



Glossary

Adaptive Server[®] Enterprise

15.7

DOCUMENT ID: DC36147-01-1570-01

LAST REVISED: September 2011

Copyright © 2011 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.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. 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 Sun Microsystems, Inc. in the U.S. and other countries.

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

IBM and Tivoli are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

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.

Glossary

This section lists and describes the terms used in Adaptive Server® documentation.

Select the letter below to go to that section of the *Glossary*.

A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | U | V | W | X

A

abstract plan

Instructions to the Adaptive Server Enterprise query processor on the access path to use to manipulate the data for a query, for example **join** order, join algorithms, index usage, and so on.

abstract plan derived table

A table used in query processing, the optimization and execution of queries. An abstract plan derived table differs from a SQL derived table in that it exists as part of an abstract plan and is invisible to the end user. See also derived table and SQL derived table.

access

The use of a select, insert, update, or delete command on a table or view.

access method

The method used to find the data needed to produce results for a query. Access methods can be serial or parallel. Serial access methods include: **table scan**, **nonclustered index** access, **clustered index** access.

account locking

An **Adaptive Server** facility that prevents a user from logging in to Adaptive Server, but permits the user's account to own databases and objects. The user can be denied access without disrupting the permissions the user may have granted to other users or the availability of database objects owned by the user. Also known as login locking.

action

(Clusters concept) A user-executed command that stops one or more instances at a specified time to initiate a planned failover, downtime, or other administrative task. An action changes the state of an instance.

adapter

A Sybase® Enterprise Event Broker (SEEB) component that provides an interface between SEEB and external applications or messaging systems. An adapter detects events, validates event contents, and passes them to an inflow processor. Adapters also receive SEEB events at outflow processors and export the events to external applications.

Adaptive Server	The server in the Sybase client/server architecture (called SQL Server in versions earlier than 11.5). Adaptive Server manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.
Adaptive Server engine	See engine .
Adaptive Server login	The name a user uses to log in to Adaptive Server . A login is valid if Adaptive Server has an entry for that user in the system table <code>syslogins</code> .
address lock	A type of lock applied to certain resources such as network buffers or internal structures. Address locks are also used to lock index pages.
ad hoc statement	A SQL statement that is not part of a stored procedure and is not dynamic SQL. Transact-SQL batches are composed of ad hoc SQL statements.
adjusted query plan	Adaptive Server may create an adjusted query plan to compensate for the available worker processes. It is generated at runtime and compensates for the lack of available worker processes.
affinity	Describes a process in which a certain Adaptive Server task runs only on a certain engine (task affinity), a certain engine handles network I/O for a certain task (network I/O affinity), or a certain engine runs only on a certain CPU (engine affinity).
aggregate function	A function that generates summary values that appear as new columns in the query results. The aggregate functions available in Transact-SQL are average (avg), maximum (max), minimum (min), sum (sum), and count of the number of items (count).
alias	A pseudonym that allows an Adaptive Server user to be known in a database as another user.
allocation page	The first page of an allocation unit , which tracks the use of all pages in the allocation unit.
allocation unit	A logical unit of 1/2 MB, or 256 2K pages. The disk init command initializes a new database file for Adaptive Server and divides it into allocation units.
allpages-locking scheme	One of the three locking schemes in Adaptive Server . In allpages locking, locks are acquired on data and index pages. See also datapages-locking scheme and datarows-locking scheme .
API	See application program interface (API) .

application	A client program that interacts with Adaptive Server. Specifically, an application name refers to the program_name column in a sysprocesses table entry. The client libraries provide a mechanism to specify the program_name of an application.
application program interface (API)	A subroutine that allows client applications to interface with Adaptive Server . Also known as libraries .
argument	A value supplied to a function or procedure that is required to evaluate the function.
arithmetic expression	An expression that contains only numeric operands and returns a single numeric value. In Transact-SQL, the operands can be of any Adaptive Server numeric datatype. They can be functions, variables, parameters, or they can be other arithmetic expressions. Also called a numeric expression .
arithmetic operator	Symbols that enable you to create an arithmetic expression in SQL statements. Addition (+), subtraction (-), division (/), and multiplication (*) can be used with numeric columns. Module (%) can be used only with integer datatypes. See also comparison operator .
association key	The association key combination of user ID, group ID and query text means that for a given user, there cannot be two queries in the same abstract plan group that have the same query text, but different query plans.
asymmetrical	A high availability system consisting of one primary companion and one secondary companion. In an asymmetrical system, only the primary companion can fail over. In this system, the secondary Adaptive Server is also known as a “hot standby.”
audit trail	Audit records stored in the sybsecurity database.
auditing	The act of recording security-related system activity that can be used to detect penetration of the system and misuse of system resources.
authorization	See roles .
automatic recovery	A process that runs every time Adaptive Server is restarted. The process ensures that all transactions that completed before Adaptive Server stopped are brought forward and all incomplete transactions are rolled back.

B

B-tree	Short for balanced tree, or binary tree. Adaptive Server uses B-tree indexing. All leaf pages in a B-tree are the same distance from the root page of the index. B-trees provide consistent and predictable performance, good sequential and random record retrieval, and a flat tree structure.
backup	A copy of a database or transaction log, used to recover from a media failure.
Backup Server	Backup Server performs local or remote backups (dumps) and restores (loads) on selected databases and transaction logs on behalf of Adaptive Server . A Backup Server must be running on the same system as Adaptive Server.
base column	A column from which a computed column or function-based index is defined.
base date	January 1, 1900; the date supplied by Adaptive Server when a user does not specify a value for a date column.
base instance	(Clusters concept) An instance assigned to a logical cluster on which a logical cluster normally runs.
base priority	The priority of a task or thread when it is created. At any time, the execution priority of a thread may be greater than or equal to its base priority, but the execution priority of a thread will never be lower than its base priority.
base tables	The permanent tables on which a view is based. Also called underlying tables.
batch	One or more Transact-SQL statements terminated by an end-of-batch signal, which submits them to Adaptive Server for processing.
bcp	See bulk copy .
Bean	A server-side component that contains the business logic of a program. At runtime, the application clients execute the business logic by invoking the enterprise Bean's functions.
binding	An association between a default or a rule and a table column or a user datatype. When a rule or a default is bound to a table column or a specified user datatype, it affects the data that can be entered in the column. Binding can also refer to an association between a named data cache and a database, table, or index. When a database, table, or index is bound to a named data cache, reads from the database, table, or index go into the named cache.
blocking	Waiting for a lock; tasks that need to acquire a lock on a row, page, or table must wait, or block, if another process holds an incompatible lock.
Boolean expression	An expression that evaluates to TRUE (1), or FALSE (0). Boolean expressions are often used in control of flow statements, such as if or while conditions.

branch processor	A point on a route that includes switching logic. A single event can enter the branch processor, and multiple events can depart to sub-branches. A typical application for a branch is to fan out copies of a message to multiple sub-branches.
buffer	A unit of storage in a memory pool. A single buffer cache can have pools configured for different I/O sizes, or buffer sizes. All buffers in a pool are the same size. If a pool is configured for 16K I/O, all buffers are 16K, holding eight data pages. Buffers are treated as a unit; all data pages in a buffer are read, written, or flushed from cache simultaneously.
buffer cache	An area of memory within Adaptive Server that contains the in-memory images of database pages, as well as the data structures required to manage the pages. Each cache is given a unique name that is used for configuration purposes. By default, Adaptive Server has a single cache named “default data cache.” Caches configured by users are called “user-defined caches.” Buffer caches are also referred to as data caches and named caches.
buffer pool	An area of memory within a buffer cache that contains a set of buffers linked together on an MRU/LRU (most recently used/least recently used) list.
buffer reuse strategy	Reading pages into the data cache at the LRU end of the cache chain, so that the same buffer is available for reuse immediately. This strategy keeps select commands that require large numbers of page reads from flushing other data from the cache.
buffer replacement strategy	see buffer reuse strategy .
built-in functions	A wide variety of functions that take one or more arguments and return results. The built-in functions include the aggregate function , mathematical function , system function , string function , text and image function , date function , and datatype conversion function .
bulk copy	The process of copying data in and out of databases. In Adaptive Server , this operation is performed with the bcp utility.
bushy parallelism	Occurs when several CPUs execute different subplans of a complex query plan in parallel.
bushy tree plan	A query plan in which some join operations have two or more direct children that are also join operations. Non-bushy plans are typically referred to as “left deep trees,” in which the right child is a scan operator, or a “join lists,” in which tables are joined in linear order.

business event	An event generated by a specific business operation, such as a stock price change transmission. In contrast, a user login is not a business event.
business rule	A rule that restricts input and updates based on real-world requirements. For example, a sales database can be required to forbid the sale of any items for which there is insufficient stock. Adaptive Server can check each row before it is added to the <code>sales_orders</code> table and enforce the requirement using the business rule on the quantity column.
C	
CA certificate	Also known as trusted root certificates, the CA certificate is a list of trusted CAs loaded by the server at a start-up.
cache	See buffer cache .
cache hit ratio	For many processes, Adaptive Server uses an in-memory cache. The cache hit ratio is the percentage of times a needed page or result was found in the cache. For data pages, the cache hit ratio is the percentage of page requests that are serviced by the data cache compared to requests that require disk I/O.
cache strategy	The optimizer specification of the characteristics of the buffer cache to be used by the execution engine for processing pages of explicit tables referenced in the query, or transient information used during internal processing of the query.
candidate key	A primary key or unique constraint column. A table can have multiple candidate keys.
Cartesian product	All the possible combinations of the rows from each of the tables specified in a join. The number of rows in the Cartesian product is equal to the number of rows in the first table times the number of rows in the second table. Once the Cartesian product is formed, the rows that do not satisfy the join conditions are eliminated.
cascading delete	A delete operation that affects related data in other tables.
cascading menu	A submenu that appears to the right of a selected pull-down menu item. An arrow next to an item on a pull-down menu indicates the existence of a cascading menu.
catalog stored procedure	A type of system procedure that returns data from the system tables in tabular format.
certauth	A Sybase utility to convert a server certificate request to a CA-signed certificate.

certpk12	A Sybase utility to convert third-party PKCS12-format certificates into a format that is understood by Adaptive Server and Open Client/Open Server.
certreq	A Sybase utility to generate public and private key pairs and certificate requests.
chained transaction mode	Determines whether or not Adaptive Server automatically starts a new transaction on the next data retrieval or data modification statement. When chained transaction mode is set on outside a transaction, the next data retrieval or data modification statement begins a new transaction. This mode is ANSI-compliant. It ensures that every SQL data retrieval and data modification statement occurs inside a transaction. Chained transaction mode may be incompatible with existing Transact-SQL programs. Chained transaction mode is off by default. Applications that require ANSI SQL (such as the Embedded SQL precompiler) should automatically set the chained option on at the beginning of each session.
character expression	An expression that returns a single character-type value. It can include literals, concatenation operators, functions, and column identifiers.
character set	A set of specific (usually standardized) characters with an encoding scheme that uniquely defines each character. ASCII and ISO 8859-1 (Latin 1) are two common character sets.
character set conversion	Changing the encoding scheme of a set of characters on the way into or out of Adaptive Server . Conversion is used when Adaptive Server and a client communicating with it use different character sets. For example, if Adaptive Server uses ISO 8859-1 and a client uses Code Page 850, character set conversion must be turned on so that both server and client interpret in the same way the data passing back and forth.
check constraint	A constraint placed on the check command that limits the values users can insert into a column of a table. A check constraint specifies the search_condition that values must pass before being inserted into the table.
checkpoint	The point at which all data pages that have been changed are guaranteed to have been written to the database device.
cell	A pair of contiguous values in a distribution or histogram . Each step is associated with the set of rows having column values that fall between the upper and lower bound of the step. Also called a step .
certificates	An official document used to verify the identity of one entity (server, router, Web site, and so on) to another.

child page pointer	A pointer in an index that points to the next lowest level of the index containing the key values. A child page pointer consists of a key and a row ID.
CIS	See Component Integration Services (CIS) .
CipherSuite	A preferential list of key-exchange algorithms, hashing methods, and encryption methods used by SSL-enabled applications.
class	In Java, a class provides the definition of an object; an object is an instantiation of a class. An application with a user interface includes forms, which are classes with specific support for the design environment.
classification	A hierarchical level of security; for example, Top Secret, which has a higher classification value than Secret.
clause	A set of keywords and options that tailor a Transact-SQL command to meet a particular need. Also called a keyword phrase.
client	The user's side of a client/server arrangement; can refer to the software making the calls to the server or to the machine running the client software.
client cursor	A cursor declared through Open Client calls or Embedded SQL. Open Client keeps track of the rows returned from Adaptive Server and buffers them for the application. Updates and deletes to the result set of client cursors can be done only through Open Client calls.
client migration	(Cluster concept) The movement of an established client connection from one instance to another. The client is migrated from the old to the new instance without the client application being aware of the migration. Client migration is used for dynamic load distribution and for administrative actions such as logical cluster failback.
client/server architecture	A computer system architecture in which clients request a service and a server provides that service. Each machine can then specialize in the tasks it is best suited for.
client task	A thread spawned to service a client request.
cluster	A collection of homogeneous nodes in a network that operate as a single system. Each node has its own CPU and memory. All nodes communicate with each other through private and high-speed communication pathways.
cluster lock manager (CLM)	(Cluster concept) The server module that provides distributed locking services for the cluster. The CLM enables sharing of buffers, global objects, and metadata among the instances.

