASE function 331

UNION

- in subqueries 701

UNIQUEIDENTIFIER data type 82

UNIQUEIDENTIFIERSTR data type

- about 69

universally unique identifiers

- SQL syntax for `NEWID` function 239

unloading

- external libraries¶ 360

UPPER function 332

usefulness hints 53

USER

- special constant 576
- special value 58

user administration

- See login management

user IDs

- Adaptive Server Enterprise 698
- case sensitivity 693
- determining from user name 316, 333

user name

- determining from user ID 317, 333

USER_ID function 333

USER_NAME function 333

user-defined data types

- about 89
- case-sensitivity 693
- CREATE DOMAIN statement 90
- Transact-SQL 92

- user-defined functions 121
 - compatibility 706
- user-supplied condition hint strings 49
- user-supplied condition hints, guidelines and usage 54
- user-supplied condition selectivity 48
- user-supplied conditions
 - for queries 48
- user-supplied hints on join equality conditions 53
- users
 - adding 363
 - displaying information about 517
 - dropping 414
 - modifying 458
- utilities
 - SQL Anywhere 714
- UUIDs
 - SQL syntax for NEWID function 239
 - SQL syntax for STRTOUUID function 311
 - SQL syntax for UUIDTOSTR function 334
- UUIDTOSTR function
 - SQL syntax 334

V

- VAR_POP function 335
- VAR_SAMP function 336
- VARBINARY data type 79
- VARCHAR data type
 - about 69, 70
 - trimming trailing blanks 71, 688
- variables
 - about 58
 - connection-level 59
 - global 58, 60
 - local 58
 - Transact-SQL 711
- VARIANCE function 338

- verifying
 - indexes 375
 - keys 375
 - partitioned tables 375
- views
 - displaying information about 426
 - updatable 707

W

- WEEKS function 340
- WEIGHTED_AVG function 341
- WHERE clause
 - Transact-SQL 702
- WIDTH_BUCKET function 343
- window functions
 - window function type 103
 - window name or specification 103
 - window partition 103
- window functions, defining 103
- window name 103
- window specification 103
- window type 103
- windows aggregate functions 102
- WITHIN GROUP clause 105

Y

- YEAR function 345
- YEARS function 346
- YMD function 347

Z

- zero-length strings
 - Transact-SQL 703

