System Administration Guide: Volume 1

SAP® Adaptive Server® Enterprise 16.0

DOCUMENT ID: DC31654-01-1600-01
LAST REVISED: May 2014

Contents

CHAPTER 1 Overview of System Administration

The system administraor is responsible for setting up and maintaining SAP® Adaptive Server® Enterprise (ASE).

SAP® ASE administration tasks include:

- Installing SAP ASE and Backup Server
- Creating and managing SAP ASE login accounts
- Granting roles and permissions to SAP ASE users
- Managing and monitoring the use of disk space, memory, and connections
- Backing up and restoring databases
- Diagnosing system problems
- Configuring SAP ASE to achieve the best performance

In addition, system administrators may assist with certain database design tasks that overlap with the work of application designers, such as enforcing integrity standards.

Although a system administrator generally concentrates on tasks that are independent of the applications running on SAP ASE, he or she is likely to have the best overview of all applications. For this reason, a system administrator can advise application designers about existing data, make recommendations about standardizing data definitions across applications, and so on.

However, the distinction between what is specific to an application is sometimes unclear. Owners of user databases might consult certain sections of this book. Similarly, system administrators and database owners will use the *Transact-SQL Users Guide* (especially information about data definitions, stored procedures, and triggers). Both system administrators and application designers will use the *Performance and Tuning Series*.

Roles Required for System Administration Tasks

Many SAP ASE commands and procedures require the system administrator or system security officer role. Other system administration information is relevant to database owners.

Various security-related, administrative, and operational tasks are grouped into the following user roles

- *system administrator* – by default, the system administrator (sa) is assigned these roles:
  - **sa_role**

    - **sso_role**
    - **oper_role**
    - **sybase_ts_role**

  The system administrator’s tasks that are related to SAP ASE include:

    - Managing disk storage
    - Monitoring the automatic recovery procedure
    - Fine-tuning by changing configurable system parameters
    - Diagnosing and reporting system problems
    - Backing up and loading databases
    - Modifying and dropping server login accounts
    - Granting and revoking the system administrator role
    - Granting permissions
    - Creating user databases and granting ownership of them
    - Setting up groups, which can be used for granting and revoking permissions
- *System security officer* – performs security-related tasks, such as:
    - Creating server login accounts, which includes assigning initial passwords
    - Changing the password of any account
    - Granting and revoking the system security officer and operator roles
    - Creating, granting, and revoking user-defined roles
    - Granting the capability to impersonate another user throughout the server
    - Setting the password expiration interval
    - Setting up network-based security services
    - Managing the audit system
- *Operator* – backs up and loads databases on a server-wide basis. The operator role allows a single user to use the **dump database**, **dump transaction**, **load database**, and **load transaction** commands to back up and restore all databases on a server without having to be the owner of each one. These operations can be performed for an individual database by the database owner or by a system administrator. However, an operator can perform them for any database.

These roles provide individual accountability for users who are performing operational and administrative tasks. Their actions can be audited and attributed to them. A system administrator operates outside the discretionary access control (DAC) protection system; that is, when a system administrator accesses objects, SAP ASE does not check the DAC permissions.

In addition, two kinds of object owners have special status because of the objects they own:

- Database owner
- Database object owner

Database Owner

The *database owner* is the creator of a database or someone to whom database ownership has been transferred. A system administrator can use the **grant** command to grant users the authority to create databases.

A database owner logs in to SAP ASE using his or her assigned login name and password, and has the "dbo" account. When this user logs in to databases they did not create, this user is known by his or her regular user name.

A database owner can:

- Run the system procedure **sp_adduser** to allow other SAP ASE users access to the database
- Use the **grant** command to give other users permission to create objects and execute commands within the database

Adding users to databases is discussed in *Security Administration Guide > Managing SAP ASE Logins and Database Users*. Granting permissions to users is discussed in *Security Administration Guide > Managing User Permissions*.

The database owner does not automatically receive permissions on objects owned by other users. However, a database owner can temporarily assume the permissions of other users in the database at any time by using the **setuser** command. Using a combination of the **setuser** and **grant** commands, the database owner can acquire permissions on any object in the database.