clusterware	(Cluster concept) Sybase software included in Adaptive Server that enables the shared-disk cluster.
clustered index	An index in which the physical order and the logical (indexed) order is the same. The leaf level of a clustered index represents the data pages themselves. A table can have only one clustered index.
code generation	The mechanism in query compilation that converts the optimizer best plan representation into the execution engine plan representation. The representation of the plan used by the optimizer is designed for efficient comparisons between competing best plans, whereas the representation of the plan used by the execution engine is designed for efficient execution.
code set	See character set .
collating sequence	See sort order .
column	A data value that describes one characteristic of an entity . A column contains an individual data item within a row or record. Also called a field .
column-level constraint	Limits the values of a specified column . Place column-level constraints after the column name and datatype in the create table statement, before the delimiting comma.
command	An instruction that specifies an operation to be performed by the computer. Each command or SQL statement begins with a keyword, such as insert, that names the basic operation performed. Many SQL commands have one or more keyword phrases , or clauses , that tailor the command to meet a particular need.
command dialog box	A dialog box that opens when you execute a command that requires information on how it should execute.
command permission	A permission that applies to commands. See also object permission .
command terminator	The end-of-batch signal that sends the batch to Adaptive Server for processing.
Common Object Request Broker Architecture (CORBA)	Common Object Request Broker Architecture. An open distributed object computing infrastructure—defined by the Object Management Group (OMG)—that automates many common network programming tasks. A key component in the architecture, the Object Request Broker (ORB), is the middleware that establishes the client/server relationships between objects and enables clients to invoke methods on a server object on the same machine or on a machine on a network.
companion server	Each Adaptive Server in a high availability system is a companion.

comparison operator	A symbol used to compare one value to another in a query. Comparison operators include equal to (=) greater than (>), less than (<), greater than or equal to (>=), less than or equal to (<=), not equal to (!=), not greater than (!>), and not less than (!<). See also arithmetic operator .
compartment	One of a set of non hierarchical values used with classifications to make up a sensitivity label. Compartments usually represent topics or work groups.
compatible datatypes	Datatypes that are automatically converted for implicit or explicit comparison.
compilation	The phase of query processing that analyzes query text and creates a query plan to provide to the execution engine.
compiled object	Any object that requires entries in the sysprocedures table. Check constraints, declarative defaults, rules, stored procedures, triggers, and views are compiled objects. These objects are described by source text . Adaptive Server uses compiled objects to contain vital information about each database and to help you access and manipulate data.
component	Packages consist of components, and a component is an application object that consists of one or more methods.
Component Integration Services (CIS)	Component Integration Services extends Adaptive Server 's capabilities and provides enhanced inter operability. It is the core inter operability feature of OmniConnect. Component Integration Services allows Adaptive Server and OmniConnect to present a uniform view of enterprise data to client applications and provides location transparency to enterprise-wide data sources.
composite indexes	Indexes that involve more than one column. Use composite indexes when two or more columns are best searched as a unit because of their logical relationship.
composite key	An index key that includes two or more columns; for example, authors(aulname, aufname).
computed column	A column defined by an expression, whether from a regular column in the same row, a function, an arithmetical operator, a path name, or other expression.
computed column index	An index that contains one or more computed columns as index key.
concatenation	Combining expressions to form longer expressions. The expressions can include any combination of binary or character strings or column names.

concurrency	Concurrent execution (multiprocessing) of independent and possibly competing processes or transactions.
consumer process	In parallel sorts, consumer processes simultaneously sort a discrete range of data received from producer processes. Consumer processes hand off the sorted data to a coordinating process .
constant expression	An expression that returns the same value each time the expression is used. In Transact-SQL syntax statements, a constant expression does not include variables or column identifiers.
constraint	A rule applied to a database object that ensures that all entries in the database object to which it applies satisfy a particular condition. For example, a column may have a constraint requiring that all values in the column be unique.
container	Determines how an event processor interprets the sequence of bytes that makes up the event body. SEEB supports the following containers: binary event, SQL rowset event, SQL typed property event, text event, rmsg, aggregate, Java object, and property.
context-sensitive protection	Protection that provides certain permissions or privileges depending on the identity of the user. This type of protection can be implemented using views and the user_id built-in function.
control-break report	A report or data display that breaks data into groups and generates summary information for each break. The breaks control the generation of summary data.
control-of-flow language	Transact-SQL's programming-like constructs (such as if, else, while, and goto) that control the flow of execution of Transact-SQL statements.
conversion	See character set conversion .
coordinating process	In parallel sorts, the coordinating process merges the results of the consumer processes into a single result set. For queries, the result set is the final, sorted data. For a create index statement, the coordinating process merges subindexes into one, final index.
CORBA	See Common Object Request Broker Architecture (CORBA) .
correlated subquery	A subquery that cannot be evaluated independently, but depends on the outer query for its results. Also called a repeating subquery because the subquery is executed once for each row that might be selected by the outer query. See also nested queries .

correlation names	Distinguish the different roles a particular table plays in a query, especially a correlated query or self-join . Assign correlation names in the from clause and specify the correlation name after the table name: <pre>select au1.au_fname, au2.au_fname from authors au1, authors au2 where au1.zip = au2.zip</pre>
cost based pruning	Query optimization technique to use estimated costs of subplans to avoid analyzing more expensive subplan alternatives.
counter	A measurable performance item that can be reported by the Windows NT Performance Monitor. This is generic Windows NT terminology for describing a mechanism for producing statistical performance information. Adaptive Server maintains a special set of counters to measure and report on Adaptive Server objects or events on the Windows NT platform.
covered query	See index covering .
covering	See index covering .
CPU cost	The optimizer's estimated amount of CPU needed to process a query.
current value	The current value is the value of the variable just before the update of the current row.
cursor	A named select statement that retrieves one or more rows from a given table, and allows you to modify or delete the rows individually. Cursors consist of two parts: the cursor result set and the cursor position .
cursor position	Indicates the current row of the cursor. You can explicitly reference that row using statements designed to support cursors, such as delete and update. Change the current cursor position through fetch, which moves the current cursor position one or more rows down the cursor result set .
cursor result set	The set of rows resulting from the execution of the select statement associated with the cursor.
cursor scan	The process of generating a cursor result set .
cursor scope	The context in which the cursor is used. A cursor's existence depends on its scope: within a particular user session, within a stored procedure, or within a trigger.

cursor stability	A locking level or isolation level in which Adaptive Server has a shared lock on the base table pages that contain a current cursor row. The page remains locked until the cursor is no longer positioned on the page (as a result of fetches). If the base table has an index, the corresponding index pages also have shared locks.
D	
DAC	See discretionary access controls (DAC) .
data cache	See buffer cache .
data definition	The process of setting up databases and creating database objects such as tables, indexes, rules, defaults, procedures, triggers, and views. See also source text .
data definition language (DDL)	Transact-SQL scripts that capture the state of a database. These scripts can re-create the database in the same location or on a different server.
data dictionary	The system tables that contain a description of each database object and how it is structured.
data flow engine	A descriptive term for the execution engine that implies one large plan with fewer materialization steps.
data integrity	The correctness and completeness of data within a database.
data modification	Adding, deleting, or changing information in the database with the insert, delete, and update commands.
data retrieval	The act of requesting data from the database and receiving the results. See also query .
data server interface (DSI)	Replication Server threads corresponding to a connection between a Replication Server and a database. DSI threads, consisting of a scheduler thread and one or more executor threads, submit transactions from the DSI outbound queue to a replicate data server. The scheduler thread groups the transactions by commit order and dispatches them to the executor threads. The executor threads map functions to function strings and execute the transactions in the replicate database.

data warehouse	A subject-oriented information store specifically designed for decision support and analytical processing. A data warehouse contains large amounts of information used by managers, strategic marketers, and merchandisers to help them understand and predict market trends and make crucial business decisions. The information in a data warehouse is not usually modified by the people who use it. See also relational database .
database	See relational database .
database administration	The tasks involved in maintaining, designing, implementing changes to, tuning, and expanding a database. See also system administration .
database device	A device dedicated to the storage of the objects that make up databases.
database event	Any transaction that causes a state change in the database. Examples include: begin, rollback, or commit transaction boundary; an insert, update, or delete; a blob (Java object, image, or text); or a stored procedure invocation that results in a change in the database. The Replication Adapter and SQL Adapter detect database events.
database file	See table .
database integrity	A characteristic of a database evidenced by the database being both valid and complete. Database integrity consists of two complementary components: validity, which guarantees that all false information is excluded from the database, and completeness, which guarantees that all true information is included in the database. See also integrity constraints .
Database Management System (DBMS)	The Adaptive Server DBMS component performs query processing and transaction management for the current task. For example, the DBMS parses, compiles, and executes SQL statements and returns any results.
database object	One of the components of a database: table, column, view, index, procedure, trigger, default, or rule. See also object .
database object owner	A user who creates a database object .
Database Owner	The creator of a database. The Database Owner has control over all the database objects in that database. The login name for the Database Owner is “dbo.”
datachange	A measure of the number of updates, inserts, and/or deletes that occur on a column of a table.

data-only locking	A term used to denote both the datapages-locking scheme and the datarows-locking scheme . In these locking schemes, only the underlying data row or page is locked, and no index row locks or page locks are acquired. Index rows or pages are implicitly locked by locking the underlying data page.
datapage cluster ratio	The percentage of page accesses that do not require an extra extent I/O when rows are accessed in index order. The database cluster ratio measures the ordering and density of data pages in extents with respect to data access via the given index.
datapages-locking scheme	One of the three locking schemes in Adaptive Server . In datapages locking, only data pages are locked. No page locks are acquired on index pages. See also allpages-locking scheme and datarows-locking scheme .
datarow cluster ratio	The percentage of row accesses for an index that do not require an extra logical or physical I/O when rows are accessed in index order. The data row cluster ratio is maintained for each clustered index on a data-only-locked table and for each nonclustered index. The data row cluster ratio is not maintained for a clustered index because the ratio is always 1, that is, 100 percent clustered, since rows are guaranteed to be in key order.
datarows-locking scheme	One of the three locking schemes in Adaptive Server . In datarows locking, only data rows are locked. No row locks are acquired on index pages. See also allpages-locking scheme and datapages-locking scheme .
datatype	Specifies what kind of information each column will hold and how the data will be stored. Datatypes include char, int, money, and so on. Users can construct their own datatypes based on the Adaptive Server system datatypes.
datatype conversion function	A function that is used to convert expressions of one datatype into another datatype, whenever these conversions are not performed automatically by Adaptive Server .
datatype hierarchy	The hierarchy that determines the results of computations using values of different datatypes.
DataWindow object	See query object .
date function	A function that displays information about dates and times, or manipulates date or time values. The date functions include getdate, datename, datepart, datediff, and dateadd.
date part	Parts of a date, such as day, month, or year, recognized by the Transact-SQL date functions.

dbcc commands	Instructions to the database consistency checker (dbcc), which checks the logical and physical consistency of a database.
“dbo” account	See Database Owner .
DDL	See data definition language (DDL)
deadlock	A deadlock occurs when two or more user processes each have a lock on a separate page or table and each wants to acquire a lock on the other process’s page or table. The transaction with the least accumulated CPU time is killed and all of its work is rolled back.
DSS (decision support system)	A database system that supports queries involving large amounts of data. Commonly used to make business decisions. DSS queries typically access entire tables or large portions of tables, involve joins between many tables, and return summaries of large result sets. DSS applications are often run on a scheduled basis; for example, to produce the same report at the close of every business day. See also data warehouse .
default	The option chosen by the system when no other option is specified.
default clause	Specifies the default value for a column in the create table statement.
default database	The database that users automatically connect to when they log in.
default language	For a user, the language that displays that user’s prompts and messages, set with <code>sp_modifylogin</code> or the <code>language</code> option of the <code>set</code> command. For Adaptive Server , the language used to display prompts and messages for all users unless a user chooses a different language.
deferred update	An update operation that takes place in two steps. First, the log records for deleting existing entries and inserting new entries are written to the log, but only the delete changes to the data pages and indexes take place. In the second step, the log pages are rescanned, and the insert operations are performed on the data pages and indexes. See also direct update .
definition statement	See source text .
degree of parallelism	The number of worker processes that the optimizer chooses to execute the query in parallel. The degree of parallelism depends on both the upper limit of parallelism for the query and the level of parallelism suggested by the optimizer.
delete/insert direct update	A type of direct update operation. The row is deleted from its original location, and inserted at a new location.

deleted table	A pseudo-table containing rows with the deleted values of a delete statement, or the pre-updated values (before image) of an update statement.
delimited identifiers	Object names, enclosed in double quotes, which avoid certain restrictions on object names.
demand lock	Prevents additional shared locks from being set on a table or data page. Any new shared lock request has to wait for the demand lock request to finish.
dense frequency count	A frequency count in which values are contiguous in the domain. For example, in the integer domain, the set of values 1, 2, 3 is dense. Compare to sparse frequency count .
density	The average fraction of all the rows in an index that have the same key value. Density is 1 if all of the data values are the same and 1/N if every data value is unique.
dependency	A relationship between objects that occurs when one object refers to another, such as a stored procedure that refers to a table.
dependent	Data is logically dependent on other data when master data in one table must be kept synchronized with detail data in another table in order to protect the logical consistency of the database.
derived statistics	Statistics that are computed and used during search space operations, and whose lifetime is for the duration of the query optimization. These statistics include modified histograms, densities, column widths and table statistics that exist after all predicates (filtering or otherwise) are applied.
derived table	A table derived directly or indirectly from one or more tables by the evaluation of a query expression whose result is in the form of one or more columns. The values of a derived table are derived from the values of the underlying tables when the query expression is evaluated. A derived table is not described in system catalogs or stored on disk. See also SQL derived table and abstract plan derived table.
derived table expression	A query expression that defines a SQL derived table. A derived table expression consists of a select statement nested in the from clause of a select, insert, or create view statement.
detail	Data that logically depends on data in another table. For example, in the pubs2 database, the salesdetail table is a detail table. Each order in the sales table can have many corresponding entries in salesdetail. Each item in salesdetail is meaningless without a corresponding entry in the sales table.

deterministic property	A property of computed columns by which an expression always returns the same results from a specified set of inputs.
development tool	Software such as PowerBuilder that helps you build specialized graphical user interface applications for accessing Adaptive Server databases.
device	Any piece of disk or a file in the file system used to store databases and their objects. See also database device .
device I/O	The action of reading to or writing from a database device .
digital signature	Digital signatures are created with a mathematical algorithm that generates a unique, fixed-length string of numbers from a text message; the result is called a hash or message digest. Digital signatures are used for tamper detection and non-repudiation.
dimension table	A table in a “star” schema that acts as a primary key that can be combined with other dimension tables in the same “star” schema to form a composite key to access or join respective information from the central fact table.
direct update	An update operation that takes place in a single step, that is, the log records are written and the data and index pages are changed. Direct updates can be performed in three ways: in-place update , on-page update , and delete/insert direct update . See also deferred update .
directory information tree (DIT)	Entries in an LDAP server are called entities. Each entity has a distinguished name (DN) and is stored in a hierarchical tree structure. The tree structure is called the directory information tree.
dirty buffer grab	Occurs when a wash area is too small for the usage in a buffer pool, and operations that need a clean buffer may have to wait for I/O to complete on the dirty buffer at the LRU end of the pool or at the victim marker.
dirty pages	Pages that have been updated since they were last written.
dirty read	Occurs when one transaction modifies a row, and then a second transaction reads that row before the first transaction commits the change. If the first transaction rolls back the change, the information read by the second transaction becomes invalid.
discretionary access controls (DAC)	Restrict your access to objects based on your identity or your group membership. The controls are discretionary in the sense that a user with a certain access permission can pass that permission onto any other user (such as with the grant command). See also permission .

disk allocation pieces	Disk allocation pieces are the groups of allocation units from which Adaptive Server constructs a new database file. The minimum size for a disk allocation piece is one allocation page .
disk initialization	The process of preparing a database device or file for Adaptive Server use. After the device is initialized, it can be used for storing databases and database objects. The command used to initialize a database device is disk initialization .
disk mirroring	A duplicate of an Adaptive Server database device. All writes to the device being mirrored are copied to a separate physical device, making the second device an exact copy of the device being mirrored. If one of the devices fails, the other contains an up-to-date copy of all transactions. The command <code>disk mirror</code> starts the disk mirroring process.
display precision	The number of significant binary digits offered by the default display format for real and float values. Internally, real and float values are stored with a precision less than or equal to that of the platform-specific datatypes on which they are built. For display purposes, Sybase real values have 9 digits of precision; Sybase float values, 17.
distribution	An ordered set of $n+1$ values from a column of a table, obtained by sorting all the values and then selecting the first value and every subsequent $(\text{row_count}/\text{nth})$ value (where <code>row_count</code> is the number of rows in the table).
distribution page	A page associated with an index that was used to store statistics about data distribution in pre-11.9 versions of SQL Server and Adaptive Server .
DIT	See directory information tree (DIT) .
DLL	See dynamic link library (DLL) .
document type definition	(DTD) A list of element and attribute definitions unique to a class, in an XML document.
driver	A Sybase library that provides an interface to an external service provider.
DSS	Applications characterized by queries that process large volumes of data.
DTD	See document type definition
DTX participant	Distributed transaction participant (DTX participant) is an internal memory structure that Adaptive Server uses to coordinate a transaction that has been propagated to a subordinate Adaptive Server.

dump	<p>Making a backup of a database, including the data and the transaction log. This action may take two forms: a “database dump,” which includes all the database’s data and its transaction log (using the <code>dump database</code> command), or a “transaction log dump,” which includes the database’s transaction log but not its data (using the <code>dump transaction</code> command). A transaction log dump also removes records from the transaction log, making more log space.</p> <p>Also, the data that results from this action. See also load.</p>
dump device	<p>A single tape, partition, or file used for a database or transaction dump. A dump can span many devices, or many dumps can be made to a single tape volume.</p>
dump file	<p>The name of a dump file used to identify a specific backup on the backup media. The name cannot exceed 17 characters and must conform to operating system conventions. If you do not enter a name, Backup Server generates a default name based on:</p> <ul style="list-style-type: none">• The last seven characters of the database name• The two-digit year• The three-digit day of the year (1 through 366)• Hexadecimal-encoded time at which the dump file was created <p>For example, the file <code>catons93059E100</code> contains a copy of the publications database created on the fifty-ninth day of 1993.</p>
dump striping	<p>Allows you to use multiple dump devices for a single dump or load command.</p>
dump volume	<p>A single tape, partition, or file used for a database or transaction dump. A dump can span many volumes, or many dumps can be made to a single tape volume.</p>
dynamic configuration parameter	<p>An Adaptive Server configuration parameter that is updated immediately when you reset it; you do not have to restart Adaptive Server for it to take effect. See also static configuration parameter.</p>
dynamic dump	<p>A dump made while the database is active.</p>
dynamic index	<p>A worktable built by Adaptive Server for the resolution of queries using <code>or</code>. As each qualifying row is retrieved, its row ID is stored in the worktable. The worktable is sorted to remove duplicates, and the row IDs are joined back to the table to return the values.</p>
dynamic link library (DLL)	<p>Software used by Microsoft Windows and IBM OS/2 to provide services to applications.</p>

dynamic load distribution	(Cluster concept) The migration of an established client connection to a different instance in an attempt to balance the workload within a logical cluster.
dynamic partition elimination	Refers to the process of eliminating partitions of a table or index at run time. This is typically true when using unbound predicates .
E	
EAServer Adapter	An adapter that provides a client interface to EAServer. The EAServer Adapter enables you to execute Enterprise JavaBean methods as part of business event processing.
early deactivation	Allows a component's methods to specify when deactivation occurs. Early deactivation prevents a client application from tying up the resources that are associated with a component instance and allows the instance to serve more clients in a given time frame.
encryption	The process wherein a cryptographic algorithm is used to encode information to safeguard it from anyone except the intended recipient.
engine	An instance of the Adaptive Server executable that can communicate with other Adaptive Server engines in shared memory. An Adaptive Server running on a uniprocessor machine always has one engine, engine 0. An Adaptive Server running on a multiprocessor machine can have one or more engines. Also called server engine .
engine group	A set of one or more engines.
Enterprise Application Server	EAServer (also known as Jaguar). An integrated development and deployment environment with a scalable, secure, and transaction-aware platform for Web and component-based applications. Sybase Enterprise Event Broker (SEEB) runs as a service of EAServer.
enterprise data	Data that exists anywhere in a networked system. Enterprise data can be stored on Sybase servers or as heterogeneous data.
entity	A database or a database object that can be identified by a unique ID and that is backed by database pages. Examples of entities are the database pubs2, the log for database pubs2, the clustered index for table titles in database pubs2, and the table authors in database pubs2. Identifying entities is the first step in the logical design of a database.
equijoin	A join based on equality.

equipartitioned	Two tables having compatible partitioning keys and partitioning criteria. If two tables have the same number of partition keys with compatible datatypes, and the partition criteria such as intervals are the same for the range partitions, the two tables are considered equipartitioned.
error handling	Techniques available to Transact-SQL programmers on which to base code and display errors and error messages.
error log	A file that stores severe error messages and the results of the start-up and recovery of databases.
error message	A message issued by Adaptive Server , usually to the user's terminal, when Adaptive Server detects an error condition.
error state number	The number attached to an Adaptive Server error message that allows unique identification of the line of Adaptive Server code at which the error was raised.
evaluated configuration	The configuration of SQL Server that was evaluated at the C2 security level by the NSA (National Security Agency) in 1996 on the HP 9000 HP-UX BLS, 9.09+ platform. Certain features of SQL Server, such as remote procedures and direct updates to system tables, were excluded from the evaluated configuration. For a complete list of features excluded from the evaluated configuration, see Appendix A in the SQL Server Installation and Configuration Guide for HP 9000 HP-UX BLS, 9.09+.
event	Within Sybase Enterprise Event Broker (SEEB) , an object that is imported, passed between processors, and exported to an external database. Outside SEEB, an event can be a message or a database operation.
event broker	An application that routes, transforms, and delivers events to a bus, queue, or comparable destination.
event content	The combination of an event schema and an event container.
event processor	The Sybase Enterprise Event Broker (SEEB) component that directs an event through a route, and processes events along the route. A route comprises a sequence of two or more event processors. The event processor acts upon the event according to event processor logic. SEEB event processors include inflow event processors, outflow event processors, transfer processors, branch processors, and transform processors.
exchange operator	An operator, which when applied to a data stream, can change the degree or data semantics of the stream. Exchange operators take part in repartitioning. An exchange operator has both a producer side and a consumer side. The producer tasks run the relational operator clone below the exchange, while the the consumer side runs in as many clones as the consumer operator requires.

exclusive lock	A type of lock acquired on an object during a write operation. It does not allow other transactions to acquire exclusive, updated, or shared locks on the object. Exclusive locks can be obtained on a table or page.
exclusive row lock	A type of lock acquired on a data row during a write operation. It does not allow other transactions to acquire exclusive, update, or shared locks on the row.
execute cursor	A cursor that is a subset of client cursors whose result set is defined by a stored procedure that has a single select statement. The stored procedure can use parameters. The values of the parameters are sent through Open Client calls.
execution class	A classification of applications, logins, and stored procedures based on the performance requirements of these entities. An execution class is defined by a set of execution attributes associated with the objects belonging to that class.
existence join	A type of join performed in place of certain subqueries. Instead of the usual nested iteration through a table that returns all matching values, an existence join returns TRUE when it finds the first value and stops processing. If no matching value is found, it returns FALSE.
existence scan	A scan algorithm based on stopping the scan of the table as soon as the first row is fully qualified. Existence scans are typically introduced by tables from a flattened exists subquery
expected row size	A user-defined average for the expected size of a data row. Specifying the expected row size for a table reserves space on the data page for rows that grow in length after they are inserted. This helps to avoid row forwarding .
expression	A computation, column data, a built-in function, or a subquery that returns values.
expression subquery	A type of subquery that is introduced with an unmodified comparison operator, must return a single value, and can be used almost anywhere an expression is allowed in SQL.
extended high-availability (HA) failover	Adaptive Server provides a list of failover addresses to high-availability-aware clients when they connect. This allows multiple clients to fail over eliminates the need for the “HAFILOVER” entry in the directory services or interfaces file.
extended stored procedure	A mechanism for calling external procedural language functions from within Adaptive Server . Users invoke an extended stored procedure using the same syntax as they use for a regular stored procedure, but instead of executing Transact-SQL statements, the ESP executes procedural language code, which is compiled in a dynamic link library (DLL) or shared library.

extent	Whenever a table or index requires space, Adaptive Server allocates a block of eight pages, called an extent, to the object. The size of the extent is determined by which logical page size your server is using, 2, 4, 8, or 16K.
extent allocation	A method of allocating new space for a table or an index. In most cases, new space is allocated one page at a time. However, commands that need to allocate large amounts of space can perform extent allocation, allocating eight pages at a time.
extent stealing	When partitions have run out of space on the device, and extents have been allocated on another physical device. Extent stealing can take place when data is being inserted into the table with insert commands or bulk copy and while clustered indexes are being created.
external login	An alternate login name and password used when a Component Integration Services client logs in to a remote server. It is created using <code>sp_addexternlogin</code> .
F	
fact table	The main table of a “star” database schema that has a composite key composed of attributes that are foreign keys of several dimension tables.
failback	The planned event during which Adaptive Server is migrated back to, and restarted on, the original machine. This involves moving the failed-over databases, devices, and client connections from the secondary companion to the restarted primary companion.
failover	During failover, Adaptive Server migrates to another machine which takes over the responsibility of managing the failed over Adaptive Server. Failover can occur because of either a scheduled maintenance or a failure of Adaptive Server or the machine running Adaptive Server.
failover group	(Cluster concept) A set of failover instances defined for a logical cluster. Failover groups let you specify preference and order for failover instances.
failover instance	(Cluster concept) An instance on which a logical cluster can run if one or more of its base instances fail.
failover mode	The mode of the primary companion after it has failed over and is running on the secondary companion.
family	The coordinating process requests four worker processes from the pool of worker processes. The coordinating process together with the worker process is called a family.

fatal errors	Errors with severity levels of 19 and above. They terminate the user's work session, and requires that he or she log in again.
Federal Information Processing Standards (FIPS) flagger	An Adaptive Server option activated with the set command. When the FIPS flagger is active, Adaptive Server returns a warning message when you use Transact-SQL extensions to entry-level SQL92. FIPS recognizes SQL89 as the base standard.
fetch	A fetch retrieves one or more rows and changes the current cursor position in the cursor result set. Also called a cursor fetch
field	See column .
filter	An operation that allows you to specify which objects to hide or display in a window or dialog.
fixed row ID	A method of ensuring that the row ID (RID) associated with a row does not change. The offset table of the row may change because of row movement within the page, but the row number itself does not change. If row expansion requires a row to move, the row migrates to a new page and the forwarding address is stored in the original location of the row. The row ID does not change for the life of the row, except when clustered indexes are created or when the reorg rebuild command is run. Row IDs are reused after a row is physically deleted from a page.
for load	Specifies that a database will be created for restoration from tape.
for trigger	The trigger currently supported by Adaptive Server. These triggers fire after an insert/update/delete statement on a table, and are sometimes called after triggers.
forcing	Including options in a query or session to override the optimizer's choice for a query plan.
foreign key	A key column in a table that logically depends on a primary key column in another table. Also, a column (or combination of columns) whose values are required to match a primary key in some other table.
form	A form is an application window on which you place buttons, text boxes, and other elements of the user interface.
format file	A file created while using bcp to copy data out from a table in an Adaptive Server database to an operating system file. The format file contains information on how the data being copied out is formatted and can be used to copy the data back into an Adaptive Server table or to perform additional copy outs.

forwarded row	A row in a data-only-locked table that has been migrated to a new data page, because of a change in the length of the row. The previous location of the row stores a pointer to the forwarded row's location.
forwarding address	The new location of a row. The forwarding address is in the original location of the row.
fragment	When you allocate only a portion of the space on a device with <code>create</code> or <code>alter database</code> , that portion is called a fragment.
freelock list	On a multiengine Adaptive Server , each engine has its own freelock list (a list of locks available to satisfy lock requests). If the engine freelock list runs out of locks, Adaptive Server moves locks from the global freelock list to the engine freelock list.
free-space threshold	A user-specified threshold that specifies the amount of space on a segment, and the action to be taken when the amount of space available on that segment is less than the specified space.
frequency cell	A cell in a histogram that represents a frequency count . The weight of a frequency cell indicates the percentage of columns that match the value for the cell. Compare to range cell .
frequency count	A method of modeling a data distribution with a low number of values in the domain. For example, a gender column can be modeled by two values, "M" and "F", which have a count of about 0.48 and 0.52, respectively.
functions	See built-in functions .
function-based index	An index that contains one or more expressions as index keys.
G	
gateway	Intermediate software that provides language translation, datatype conversion, and protocol conversion between the client and the data source.
generalized index key	An index key generated by invoking a user-defined generalized function.
generic column	Usually, a column of a table referenced in a query , but also an abstraction that includes expressions that can be used in an expression join.
generic table	Usually, a table referenced in a query , but also an abstraction to represent any object that is permuted in the join order by the optimizer. For example, a subquery is modeled as a generic table

global index	An index on a partitioned table. A global index results when an index and the table have different partitioning strategies such that index leaf pages of global indexes point to more than one partition.
global statistics	Statistics that apply to all data values of a table .
global variable	A system-defined variable that Adaptive Server updates on an ongoing basis. For example, @@error contains the last error number generated by the system. See also local variable .
greedy search strategy	Any optimizer permutation strategy to obtain a query plan very quickly. Likely to be a suboptimal strategy, because very coarse criteria is used, rather than looking at all query plans
group	A uniquely named set of users assigned a set of permissions for the objects and operations within a database.
grouped aggregate	See vector aggregate .
guest	A user name in the sysusers table of the model database, which enables a user with a valid Adaptive Server login to use databases created from model, with limited privileges.
H	
Halloween problem	An anomaly associated with cursor updates, whereby a row seems to appear twice in the result set. This happens when the index key is updated by the client and the updated index row moves farther down in the result set.
hard fault	A persistent corruption of Adaptive Server found during a dbcc checkstorage operation. Hard faults cannot be corrected by restarting Adaptive Server. See also soft fault .
hash buckets	A hash table is composed of hash buckets, each of which is selected by a hash key .
hash key	A value derived by a hash function on the object being managed.
hash partitioning	Partitioning of tables or indexes according to values in the partition keys . The partitioning key or keys processed by a system-defined hashing function to determine the partition to which each row belongs.
hash table	Used to support hashing algorithms, which quickly access objects based on identical hash keys .
hash-based aggregation	A strategy for evaluating group by aggregates in which the group is looked up by a hash key on the grouping columns.

heap table	A table without a clustered index , where data is stored as a heap. The data rows are not stored in any particular order, and new data is always inserted on the last page of a page chain. A typical heap table is an unpartitioned table without a clustered index, but a partitioned table without a clustered index is also classified as a heap table.
heterogeneous data	Any data from non-Sybase data sources such as Oracle, Informix, or DB2.
heterogeneous server	A remote server that is not an Adaptive Server installation, for example, an Informix or DB2 database.
heuristic-based pruning	Optimization techniques where portions of search space (for example, tree shapes and permutations) are skipped based on a set of predetermined rules applicable to a query
hexadecimal string	A hexadecimal-encoded binary string that begins with the prefix 0x and can include the digits 0 through 9, and the uppercase and lowercase letters A through F. The interpretation of hexadecimal strings is platform specific. For some systems, the first byte after the prefix is the most significant; for others, the last byte. For example, the string 0x0100 is interpreted as 1 on some systems and as 256 on others.
high availability	A system that is designed to reduce the amount of system downtime.
histogram	A set of cells in which each cell has a weight. Each cell has an upper limit, a lower bound, and a float value between 0 and 1 representing the percentage of rows within the range. Adaptive Server statistics for column values are stored as histograms in the sysstatistics system table. These statistics are used by the query optimizer.
histogram tuning factor	A factor used, only when frequency cells exist, to increase the number of steps in a histogram over the default or specified step count. For example, a factor of 3 can increase a default step count of 20 to 60 if frequency cells exist in the distribution.
histogram weight array	An array of float values associated with a histogram which gives either the percentage of the table selected by that cell (for table-normalized histograms), or the percentage of a cell selected by a particular predicate (for cell-normalized histograms).
home page	The original location of a row. If the row has been moved by row forwarding , the home page contains the forwarding address of the row.

horizontal parallelism	Partitioned parallelism and independent parallelism are classified into horizontal parallelism . The ability to run multiple instances of operators on different data sets located across different storage units is also called horizontal or partitioned parallelism .
hysteresis	A value used to control the spacing of thresholds on a segment and to prevent the stored procedure associated with a threshold from being triggered too frequently.
I	
identifier	A string of characters used to identify a database object , such as a table name or column name.
IDENTITY column	A column that contains system-generated values that uniquely identify each row in a table. IDENTITY columns store unique numbers, such as invoice numbers or employee numbers, that are generated automatically by Adaptive Server . The value of the IDENTITY column uniquely identifies each row in a table.
image function	See text and image function .
implicit conversions	Datatype conversions that Adaptive Server automatically performs to compare datatypes .
independent parallelism	See bushy parallelism .
index	A database object that consists of key values from the data tables and pointers to the pages that contain those values. Indexes speed up access to data rows by pointing Adaptive Server to the location of a table column's data on disk.
index covering	A data access condition where the leaf-level pages of a nonclustered index contain the data needed to satisfy a query. The index must contain all columns in the select list as well as the columns in the query clauses, if any. The server can satisfy the query using only the leaf level of the index. When an index covers a query, the server does not access the data pages.
index intersection	An access path in which several RIDs from two or more indexes of a table are joined to obtain the set of RIDs that qualify the result set for the scan of the table with several SARGs that are joined together using "AND."
index page cluster ratio	The percentage of index leaf page accesses via the page chain that do not require extra extent I/O. The index page cluster ratio is maintained for clustered indexes on data-only-locked tables and for nonclustered indexes .

index selectivity	The ratio of duplicate key values in an index . An index is selective when it lets the optimizer pinpoint a single row, such as a search for a unique key. An index on non unique entries is less selective. An index on values such as “M” or “F” (for male or female) is extremely nonselective.
index union	An access path in which several row ID (RID) s from two or more indexes of a table are joined with duplicate removal to obtain a set of RIDs that qualify the result set for the scan of the table with several SARGs that are joined together using “OR.”
initial response time	The time required to return the first result row of a query to a user. For some queries, initial response time can be very brief, even though the time to return the full result set can be much longer.
initialization	See disk initialization .
in-memory database	A database stored entirely in a named cache without using disk storage for data or logs. In-memory databases have a durability level of <code>no_recovery</code> (the database is re-created when you restart Adaptive Server). In-memory databases take advantage of residing entirely in memory and their low durability to offer improved performance over disk-resident databases, and can manage high transaction throughput.
inner join	Joins that produce a result set that includes only the rows of the joining tables that meet the restriction. Rows that do not meet the join restriction are not included in the joined table.
inner query	See subquery .
inner table	Describes the placement of a table in an outer join . For example, in a left join, the inner table is the right table.
in-order join	A join operation where some (or all) of the joining attributes from the outer join are ordered, as would occur from a sort or index scan.
in-place update	A type of direct update operation. An in-place update does not cause data rows to move on the data page. Compare to on-page update and delete/insert direct update .
instance	(Cluster concept) An Adaptive Server that participates in a shared-disk cluster.
instance number	(Cluster concept) A number that uniquely identifies a named instance in the Adaptive Server shared-disk cluster.

instance state	(Cluster concept) The state of an instance in a logical cluster as it is perceived by a logical cluster. Thus, an instance can be physically online, but offline to a given logical cluster.
insert table	A pseudo-table containing rows with the inserted values of an insert statement, or the updated values (after image) of an update statement.
instance pooling	Allows a single component instance to service multiple clients. Instance pooling eliminates resource drain from repeated allocation of component instances
instead of trigger	Trigger that can be defined on a view, replacing the standard action of the update/insert/delete statement.
int	A signed 32-bit integer value.
integrity constraints	A model to describe the database integrity in the create table statement.
integrity rules	Rules that describe how data will be kept accurate and consistent in the relational model. See also database integrity .
intent lock	A type of table-level lock. A shared intent lock indicates that there is a shared lock on a page or row of that table . It prevents other transactions from acquiring an exclusive table-level lock. An exclusive intent lock indicates that there is an exclusive or update lock on a page or row of that table. It prevents other transactions from acquiring an exclusive or shared table-level lock.
interfaces file	<p>The interfaces file is an operating system file that must be available on each machine from which connections to Adaptive Server are made. By default, this file is located in the directory that is specified in the SYBASE environment variable. The name of the file varies, but is usually <i>interfaces</i> on a UNIX system and <i>sql.ini</i> on a PC.</p> <p>Each entry in the interfaces file tells the host machine how to connect to Adaptive Server. An interfaces file entry contains the name of an Adaptive Server installation and a list of services provided by the server.</p>
internationalization	The process of enabling an application to support multiple languages and cultural conventions.
I/O	See device I/O .

isolation level	Specifies the kinds of actions that are not permitted while the current transactions execute. The ANSI standard defines four levels of isolation for SQL transactions. Level 0 prevents other transactions from changing. The user controls the isolation level with the set option <code>transaction isolation level</code> or with the <code>isolation</code> clause of <code>select</code> or <code>readtext</code> . Level 3 is equivalent to doing all queries with <code>holdlock</code> . The default is level 1. Also called “locking level.”
isql	A command-line interface to Adaptive Server .
iterator	An execution engine operator. Query results are encapsulated using iterators that are self-contained software objects that accept a stream of rows from the data manager (such as an index or table scan), or other iterators, and subsequently can produce a stream of rows that can be requested by the client or other iterators. The role of an iterator is to process many iterations of a data set across many nodes in serial or parallel. Iterators do not know the source of the data stream if it is external (another iterator); iterators do know the source if it is internal (the iterator itself produces it). For each iteration of a data set, the iterator applies a predefined behavior to the data set being processed, manipulating the data according to the specification of that iterator. For example, scanning rows from a table on disk can be the behavior of one type of iterator. A key feature of iterators is that regardless of what type the iterator is and what behavior is associated with it, all iterators follow the same mode and have the same external interface. They all open data streams, iteratively read the streams, process the data and close the streams.
J	
Java application	In Java terminology, an application is a Java program that does not require a host server or browser to run. The user sees windows and menus and interacts with controls. The application can connect with a middle-tier or database server.
job	A series of one or more actions performed on a database in a single operation, such as: backups, update statistics, dump database. Used specifically with Job Scheduler.
join	A basic operation in a relational system that links the rows in two or more tables by comparing the values in specified columns .
join density	See total density .
join histogram	An intermediate histogram created during optimization only and then discarded. A join histogram is the result of taking the histograms of two columns that are equijoin and producing a histogram that models the data distribution after the join has occurred

join selectivity	An estimate of the number of rows from a particular table that will join with a row from another table. If histograms are available on the joining columns then the selectivity is based on the join histogram. If no histograms are available then the selectivity is $1/N$, where N is the number of rows in the smaller side of the join.
joined table	The result of the join between the left and right tables.
K	
kernel	A module within Adaptive Server that acts as the interface between Adaptive Server and the operating system. The kernel manages tasks associated with Adaptive Server's clients. For example, it tracks the state of the task, which is necessary in a multithreaded environment because of context switching.
key	A field used to identify a record, often used as the index field for a table.
key value	Any value that is indexed.
keyword	A word or phrase that is reserved for exclusive use by Transact-SQL. Also known as a reserved word .
keyword phrase	See clause .
L	
language cursor	A cursor declared in SQL without using Open Client. As with Adaptive Server cursors, Open Client is completely unaware of the cursors and the results are sent back to the client in the same format as a normal select.
language group	The language or languages that are covered by a character set. A language group can contain many languages or only one language.
Language Module	A set of files, including localization files, that provide alternate language, sort order, and character sets for Adaptive Server .
large I/O	An I/O that reads or writes multiple data pages.
last-chance threshold	A default threshold in Adaptive Server that suspends or kills user processes if the transaction log has run out of room. This threshold leaves just enough space for the deallocation records for the log itself. The last-chance threshold calls a procedure named <code>sp_thresholdaction</code> . This procedure is not supplied by Sybase. It must be written by the System Administrator.
latch	A lightweight, nontransactional synchronization mechanism held for a very short duration to ensure the physical consistency of the page.

lava operator	A self-contained software object that implements a basic operation and may be chosen by the optimizer as part of a lava query plan. Some examples of lava operators include the ScanOp that reads rows from database tables , the MergeJoinOp that implements the merge join and the InsertOp that inserts rows into tables. There are 32 lava operators. See iterator .
lava query execution engine	The Adaptive Server module that executes the lava query plan that was chosen by the optimizer. Query plans are executed by calling methods of the top operator in the plan (the RootOp), which calls methods on its child operators, which in turn call methods on their child operators, down to the leaf operators if necessary, to generate a result row. Result rows are generated at the leaf operator nodes and passed up the operator tree to the RootOp, which consumes them, in other words, sends them to the client, or assigns values to variables, and so on
lava query plan	An “upside- down” tree of lava operators: The top operator can have one or more child operators, which in turn can have one or more child operators, and so on, thus building the upside-down tree of operators. The optimizer chooses the shape of the tree and the operators in a lava query plan, and the plan is executed by the lava query execution engine.
LDAP (Lightweight Directory Access Protocol)	An industry standard for accessing directory services, which allow components (applications, other servers, and so on) to look up information by a distinguished name (DN) from an LDAP server that stores and manages server, user, and software information used throughout the enterprise or over a network. LDAP defines the communication protocol and the contents of messages exchanged between clients and servers.
leaf level	The level of an index at which all key values appear in order. For Adaptive Server clustered indexes, the leaf level and the data level are the same. For nonclustered indexes, the last index level above the data level is the leaf level, since key values for all of the data rows appear there in sorted order.
left deep tree plan	This is an alternative “tree-based” description of a “join order” of a set of tables. It is a tree shape of query plan structure where right nodes are always leaf nodes. The order of tables in this tree shape can be influenced by the set forceplan option.
legacy data	Data from older, perhaps out-of-date, data sources or from data sources that are no longer supported by current standards.
libraries	See application program interface (API) .
lightweight directory access protocol	See LDAP (Lightweight Directory Access Protocol) .

lightweight procedure	A procedure that can be created and invoked internally by Adaptive Server . A lightweight procedure has no system table entries and consists of a plan saved in the procedure cache and some memory data structures that identify it.
lightweight process	See worker process .
LIO	See logical I/O cost .
list partitioning	Partitioning of tables or indexes according to values in the partitioning keys. The partitioning key or keys are compared with sets of user-supplied values to determine the partition to which each row belongs.
livelock	A request for an exclusive lock that is repeatedly denied because a series of overlapping shared locks keeps interfering. Adaptive Server detects the situation after three denials and grants a demand lock to the update transaction, queueing further shared lock requests after the demand lock.
load	To restore data stored in a backup created during a dump .
load profile	(Cluster concept) A set of weighted metrics used to determine the relative workload on an instance in a logical cluster. You can create your own load profiles or use one of the profiles provided by Sybase.
load score	(Cluster concept) A computed value of the overall load on an instance; a unitless number that can be used to compare relative workloads on different instances in a logical cluster, or on the same instance at different times.
LOB	Large objects. LOB refers to the text, image, and untext datatypes. LOB values can be very large and require special considerations when they are stored, in the database or in memory.
local index	When a table's index is partitioned the same way as its data.
local server	The server or node where a query originates.
local statistics	Statistics that apply to data values of a specific partition of a partitioned table.
local system temporary database	(Cluster concept) Space for temporary tables and worktables. Each instance in the cluster has a local system temporary database that it alone can access.
local variable	A user-defined variable defined with a declare statement. See also global variable .
localization	The process of adapting an internationalized product to meet the requirements of one particular language or region, for example Spanish, including providing translated system messages; translations for the user interface; and the correct formats for date, time, and currency.

localization files	Files that contain translated error messages, character set and sort order definitions, and utilities for converting Adaptive Server 's character set into the appropriate character set for a particular terminal.
location transparency	A feature unique to Component Integration Services that makes remote data appear as if it were local. Users do not need to know where the data resides to access it.
lock	A concurrency control mechanism that protects the integrity of data and transaction results in a multiuser environment. Adaptive Server applies page or table locks to prevent two users from attempting to change the same data at the same time, and to prevent processes that are selecting data from reading data that is in the process of being changed.
locking	The process of restricting access to resources in a multiuser environment to maintain transactional isolation and prevent concurrent access problems. Adaptive Server automatically applies locks to tables, pages, or rows.
locking level	See isolation level .
lock promotion	Promotion of the row or page locks acquired by a scan to a table lock. Lock promotion takes place when a scan has exceeded the user-configurable setting for the lock promotion threshold. Locks are promoted from row locks to table locks, or page locks to table locks.
lock promotion threshold	The number of page locks allowed in a table before Adaptive Server attempts to issue a table lock. If the table lock is successful, Adaptive Server releases the page locks.
logical cluster	(Cluster concept) A method of abstracting the physical cluster so that multiple application services can be established. A logical cluster supports fine-tuned management of the workload within the cluster by enabling application- or user-specific service level agreements, resource assignments, and failover rules. Applications connect directly to a logical cluster.
logical design	A design in which you define the tables, relations, and keys of a relational database. See also physical design .
logical expression	An expression that evaluates to TRUE (1), FALSE (0) or UNKNOWN (NULL). Logical expressions are often used in control of flow statements, such as if or while conditions.
logical I/O cost	Optimizer's estimate of the number of logical reads.

logical key	The primary, foreign, or common key definitions in a database design that define the relationship between tables in the database. The logical key is not necessarily the same as the physical key (the key used to create index) on the table.
logical operator	The operators and, or, and not. All three can be used in where clauses. The operator and joins two or more conditions and returns results when all of the conditions are true; or connects two or more conditions and returns results when any of the conditions is true. In the context of optimization, a logical operator is an operation in query processing that does not specify a specific algorithm, for example, join, scan, or sort.
logical page size	Specified when you build a new master device. All databases, and all objects in every database, use the same logical page size. The size of Adaptive Server 's logical pages (2, 4, 8, or 16K) determines the server's space allocation.
logical partitioning	Any partitioning of rows into 2 or more partitions, which, in Adaptive Server means using either semantic partitioning, for which a license is required; or round-robin partitioning for which no license is required.
logical property	A property that is common to a set of subplans associated with a set of tables (equivalence class). An example is the row count, since no matter how the set of tables are joined, the same row should exist after the same predicates are applied.
logical read	The process of accessing a data or index page already in memory to satisfy a query. See also physical read .
login redirection	(Cluster concept) The mechanism by which an instance can direct an incoming client connection to a different instance in the cluster. Login redirection is used to route inbound connections to instances in a logical cluster and for load balancing.
login	The name a user uses to log in to Adaptive Server . A login is valid if Adaptive Server has an entry for that user in the system table syslogins.
look-ahead set	The set of pages for a particular operation to be fetched by asynchronous prefetch. Each Adaptive Server operation that uses asynchronous prefetch builds a look-ahead set based on the known or expected set of pages that will be needed by the operation in the near future.
LRU (least recently used) buffer reuse strategy	A caching strategy for replacing the least-recently used buffers in the data cache. A clean data page is taken from the LRU end of the data cache to store a page read from disk. The new page is placed on the data cache's page chain at the most recently used (MRU) end of the cache, so that it stays in memory.

magic values Values that are substituted or used as defaults when actual values are not known. For example, the optimizer uses magic numbers for the selectivity of a value in a parameter that cannot be determined when the query is optimized.

M

main form The first window displayed by the application. When you create an applet or application target, a main form is created for you.

mantissa The mantissa is the decimal part of a number, for example .654 is the mantissa of 1.654.

MAPI See **Messaging Application Programming Interface (MAPI)**.

master database The system database that controls the user databases and the operation of **Adaptive Server** as a whole. Known as master, it keeps track of such things as user accounts, ongoing processes, and system error messages.

master table A table that contains data on which data in another table logically depends. For example, in the pubs2 database, the sales table is a master table. The salesdetail table holds detail data that depends on the master data in sales. The detail table usually has a foreign key that joins to the primary key of the master table.

master-detail relationship A relationship between sets of data where one set of data logically depends on the other. For example, in the pubs2 database, the sales table and salesdetail table have a master-detail relationship. See **detail** and **master table**.

matching index scan A scan using a nonclustered index when the query has a where clause (search argument) on a set of columns, and the columns form a **prefix subset** of keys on the index. The index is used to position the search at the first matching key, and then scanned forward for additional matches on the specified index key columns. The scan stops at the first row that does not match. Matching index scans are quite fast and efficient. Compare to **nonmatching index scan**.

materialized computed column A computed column whose result is pre-evaluated and stored in the specified table or index page.

mathematical function A function that returns values commonly needed for operations on mathematical data.

media failure A media failure occurs when the information on a medium (typically a hard disk drive) becomes unusable.

MenuBar object In building a **Java application** with **PowerJ**, a MenuBar object represents all the menus displayed by the form. MenuBar objects can be added only to forms that are based on the Frame class.

message number	The number that uniquely identifies an error message.
messaging	Messaging refers to the process of sending events, or strings of bytes, from a source to a destination over a communications pathway.
Messaging Application Programming Interface (MAPI)	An e-mail application programming interface developed by Microsoft. Adaptive Server integrates with MAPI to provide an interface for sending and receiving data by e-mail rather than through the traditional client/server connection-based facility.
metadata	Data about data. Metadata is stored in local proxy tables by Component Integration Services. The metadata stored by Component Integration Services represents schemas with information about remote tables.
metadata cache	A reserved area of memory used for tracking information on indexes, objects, or databases. You can configure the size of the metadata caches based on the number of metadata descriptions used by indexes, objects, or databases.
metadata description	A memory data structure that represents the state of an index, an object, or a database while it is in use or cached between uses.
MIME	Multipurpose Internet Mail Extensions (MIME).
minor column	A non-leading column of a composite index. For an index on (A, B, C), B and C are minor columns.
mirror	See disk mirroring .
mirror device	A duplicate Adaptive Server database device. All writes to the primary device are copied (mirrored) to a second physical device. Writes can be either serial (consecutive) or parallel (simultaneous). If one device fails, the other contains an up-to-date copy of all transactions.
mixed workload	Relational queries are broadly classified into simple transactional queries found in OLTP environments and complex queries found in DSS environments. In a production environment, database systems are configured to run either transactional or complex queries at the same time or at different times. Installations that support both of the above are referred to as “mixed-workload” systems. Since it is not always possible to predict the type of workload, it is important to support both OLTP and DSS queries in the same configuration of the data processing system to efficiently support workloads of all types
model database	A template for new user databases. The installation process creates model when Adaptive Server is installed. Each time create database is issued, Adaptive Server makes a copy of model and extends it to the size requested, if necessary.

modulo	An arithmetic operator represented by the percent (%) sign that gives the integer remainder after a division operation on two integers. For example, $21 \% 9 = 3$ because 21 divided by 9 equals 2 with a remainder of 3.
most recently used (MRU) replacement strategy	A caching strategy for table scans and nonclustered index scans. The optimizer chooses this strategy when it determines that the pages need to be accessed only once for a particular query. Instead of adding all of the pages to the MRU/LRU chain, the pages are immediately flushed as soon as the query finishes examining them, and the next page for the query is read into the buffer.
multibyte character set	A character set that includes characters encoded using more than one byte. EUC JIS and Shift-JIS are examples of character sets that include several types of characters represented by multiple bytes in a Japanese language environment.
multi-database transaction	When a transaction accesses, uses, and/or modifies data from multiple databases.
multiprocessing	Multiple processes that share memory or use some type of synchronized method for passing messages between them are performing multiprocessing.
Multipurpose Internet Mail Extensions (MIME)	Multipurpose Internet Mail Extensions. A format for exchanging complex messages through the Internet. Web browsers and servers use MIME types to describe the messages transmitted from the server to the browser.
multi-way joins	Join queries in which some tables join to two or more tables resulting in star join and snowflake join configurations.
N	
named cache	See buffer cache .
national character sets	See native character sets .
native character sets	Also referred to as national character sets. Character sets that are platform-specific and support a subset of languages, for example, the Western European languages.
native formats	Operating system-specific formats created using the -n option. Native formats usually create a more compact operating system file.
natural join	A join in which the values of the columns being joined are compared on the basis of equality, and all the columns in the tables are included in the results, except that only one of each pair of joined columns is included.
nested queries	select statements that contain one or more subqueries.

nested select statements	See nested queries .
network affinity migration	The process of moving network I/O from one engine to another. SMP systems that support this migration allow Adaptive Server to distribute the network I/O load among all of its engines.
next-key locking	A strategy to protect repeatable read transactions from reading phantom rows .
node	A machine in a high availability system. For clusters, A machine (hardware) that hosts an Adaptive Server instance.
nonclustered index	An index that stores key values and pointers to data. The leaf level points to data pages rather than containing the data itself.
noncorrelated subquery	A noncorrelated subquery can be evaluated as if it were an independent query. Noncorrelated subqueries can alternatively be stated as joins and are processed as joins by Adaptive Server .
nonmatching index scan	A scan using a nonclustered index when the search arguments do not form a prefix subset of the index key columns, although they match some parts of the composite key . The scan is performed using the index from the lowest key value to the highest key value, searching for the matches specified in the query. This type of scan is performed on nonclustered indexes when all columns for a table referenced in the query are included in the index. Although cheaper than a table scan, a nonmatching index scan is more expensive than a matching index scan .
nonrepeatable reads	Such reads occur when one transaction reads a row and a second transaction modifies that row. If the second transaction commits its change, subsequent reads by the first transaction yield different results than the original read.
non-repudiation	The inability to deny that a document is valid. In regards to Internet communications, digital signatures provide non-repudiation; the sender cannot deny that he sent the message.
normal companion mode	The mode during which two Adaptive Servers in a high availability system are functioning as independent servers and are configured to fail over during a scheduled maintenance or system failure.
normalization rules	The standard rules of database design in a relational database management system.
not-equal join	A join on the basis of inequality.

null	A value given to a parameter or option that has no explicitly assigned value. NULL is not equivalent to zero, or to blank. A value of NULL is not considered to be greater than, less than, or equal to any other value, including another value of NULL.
numeric expression	See arithmetic expression .
O	
OAM	See object allocation map (OAM) .
object	A passive entity that contains or receives information, and that cannot change the information it contains. In Adaptive Server , objects include rows, tables, databases, stored procedures, triggers, defaults, and views. See also database object .
object access permission	Regulates the use of certain commands that access certain database objects. Object access permissions are granted and revoked by object owners.
object allocation map (OAM)	One or more pages within a table or an index on a table that contain pointers to the allocation pages for each allocation unit where the object uses space.
object allocation map (OAM) pages	Each table, and each index on a table has an OAM (object allocation map) page. The OAM page holds information on pages allocated to the table or index and is checked when a new page is needed for the index or table.
object allocation map (OAM) scan	A method of accessing data by first reading the OAM pages for a table, then the allocation pages, and finally the data pages where the data is stored.
object creation permission	Regulates the use of commands that create objects. These permissions can be granted only by a System Administrator or Database Owner.
object owner	User who owns an object either by having created the object or by being given ownership. For example, a System Administrator can designate a user as a Database Owner, a Database Owner can designate a user as a table owner or give a user permission to create a table.
object permission	A permission that regulates the use of certain commands (data modification commands, plus select , truncate table and execute commands) to specific tables, views or columns. See also command permission .
ODBC	See open database connectivity (ODBC) .
offset table	A set of fields at the end of data and index pages that store pointers to the byte location of the rows on the page. Indexes on allpages-locked tables do not have offset tables; all other indexes and all data pages do.

OLTP	See online transaction processing (OLTP) .
online transaction processing (OLTP)	Characterized by many short transactions containing queries that have minimal resource usage.
on-page update	A type of direct update operation, performed when the length of the data row changes. The changed data row remains on the same data page, but other rows on the page may move. Contrast to in-place update and delete/insert direct update .
open database connectivity (ODBC)	The ODBC interface, defined by Microsoft Corporation, is a standard interface to database management systems in the Windows and Windows NT environments.
open logical cluster	(Cluster concept) A logical cluster that accepts connections that have no defined route. By default, the system logical cluster has the open property, but you can grant the open property to another logical cluster. Only one logical cluster can have the open property at a time.
operating system	A group of programs that translates your commands to the computer, so that you can perform such tasks as creating files, running programs, and printing documents.
operating system file	A collection of data named and recognized by the operating system . Adaptive Server data is not stored in operating system files, but can be exported to operating system files by using the bulk copy operation.
Operator	The role required to allow the ability to back up and restore databases on a server-wide basis. See also roles .
operators	Symbols that act on two values to produce a third. See also comparison operator , logical operator , and arithmetic operators .
optimization goals	A set of user-defined goals that can be specified to influence optimization techniques considered, so as to generate plans that are suitable for a specific query or application.
optimization rules	When determining the best optimization plan, most decisions made by the optimizer are based on estimated costs, but some decisions are based on specific characteristics of the query predicate and/or tables. For example, it is best to use a join predicate between two tables, except in the case of a star join, so that the optimizer does not evaluate other join order permutations.

optimization timeout	The mechanism by which the optimizer terminates searching for a better plan than the current best plan, since compilation time has exceeded the specified criteria. The current best plan is then used for processing the query in the execution engine.
optimizer	Adaptive Server code that analyzes queries and database objects and selects the appropriate query plan . The Adaptive Server optimizer is a cost-based optimizer. It estimates the cost of each permutation of table accesses in terms of CPU cost and I/O cost.
OR strategy	An optimizer strategy for resolving queries using or and queries using in (values list). Indexes are used to retrieve and qualify data rows from a table. The row IDs are stored in a worktable. When all rows have been retrieved, the worktable is sorted to remove duplicates, and the row IDs are used to retrieve the data from the table.
ordering	Refers to a specific sequence (ascending or descending) of attributes in a result set from an index scan or sort.
outer join	A join in which both matching and nonmatching rows are returned. The operators *= and =* are used to indicate that all the rows in the first or second tables should be returned, regardless of whether or not there is a match on the join column.
outer query	Another name for the principal query in a statement containing a subquery .
outer table	Describes the placement of a table in an outer join . For example, in a left join, the outer table is the left table.
overflow page	A data page for a table with a non unique clustered index, which contains only rows that have duplicate keys. The key value is the same as the last key on the previous page in the chain. There is no index page pointing directly to an overflow page.
ownership chain	Dependencies of views on other views and/or tables, and of procedures on other procedures, views, and/or tables.
P	
package	A package is a collection of components that work together to provide a service or some aspect of your application's business logic. A package defines a boundary of trust within which components can easily communicate. Each package acts as a unit of distribution, grouping together application resources for ease of deployment and management.

page	A 2K, 4K, 8K, or 16K block of data that is the minimal unit that can be read from or written to disk.
page lock	A page lock locks an entire data or index page. Adaptive Server uses page locks as frequently as possible to reduce the contention for data among users and to improve concurrency. A page lock is less restrictive than a table lock.
page split	An action performed by Adaptive Server when new data or index rows need to be added to a page , and there is not enough room for the new rows. Usually, the data on the existing page is split approximately evenly between the newly allocated page and the existing page.
parallel processing	The simultaneous execution of Adaptive Server tasks or subtasks. Multiple processes share memory or use some type of synchronized method for passing messages between them.
parallel query optimizer	A component of the optimizer that adds scheduling information to a plan, re-evaluates the plan, and then generates the best parallel plan. The parallel query optimizer is a scheduler that looks at a set of parallel plans annotated with resource usage. Based on the total resources available for a query, it finds the best plan based on its response time.
parallel sort	A technique that employs multiple worker processes to sort data in parallel, either in response to a create index command or to a query that requires an internal sort. A single process partitions the input data into discrete ranges; multiple processes simultaneously sort each individual range of data rows and create subindexes; then a single process merges the sorted ranges and indexes into one sorted index or result set.
parameter	An argument to a stored procedure or system procedure .
Partial character	An incomplete (or truncated) series of bytes for a multibyte character in string data that uses a multibyte character set. Usually appears at the end of a string value.
partition elimination	In a query that has predicates on the partitioning keys, partition elimination can find out which partitions will not be referenced for a given predicate. However, predicates that are currently useful for partition elimination must qualify as conjunctive (clauses in a SQL query connected by an AND keyword) or disjunctive (clauses in a SQL query connected by an OR keyword) predicates on a single table of the form-column <relop> <literal>.
partition keys	Values in user-defined partitioning columns that determine the partition assignment for each table row. Only semantic partitioning strategies— hash partitioning , list partitioning , and range partitioning —have partition keys.

partition	User-created database objects that store parts of tables and indexes, and can be managed independently. You can divide a table or index into multiple partitions that reside on separate physical devices to improve the performance of concurrent inserts and to facilitate parallel queries and sorts. Table rows are assigned to partitions according to one of four strategies: round-robin, list, range, and hash.
partition skew	Uneven distribution of data across partitions .
partitioned parallelism	Data is divided into more than one physical partition so each partition can be accessed in parallel and managed by a worker thread. The IO and CPU parallelism resulting from such configuration speed up the SQL queries proportional to the number of partitions .
partitioning key	A search condition that evaluates to partition specification. The set of columns participating in the key specification is known as the partitioning key.
passthrough mode	A mode that allows clients to communicate with remote servers in native syntax. The Transact-SQL parser and compiler are bypassed in this mode, and each language “batch” received from the user is passed directly to the server to which the user is connected.
password encryption	The process of storing a password in non decipherable encrypted form.
performance	The speed with which Adaptive Server processes queries and returns results. Performance is affected by several factors, including indexes on tables, use of raw partitions compared to files, and segments.
permission	The authority to perform certain actions on certain database objects or to run certain commands.
phantom rows	Occur when one transaction reads a set of rows that satisfy a search condition, and then a second transaction modifies the data (through an insert, delete, update, and so on). If the first transaction repeats the read with the same search conditions, it obtains a different set of rows. Also known as phantoms.
physical cluster	(Cluster concept) The shared-disk cluster, with a specific quorum disk, member instances, and interconnection information. All instances in the physical cluster have direct access to a single installation of the databases and are monitored and managed by the cluster membership service.
physical design	Mapping the logical design to the Transact-SQL data definition commands that actually create the databases on the server.
physical I/O cost	The optimizer’s estimated number of physical reads.

physical key	A column name, or set of column names, used in a create index statement to define an index on a table. A physical key on a table is not necessarily the same as the logical key .
physical operator	An algorithm implementing a logical operator, for example, the logical operator “join” is implemented by the physical operator algorithms “merge join,” “hash join,” and “nested loop join.”
physical property	A property that is associated with a physical operator and depending on the actual algorithm implemented by that operator and on the physical properties of its children (hence, recursively, on the physical operators in the sub-plan). For example, the ordering (from an index scan or sort) of the outer child is usually inherited after subsequent join operators are evaluated, but each plan in an equivalence class has potentially different orderings depends on the underlying operators used in the sub-plan.
physical read	A disk I/O to access a data, index, or log page. Adaptive Server estimates physical reads and logical reads when optimizing queries. See also logical read .
PIO	See physical I/O cost .
pipelined parallelism	In a multi-step SQL operation, each independent step can begin execution before the preceding step is completed. More than one processor can work on a single query, resulting in shorter response times.
pipelining	Two sets of threads serve as producers and consumers. While producers put data in a shared buffer, consumers can process the data in the shared buffer concurrently with producers.
plan cache	In the context of the optimizer, a usage of the procedure cache in which it stores useful partial plans (plan fragments) that may be necessary in future construction of complete plans. The plan cache exists only during query compilation and is released before query execution.
platform	The underlying computer system on which application programs run. For example, both HPUX and Solaris are platforms. The platform is distinct from the operating system, which runs on the platform.
plug-ins	Software that enhances the functionality of the base application.
point query	A query that restricts results to a single specific value, usually using the form “where <i>column_value</i> = <i>search_argument</i> ”.
position number	The position of a column in the select list. You can use this instead of the column name.

positioned delete	A delete performed through a cursor, using the where current of clause. Compare to searched delete .
positioned update	An update performed through a cursor, using the where current of clause. Compare to searched update .
PowerJ	A Java component used with Adaptive Server to develop applications for the Web.
precision	The maximum number of decimal digits that can be stored by numeric and decimal datatypes. The precision includes all digits, both to the right and to the left of the decimal point.
predicate	A clause that is based on a comparison operator . It does not apply to and, or, or not.
prefetch	The process of performing multipage I/Os on a table, nonclustered index, or the transaction log. For logs, the server can fetch up to 256 pages, for nonlog tables and indexes, the server can fetch up to 8 pages.
prefetch strategy	See prefetch .
prefix subset	Used to refer to keys in a composite index. Search values form a prefix subset when leading columns of the index are specified. For an index on columns A, B, and C, these are prefix subsets: A, AB, ABC. These are not: AC, B, BC, C. See also matching index scan and nonmatching index scan .
primary companion	The Adaptive Server whose databases and connections are migrated to the secondary Adaptive Server during failover.
primary key	The column or columns whose values uniquely identify a row in a table.
primary key constraint	A primary key constraint is a unique constraint that does not permit null values for the component key columns. There can only be one primary key constraint per table. The primary key constraint creates a unique index on the specified columns to enforce this data integrity.
privilege	See permission .
process	An execution environment scheduled onto physical CPUs by the operating system.
process affinity	A process in which a certain Adaptive Server task runs only on a certain engine, or a certain engine runs only on a certain CPU.

procedure	A collection of SQL statements and optional control-of-flow statements stored under a name. Adaptive Server-supplied procedures are called system procedures.
procedure cache	Long term memory used for stored procedures, batch query plans, triggers, statement cache, and datachange tracking, plus numerous short term memory usages, for example, query compilation and update statistics.
producer process	In a parallel sort , producer processes read data simultaneously from the input table, determine the range to which each data value belongs, and distribute data values to consumer processes associated with the proper ranges.
projection	The set of attributes available on the output of an operator. This implies a minimal set of attributes in which each attribute is needed by some parent of the respective operator.
proxy authorization	The ability to impersonate another user in the server. A System Security Officer can grant proxy authorization to a user. Proxy authorization allows administrators to check permissions for a particular user or to perform maintenance on a user's database objects. Application servers can log in to the server and execute procedures and commands on behalf of several users.
proxy databases	Placeholder databases created on the secondary companion for every user database on the primary companion. Proxy databases reserve the database names so that during failover, all database names are unique on the system.
proxy table	A local table that has been mapped to a table on a remote server. The proxy table contains metadata and is used to access the remote table as if it were a local table.
pruning	An optimizer technique of searching for the best execution plan, in which only the subplans that could be part of the best total plan are retained. The optimizer uses cost-based pruning and heuristics-based pruning.
public	All registered users of a database are members of the group "public." Users at this level of authority can create a temporary table and have access to objects whose owners have granted permissions to "public."
public-key cryptography	Public-key cryptography consists of several mechanisms, including encryption, key exchange, digital signatures, and digital certificates, used to secure transmission of sensitive data over the Internet.
public/private keys	Also known as asymmetric keys, public/private are a pair of keys, one public one private, used to encrypt and decrypt messages. Public/private keys are used in public-key cryptography.

Q

QP (query processing) metrics

Query processing (QP) metrics identify and compare empirical metric values in query execution. When a query is executed, it is associated with a set of defined metrics that are the basis for comparison in QP metrics.

qualified

The name of a database object can be qualified, or preceded by, the name of the database and the object owner.

quantified predicate subquery

A subquery that operates on lists introduced with `in` or with a comparison operator modified by `any` or `all`. Quantified predicate subqueries return 0 or more values. This type is also used as an existence test, introduced with `exists`.

query

A SQL statement or group of SQL statements that access and/or manipulate data in a database. See also **data retrieval**.

query object

Makes the SQL query and manages the result set.

query plan

The ordered set of steps required to carry out a **query**, complete with the access methods chosen for each table.

query processing (QP) metrics

Query processing (QP) metrics identify and compare empirical metric values in query execution. When a query is executed, it is associated with a set of defined metrics that are the basis for comparison in QP metrics.

query processor

A query processor is designed to minimize the cost of running queries by examining parsed and normalized queries, and using statistics on database objects to determine the best possible query plan.

query tree

An internal tree structure that represents the user's **query**. A large portion of query processing and compilation is built around the shape and structure of this internal data structure. For stored procedures, views, triggers, rules and defaults these tree structures are stored in the `sysprocedures` table on disk, and read back from disk when the procedure or view is executed. See also **source text**.

quorum device

(Cluster concept) This device provides important information that defines the cluster, including the name of the cluster, the names of the instances in the cluster, the number of nodes, and their names. In addition, the quorum device holds state information about the instances in the cluster and defines cluster membership

R

range cell	A cell of a histogram which represents a range of values. The weight of the cell represents the percentage of column values that are greater than the lower bound of the cell, and less than or equal to the upper bound. Compare to frequency cell .
range cell density	A measure of duplicates in a column or set of columns which has some duplicate values removed when more than one distribution cell is spanned. Compare to total density .
range lock	A lock acquired by a range query at transaction isolation level 3, for the purpose of preventing phantom rows. Range locks prevent inserts into the range by other users until the repeatable read transaction completes. Range locks are used on data-only locked tables.
range partitioning	Data set for one or more attributes is partitioned on the value range. Thus, every row can be pinpointed to a given partition.
range query	A query that requests data within a specific range of values. These include greater than and less than queries, queries using between, and some queries using like.
RDBMS	See Relational Database Management System (RDBMS) .
read access	Permission to read an object (for example, to select rows from a table). See also permission .
record	See row .
recovery	The process of rebuilding one or more databases from database and log dumps . See also automatic recovery .
referential integrity	The rules governing data consistency, specifically the relationships among the primary keys and foreign keys of different tables. Adaptive Server addresses referential integrity with user-defined triggers and with referential integrity constraints .
referential integrity constraint	A requirement that data inserted into a “referencing” table, the table that defines the constraint, must have matching values in a “referenced” table. You cannot delete rows or update column values from a referenced table that match values in a referencing table. Also, you cannot drop the referenced table until the referencing table is dropped or the referential integrity constraint is removed.

reformatting strategy	A strategy used by Adaptive Server to resolve join queries on large tables that have no useful index. Adaptive Server builds a temporary clustered index on the join columns of the inner table, and uses this index to retrieve the rows. Adaptive Server estimates the cost of this strategy and the cost of the alternative—a table scan—and chooses the least expensive method.
relational database	A set of related data tables and other database objects that are organized and presented to serve a specific purpose. See also data warehouse .
Relational Database Management System (RDBMS)	A system for storing and retrieving data from two-dimensional tables in which the use of SQL is standard.
relational expression	<p>A type of Boolean or logical expression of the form:</p> <pre>arith_expression relational_operator arith_expression</pre> <p>In Transact-SQL, a relational expression can return TRUE, FALSE, or UNKNOWN. The results can evaluate to UNKNOWN if one or both of the expressions evaluates to NULL.</p>
relational operator	An operator that compares two operands and yields a truth value, such as “5 < 7” (TRUE), “ABC” = “ABCD” (FALSE) or “@value > NULL” (UNKNOWN).
relationship	A description of how entities are related. A basic step in logical design of a database is to identify the relationships between the entities you have identified. See also entity .
relaxed-durability database	A disk-resident database with a durability level of no_recovery or at_shutdown. Relaxed-durability databases support all features that traditional disk-resident database with full durability support (such as cache bindings). Relaxed-durability databases take advantage of the reduced durability requirements to offer significant run-time performance improvements for transactional work loads over traditional disk-resident databases.
remote login	A login to a remote server.
remote procedure call (RPC)	A stored procedure that is executed on an Adaptive Server other than the server the user is logged in to.
repeatable read	A query in which the results are not affected by changes made by other transactions. The transaction holds locks on the pages or rows it has read until the end of transaction to ensure other transactions do not update the pages or rows. Also known as transaction isolation level 2. This level does not provide phantom protection, which is provided at transaction isolation level 3.

repeating subquery	See correlated subquery .
replication	For databases, a process by which the changes to data in one database (including creation, updating, and deletion of records) are also applied to the corresponding records in other database.
reserved word	See keyword .
reserve page gap	A ratio to be used during large scale allocations performed by bulk copy, create index and select into. Setting a reserved page gap value allows future page allocations for forwarded rows and index page splits to be made close to the original location of the data.
resource reservation	(Cluster concept) The practice of setting aside an instance for a specific logical cluster and only allowing clients routed to that logical cluster to connect to it. To practice resource reservation, you must assign the open property to a logical cluster other than the system logical cluster.
response time	The time it takes for a single task, such as sending a Transact-SQL query to Adaptive Server , to complete. See also initial response time .
restriction	See selection .
return status	A value that indicates that the procedure completed successfully or indicates the reason for failure.
RID	See row ID (RID) .
roles	Titles recognized by Adaptive Server that provide individual accountability for users performing system administration and security-related tasks in Adaptive Server. The System Administrator , System Security Officer , and Operator roles can be granted to individual server login accounts. In addition to these system roles, a System Security Officer can create user-defined roles, such as “financial analyst” and “salary administrator.”
rollback record	A log record generated by Adaptive Server to describe the undoing of an earlier change of the transaction. The rollback record contains the address of the log record that is being undone. Rollback records are redo-only log records. They are never undone.
rollback transaction	A Transact-SQL statement used with a user-defined transaction , before a commit transaction has been received, that cancels the transaction and undoes any changes that were made to the database.

round-robin partitioning	A scheme that is best suited for load balancing. The data set is distributed in a round-robin fashion without any attention to where a data value ends up. There are no semantics associated with such a partitioning scheme, in other words, there is no certainty as to which partition a given row will belong.
row	A set of related columns that describes a specific entity . Also called a record .
row aggregate function	Functions (sum, avg, min, max, and count) that generate a new row for summary data when used with compute in a select statement .
row forwarding	Movement of a row to another page. Row forwarding takes place when a row is updated so that it no longer fits on the page. A pointer to the forwarded row is maintained at the original row location. The row ID of the row does not change.
row ID (RID)	A unique, internal identifier for a data row . The row ID, or RID, is a combination of the data page number and the row number on the page.
RPC	See remote procedure call (RPC) .
rule	A specification that controls what data can be entered in a particular column or in a column of a particular user-defined datatype .
run value	The value of the configuration parameter currently in use.
runserver file	The file used as a reference in restarting Adaptive Server or Backup Server. By default, the runserver file is named <i>RUN_servename</i> and is created when you install Adaptive Server. Runserver files are created in the <i>\$\$SYBASE/install</i> directory.
runtime adjustment	Adaptive Server creates an adjusted query plan to execute the query using fewer worker processes when the number of worker processes specified in the query plan at runtime is not available.
S	
“sa” login account	See System Administrator .
SARG	See search argument .
savepoint	A marker that the user puts inside a user-defined transaction . The user can later use the rollback transaction command with the savepoint name to cancel any commands up to the savepoint, or commit transaction to actually complete the commands. See also transaction and rollback transaction .
scalar	A SQL expression that produces a single value.

scalar aggregate	An aggregate function that produces a single value from a select statement that does not include a group by clause. This is true whether the aggregate function is operating on all the rows in a table or on a subset of rows defined by a where clause. See also vector aggregate .
scale	The maximum number of digits that can be stored to the right of the decimal point by a numeric or decimal datatype. The scale must be less than or equal to the precision .
scan descriptor	An internal data structure that manages the scan of tables being queried, particularly tables that have references to other tables.
schedule	A definition of a time scale for how and when jobs are executed and re-executed. Used specifically with Job Scheduler.
scheduled job	A job that has been bound to a schedule. Only scheduled jobs are executed. Used specifically with Job Scheduler.
schema	A collection of objects associated with a particular schema name and schema authorization identifier . The objects can be tables, views, domains, constraints, assertions, privileges, and so on. A schema is created by a create schema statement.
schema authorization identifier	All the objects are said to be owned by, or to have been created by, the associated schema authorization identifier for the schema.
scope	The context in which a feature, such as a cursor or a global variable , is used. In Adaptive Server , the scope can be an individual user session, a stored procedure , or a trigger .
script	A writing system, a collection of all the elements that characterize the written form of a human language—for example, Latin, Japanese, or Arabic.
scrollable cursor	A cursor that can fetch one or more rows, in either a forward or backward direction.
search argument	A predicate in a query's where clause that can be used to locate rows via an index. Also referred to as SARG.
search criteria	A user-specified or system-determined criteria used to influence the optimization techniques used to generate plans.
search engine	A component of the query optimizer that generates and evaluates the alternative execution plans, and selects the most optimal one. The search engine comprises three key components— search criteria , search space , and search strategy.

search space	The exhaustive set of plans considered for selection by the search engine.
search strategy	A module of the search engine that generates a specific search space and selects among the alternatives available in that search space.
searched delete	A delete performed while a cursor is active, but not through the cursor, that is, not using the where current of clause. Compare to positioned delete .
searched update	An update performed while a cursor is active, but not through the cursor, that is, not using the where current of clause. Compare to positioned update .
secondary companion	The Adaptive Server configured to accept the failed over primary Adaptive Server during failover.
secure sockets layer (SSL)	SSL is the standard for securing the transmission of sensitive information, such as credit card numbers, stock trades, and banking transactions, over the Internet.
secure sockets layer (SSL) filter	A filter that is appended to the master and query lines in the Adaptive Server interfaces file to specify that all connections to the specific port must support the SSL protocol before connections are established.
secure sockets layer (SSL) handshake	A series of round trip I/O between server and client, wherein the server presents its certificate, and the server and client negotiate and agree upon security mechanism before any data is transmitted.
SEEB	See Sybase Enterprise Event Broker (SEEB) .
segment	A named subset of database devices available to a particular database. It is a label that points to one or more database devices. Segments can be used to control the placement of tables and indexes on specific database devices.
select list	The columns specified in the main clause of a select statement. In a dependent view, the target list must be maintained in all underlying views if the dependent view is to remain valid.
selection	A subset of the rows in a table. Also called a restriction, it is one of the basic query operations in a relational system. See also producer process and view .
selectivity	See index selectivity and join selectivity .
self-join	A join used for comparing values within a column of a table . Since this operation involves a join of a table with itself, you must give the table two temporary names, or correlation names , which are then used to qualify the column names in the rest of the query.

semantic partitioning	Partitioning of tables or indexes according to values in the partitioning keys. See hash partitioning , list partitioning , and range partitioning .
semaphore	A semaphore is a simple internal locking mechanism that prevents a second task from accessing the data structure currently in use. Adaptive Server uses semaphores to protect transaction logs, user log caches, and I/O devices. Like a spinlock , a semaphore is relevant only in SMP environments.
semijoin	A join algorithm that terminates the inner scan for each outer row, as soon as the first inner row qualifies.
serializable read	A query within a transaction that always returns the same set of rows. Prevents phantom rows . Also called transaction isolation level 3.
server certificate	A server certificate authenticates the server that holds it to clients that attempt to connect to it.
server class	<p>An interface specification defining remote access for use with Component Integration Services. For each server class there is a separate access method that handles interaction between the server and remote servers of the same class. Each server class has a set of unique characteristics that System Administrators and programmers use to configure the server for remote data access. The characteristics are:</p> <ul style="list-style-type: none">• Mapping between local tables (proxy tables) and external objects• Datatype conversions specific to the class or access method• Special considerations for each class• Restrictions, if any, on Transact-SQL statements applied to the class
server cursor	A cursor declared inside a stored procedure . The client executing the stored procedure is not aware of the presence of these cursors. Results returned to the client for a fetch appear exactly the same as the results from a normal select .
server engine	See engine .
server user ID	The ID number by which a user is known to Adaptive Server .
session	Sessions allow a browser to maintain a connection with the server across a multiple request-response cycle
severity level number	A number between 10 and 24 which indicates the severity of an error condition.

shared-disk cluster	(Cluster concept) A cluster configuration where all instances have direct access to all data on all shared disks. In the Cluster Edition, all instances have direct access to database devices and jointly manage the single installation of the databases.
shared lock	A type of lock acquired on an object for a read operation. It does not allow other transactions to acquire exclusive locks on the object but allows them to acquire shared locks. Shared locks can be obtained on a table or page.
shared row lock	A type of lock acquired on a data row for a read operation. It does not allow other transactions to acquire exclusive locks on the row, but allows them to acquire shared locks.
single-user mode	Starting Adaptive Server in single-user mode allows only one System Administrator to log in, and turns on the allow updates configuration variable. Use this mode to restore the master database. This option creates a <i>m_RUN_servername</i> file and overwrites any existing <i>m_RUN_servername</i> file.
site handler	The default method of transmitting RPCs from a local server to a remote server. Site handler establishes a single physical connection between local and remote servers and multiple logical connections as required by RPCs.
SMP	symmetric multiprocessing (SMP).
snowflake schema joins	Queries where several dimension tables are joined to local fact tables that in turn are joined to a central fact table. There are no join clauses between the dimension tables , and the fact tables are large compared to their respective dimension tables
soft fault	A corruption in Adaptive Server that is usually not persistent. Most soft faults result from temporary inconsistencies in the target database caused by users updating the target database during the dbcc checkstorage operation or by dbcc checkstorage encountering data definition language (DDL) commands. These faults are not repeated when you run dbcc checkstorage a second time. See also hard fault .
sort order	A specification used by Adaptive Server to determine the order in which to sort character data. Also called collating sequence .
source text	SQL statements that define a compiled object . Source text is stored in the text column of the syscomments table. It is used during upgrade processes; therefore, it must not be deleted. Source text can be encrypted using sp_hidetext.

space management property	One of the properties of a table or index that helps manage how space for the table is allocated or used: <code>fillfactor</code> , <code>max_rows_per_page</code> , <code>exp_row_size</code> and <code>reservepagegap</code> . Space management properties are specified for tables with <code>create table</code> or <code>alter table</code> and for indexes with <code>create index</code> or the <code>alter table</code> commands that generate indexes to enforce unique or primary key constraints. Some of the space management properties can be changed with <code>sp_chgattribute</code> .
space reclamation	The process of recovering space from data-only locked tables after deletes and updates that shrink rows. Space reclamation is performed by tasks inserting or updating rows on a page, by the housekeeper, and by the <code>reorg</code> command.
sparse frequency count	A frequency count in which values are not contiguous in the domain. For example, the set of values 1, 10, 100 is sparse in the integer domain, but 1, 2, 3 is not. Compare to dense frequency count .
spinlock	A simple internal locking mechanism that prevents a process from accessing the resource currently used by another process. All processes trying to access the resource must wait (or “spin”) until the lock is released. Spinlocks protect internal data structures such as a data cache.
SQL	See Structured Query Language (SQL) .
SQL derived table	A table defined by one or more tables through the evaluation of a query expression. A SQL derived table is used in the query expression in which it is defined and exists only for the duration of the query. See also derived table and abstract plan derived table .
SQL Server	See Adaptive Server .
sql.ini file	The interfaces file containing definitions for each Adaptive Server to which your workstation can connect. The file must be on each machine from which clients connect to Adaptive Server. Each <code>sql.ini</code> file entry tells a client or host machine how to connect to a specific Adaptive Server. The file contains the name of Adaptive Server, a list of services provided by Adaptive Server, and the port to use for connecting to Adaptive Server for each service.
SSL	See secure sockets layer (SSL) .
stable mode	A stable mode is a system state in which Adaptive Server can exist for an extended period of time, such as the day-to-day operation of Adaptive Server.

standby access server	A server that is an exact duplicate of an Adaptive Server production server. A standby server is created by loading transaction log dumps from the production server to the standby server as the log dumps are made. Standby servers can provide emergency coverage in case the main server fails, or they can be read-only servers used for reporting or decision support.
star schema joins	Queries where several dimension tables are joined to a central fact table . The dimension tables do not have join clauses between them, and the fact table is large compared to its dimension tables
stateful component	A component that can remain active between consecutive method invocations.
stateful session Beans	Session Bean instances are transient and maintain a one-to-one relationship with the client. They perform tasks and store information in the database on the client's behalf. Stateful session Beans manage complex tasks that require the accumulation of data.
stateless component	A component that is deactivated after each method call and supports instance pooling. The component's state is reset across the boundary of a transaction and activation.
stateless session Beans	Session Bean instances are transient and maintain a one-to-one relationship with the client. They perform tasks and store information in the database on the client's behalf. Stateless session Beans manage tasks that do not store data between method calls.
statement	One or more complete Transact-SQL commands . A statement usually consists of a keyword , for example, select , followed by clauses such as from or where .
statement block	A series of Transact-SQL statements enclosed between the keywords begin and end so that they are treated as a unit.
static configuration parameter	An Adaptive Server configuration parameter that requires you to restart Adaptive Server after resetting the value. See also dynamic configuration parameter .
step	A pair of contiguous values in a distribution or histogram . Each step is associated with the set of rows having column values that fall between the upper and lower bound of the step. Also called a cell .
store operator	An operator that creates a fully materialized table, usually in support of the reformatting strategy.
stored procedure	A collection of Transact-SQL statements and optional control-of-flow statements stored under a name. See also system procedure .

string function	A function that operates on strings of characters or binary data. <code>substring</code> and <code>charindex</code> are Transact-SQL string functions.
striping	See dump striping .
Structured Query Language (SQL)	The language used to communicate with a relational database and that is the subject of standards set by several standards bodies.
subject	A server process acting on behalf of a user. In Adaptive Server , subjects include users, as well as stored procedures and trusted triggers while executing.
surrogate pairs	A coded character representation for a single abstract character that consists of a sequence of two code values. Surrogate pairs are designed to allow additional 2^{20} code values to be represented in the Unicode standard. Applies only to UTF-16 encoding
subquery	A select statement that is nested inside another select, insert, update, or delete statement or inside another subquery.
suspect page	A page that recovery has marked as suspect because of corruption. Suspect pages are normally inaccessible to users unless they have been forced online by special procedures.
suspended companion mode	The mode of Adaptive Server after companion mode has been suspended. During this mode, Adaptive Server cannot fail over; it is working independently of the other Adaptive Server.
Sybase Enterprise Event Broker (SEEB)	Enterprise Event Broker provides an infrastructure for enabling Sybase servers to use messaging.
symmetric key encryption	A process wherein the same algorithm (key) is used to encrypt and decrypt a message.
symmetric multiprocessing (SMP)	An environment in which multiple CPUs simultaneously process client tasks. Any CPU can execute any task. The CPUs symmetrically serve and run all functionality of the operating system and applications. This is the non-clustered Adaptive Server environment
symmetrical	A high availability system in which two independent Adaptive Servers act as failover servers for each other. That is, each Adaptive Server acts as both a primary and a secondary companion.
system administration	An assortment of tasks including but not limited to managing Adaptive Server's physical storage, creating and backing up databases, creating user accounts, granting permissions, and running diagnostic and repair functions. See also database administration .

System Administrator	A user in charge of Adaptive Server system administration, including managing disk storage, granting and revoking the System Administrator role, and creating new databases. The “sa” account, a single login , is created when Adaptive Server is installed. This login is configured with both the System Administrator and System Security Officer roles. To increase individual accountability, use the “sa” account to assign roles to individual logins.
system database	A database provided Sybase, for example, master, tempdb, model, and sybssystemprocs.
system function	A function that returns special information from the database, particularly from the system tables .
system logical cluster	(Cluster concept) A logical representation of the physical cluster. The system logical cluster is automatically created when the physical cluster is created, and it has the same name as the physical cluster. All background tasks run on the system logical cluster.
system procedure	A stored procedure residing in sybssystemprocs, which is provided by Sybase. System procedures provide shortcuts for retrieving information from the system tables , or mechanisms for accomplishing database administration and other tasks that involve updating system tables.
system procedure tables	Tables in the master database that the system procedures use to convert internal system values (for example, status bits) into human-readable format.
System Security Officer	A user who controls security-related operations in Adaptive Server , including auditing, locking and unlocking login accounts, creating user-defined roles, and password management. See also System Administrator .
system table	One of the data dictionary tables. System tables keep track of information about the Adaptive Server as a whole and about each user database. The master database contains some system tables that are not in user databases.

T

table	A collection of rows that have associated columns . The logical equivalent of a database file.
table scan	A method of accessing a table by reading every row in the table. Table scans are used when there are no conditions (where clauses) on a query, when no index exists on the clauses named in the query, or when the Adaptive Server optimizer determines that an index should not be used because it is more expensive than a table scan. See also access method .

table-level constraint	A constraint that limits values on more than one column of a table . Enter table-level constraints as separate comma-delimited clauses in the create statement. You must declare constraints that operate on more than one column as table-level constraints.
table-normalized histogram	a histogram computed by update statistics in which the weight array values (in other words, fractions of table rows) always sum to 1.0.
target	A target is an application , applet, class , or collection that you create with PowerJ .
target database	The database that you want to check with dbcc.
target segment	The target segment for a sort is the segment where the index is stored when the create index command completes.
task	A unit of execution scheduled by the kernel to fulfill a request for service.
TCP/IP	The primary transport protocol used in client/server computing. TCP/IP is the protocol that governs the transmission of data over the Internet.
temporary database	A database that provides a storage area for temporary tables and other temporary working storage needs (for example, intermediate results of group by and order by). In Adaptive Server , the temporary database is tempdb.
terminator	The characters that separate columns .
text and image function	A function that operates on text and image data. The text and image functions include patindex, textptr, and textvalid.
text chain	A special data structure used to store text and image values for a table . Data rows store pointers to the location of the text or image value in the text chain.
theta join	A join that uses a comparison operator as the join condition.
thread	See worker process .
thread of execution	See worker process .
threshold	The estimate of the number of log pages required to back up the transaction log, and the action to be taken when the amount of space falls below that value.
throughput	The volume of work completed in a given time period. It is usually measured in transactions per second (TPS).
time of enforcement	The phase of query processing during which Adaptive Server applies a given resource limit.

time slice	The amount of time a task is allowed to run on an Adaptive Server engine before voluntarily yielding the engine to another task. To see the value on your server enter: <pre>sp_configure time slice</pre>
total density	A measure of duplicates in a column or set of columns which represents all duplicates even when more than one histogram cell is spanned. Compare to range cell density .
TPS	Transactions per second. See also throughput .
transaction	A group of Transact-SQL statements that is treated as a single unit of work. Either all statements in the group are executed or no statements are executed. The tables queried during the transaction are locked until the transaction is complete.
transaction attribute	Determines how the component participates in transactions.
transaction coordinator	The transaction coordinator manages the flow of transactions that involve more than one connection and sometimes more than one data source.
transaction descriptor	An internal memory structure that Adaptive Server uses to represent a transaction.
transitive closure	a set of attributes connected by equijoins .
translation isolation level	See isolation level .
transaction log	A system table (syslogs) in which all changes to the database are recorded.
transaction log options	You have the following options for the transaction log: <ul style="list-style-type: none">• Backup, truncate, and log – back up the transaction log, remove the inactive part of the log, and create a new transaction log entry recording the backup.• Backup and log – back up the transaction log and create a new transaction log entry recording the backup. This option retains the transaction log entries.• Truncate and log – remove the inactive part of the log without backing it up and create a new transaction log entry recording the dump.• Truncate only – remove the inactive part of the log without backing it up and without creating a new transaction log entry to record the dump.

transaction mode	In a SQL transaction, allow you to set whether transactions begin with or without an implicit begin transaction statement.
transaction object	The Java client application uses a transaction object to make a database connection.
Transact-SQL	The SQL dialect used in Sybase Adaptive Server .
transitional mode	The failback transitional mode occurs when Adaptive Server shifts from failed over mode to normal companion mode. The failback transitional mode is typically of a very short duration.
trigger	A special form of stored procedure that goes into effect when a user gives a change command such as insert, delete, or update to a specified table or column. Triggers are often used to enforce referential integrity .
trigger actions	The action for which a trigger is specified.
trigger conditions	The conditions that cause a trigger to take effect.
trigger table	The table to which a trigger is attached.
trigger test tables	When a data modification affects a key column, triggers compare the new column values to related keys by using temporary worktables called trigger test tables.
trusted root file	A list of trusted Certificate Authorities (CAs) loaded by the server at start-up.
trusted user	A user who is authorized to operate over a range of sensitivity levels in a session . The user logs in by specifying a range of levels over which he or she would like to operate. The user can change current read and write levels during a session.
tuple	A row in a stored or derived table .
tuple filtering	An execution operator with a single input stream that assumes all referenced attributes are ordered, and eliminates duplicate tuples based on this assumption.
U	
unbound predicate	Predicates where the constants are not known at compile time, such as local variables.
unchained transaction mode	The transaction mode that requires explicit begin transaction statements paired with commit transaction or rollback transaction statements to complete a transaction. See also chained transaction mode .

underlying tables	See base tables .
ungrouped aggregate	See scalar aggregate .
Unicode	An international character set that supports over 650 of the world's languages, such as Japanese, Chinese, Russian, French, and German. Unicode allows you to mix different languages from different language groups in the same server, no matter what the platform.
unified login	The capability for users to be preauthenticated by a security mechanism before logging in to the server. This capability enables a user to log in to several servers without having to supply a login name and password for every connection.
unique constraint	A constraint requiring that all non-null values in the specified columns must be unique. No two rows in the table are allowed to have the same value in the specified column. The unique constraint creates a unique index on the specified columns to enforce this data integrity.
unique indexes	Indexes that do not permit any two rows in the specified columns to have the same value. Adaptive Server checks for duplicate values when you create the index (if data already exists) and each time data is added.
update	An addition to, deletion from, or change in data, involving the insert, delete, truncate table, or update statements.
update inplace	See in-place update .
update lock	A type of lock acquired on an object during the scanning phase of an update or delete operation. It does not allow other transactions to acquire update or exclusive locks on the object but allows them to acquire shared locks . Update locks can be obtained only on a page.
update row lock	A type of lock acquired on a data row during the scanning phase of an update or delete operation. It does not allow other transactions to acquire update or exclusive locks on the row but allows them to acquire shared locks .
user database	A database created by a user. See also system database .
user ID	The ID number by which a user is known in a specific database. See also server user ID .
user table	A database table that stores user data.
user-defined cache	See buffer cache .

user-defined datatype	A definition, created by the user, of the type of data (the datatype) contained in a column . The user defines these datatypes in terms of the existing system datatypes. Rules and defaults can be bound to user-defined datatypes, but not to system datatypes.
user-defined transaction	See transaction .
UTF-16	Each UCS-2 code value represents itself. Code values beyond the BMP (Basic Multilingual Plane: 0..0xFFFF) are represented using pairs of special 16-bit codes called surrogate pairs. This allows an additional 2^{20} (1MB) of code values to be represented, at the expense of using 4 bytes to do so.
UTF-8	length encoding of the Unicode Standard using 8-bit sequences, where the high bits indicate which part of the sequence a byte belongs to.
V	
variable	An entity that is assigned a value. Adaptive Server has two kinds of variables, local variable and global variables .
vector aggregate	A value that results from using an aggregate function with a group by clause. See also scalar aggregate .
vertical parallelism	the ability to employ multiple CPUs simultaneously on more than one operator in a single plan fragment. Also called pipelined parallelism .
view	A named select statement that is stored in the database as an object. It allows users to see a subset of rows or columns from one or more tables. See also producer process and selection .
view resolution	In queries that involve a view , view resolution is the process of verifying the validity of database objects in the query and combining the query and the stored definition of the view.
virtual column	Any column in the output of an execution engine operator that does not map directly back to a persistent column of a table . Usually an expression involved on one side of a join .
virtual computed column	A computed column whose result is not stored in a specified table or index page.

Virtual Server Architecture (VSA) A database architecture designed to make efficient use of hardware and operating system resources. It is Sybase's implementation of SMP processing. VSA is based on the use of **Adaptive Server** engines to achieve fully symmetric processing and Sybase's own database kernel—a multithreaded operating system that handles typical operating system functions, including scheduling and dispatching tasks.

W

wash area An area of a **buffer pool** near the LRU end of the MRU/LRU page chain. After pages enter the wash area, **Adaptive Server** initiates an asynchronous write on the pages. The wash area provides clean buffers at the LRU for any query that needs to perform a disk I/O.

wash marker A point in the cache on the MRU/LRU chain. **Adaptive Server** data cache is managed as an MRU/LRU (most recently used/least recently used) chain of buffers. As buffers age in the cache, they move from the MRU end toward the LRU end.

wildcard character A special character used with the Transact-SQL like keyword that can stand for one or any number of characters in pattern matching.

with cascade option An option, when revoking permission, that revokes permission from a specified user or group and also from anyone to whom the specified user or group has granted permission.

with grant option An option, when granting permission, that allows a specified user or group to grant the same permission to other users or groups.

worker process An **Adaptive Server** subtask spawned by an Adaptive Server task in a parallel query processing environment. Together with other worker processes, executes parts of a **query** simultaneously to reduce response time. Worker processes use only a small amount of the operating system resources that a process uses. They do not need any operating system resources after they are launched. They can share memory space with each other. Also called a **thread** or lightweight process.

workload manager (Cluster concept) An Adaptive Server module that provides application-level management of resource allocation, availability, and load distribution.

workspace A collection of one or more **targets**. Your workspace can include all the targets involved in the complete application.

worktable See **dynamic index**.

- write access** Permission to write an object, for example, to update a **row** or to add a row to a **table**.
- write-ahead log** A log, such as the transaction log, that **Adaptive Server** automatically writes to when a user issues a statement that would modify the database. After all changes for the statement have been recorded in the log, they are written to an in-cache copy of the datapage.
- X**
- XP Server** An Open Server™ process that executes extended **stored procedures** on behalf of **Adaptive Server**. XP Server and Adaptive Server run on the same machine and communicate with each other through a **remote procedure call (RPC)**.