Note: Because the database owner role is so powerful, the system administrator should carefully plan is allowed to own databases in the server. The system security officer should consider auditing the database activity of all database owners.

Database Object Owner

A *database object owner* is a user who creates a database object.

Database objects include tables, indexes, views, defaults, triggers, rules, constraints, and procedures. Before a user can create a database object, the database owner must grant the user permission to create objects of a particular type. There is no special login name or password for a database object owner.

The database object owner creates an object using the appropriate **create** statement, and then grants permission to other users.

The creator of a database object is automatically granted all permissions on that object. The system administrator also has all permissions on the object. The owner of an object must explicitly grant permissions to other users before they can access the object. Even the database owner cannot use an object directly unless the object owner grants him or her the appropriate permission. However, the database owner can always use the **setuser** command to impersonate any other user in the database, including the object owner.

Note: When a database object is owned by someone other than the database owner, the user (including a system administrator) must qualify the name of that object with the object owner's name—*ownername.objectname*—to access the object. If an object or a procedure must be accessed by a large number of users, particularly in ad hoc queries, having these objects owned by "dbo" greatly simplifies access.

Performing System Administration Tasks

You can enter SAP ASE commands using either **isql** or the SAP® Control Center graphic tool.

For complete information about using **isql**, see the *Utility Guide*. For complete information about using SCC, see the online help.

You can use isql to enter many Transact-SQL.

To start isql on most platforms, enter this command at an operating system prompt, where *username* is the system administrator:

```
isql -Uusername
```

SAP ASE prompts you for your password. To prevent other users from seeing your password, do not use the **isql -P** option.

The statements that you enter in **isql** can span several lines. Statements are not processed until you enter `go` on a separate line.

For example:

```
1> select *
2> from sysobjects
3> where type = "TR"
4> go
```

The examples used throughout the SAP ASE documentation do not include the `go` command between statements. If you are typing the examples, you must enter the `go` command to see the sample output.

You can save the Transact-SQL statements you use to create or modify user databases and database objects.

The easiest way to do this is to create or copy the statements to an ASCII-formatted file. You can then use the file to supply statements to isql to re-create databases or database objects later.

The syntax for using isql with an ASCII-formatted file is the following, where filename is the full path and file name of the file that contains Transact-SQL statements:

```
isql -Uusername -ifilename
```

On UNIX and other platforms, use the "less than" symbol (<) to redirect the file.

The Transact-SQL statements in the ASCII file must use valid syntax and the go command.

When reading commands from a file, you must:

- Supply the **-P***password* option at the command line, or,
- Include the named user's password on the first line of the input file

.

System Tables

The `master` database contains *system tables* that keep track of information about SAP ASE. In addition, each database (including the `master` database) contains system tables that keep track of information that is specific to that database.

All the SAP ASE-supplied tables in the `master` database (the SAP ASE controlling database) are considered system tables. Each user database is created with a subset of these system tables. The system tables may also be called the *data dictionary* or the system catalogs.

A `master` database and its tables are automatically created when SAP ASE is installed. The system tables in a user database are created when the **create database** command is issued. The names of all system tables start with "sys". You cannot create tables in user databases that have the same names as system tables. See *Reference Manual: Tables* for detailed descriptions of system tables and their columns.

Querying the System Tables

Query system tables in the same manner as any other tables.

For example, the following statement returns the names of all triggers in the database:

```
select name
from sysobjects
where type = "TR"
```

In addition, SAP ASE supplies *stored procedure*s (called *system procedures*), many of which provide shortcuts for querying the system tables.

These system procedures provide information from the system tables:

• **sp_commonkey**	• **sp_helpremotelogin**
• **sp_configure**	• **sp_help_resource_limit**
• **sp_countmedatada**	• **sp_helprotect**
• **sp_dboption**	• **sp_helpsegment**

• **sp_estspace**	• **sp_helpserver**
• **sp_help**	• **sp_helpsort**
• **sp_helppartition**	• **sp_helptext**
• **sp_helpcache**	• **sp_helpthreshold**
• **sp_helpconfig**	• **sp_helpuser**
• **sp_helpconstraint**	• **sp_lock**
• **sp_helpdb**	• **sp_monitor**
• **sp_helpdevice**	• **sp_monitorconfig**
• **sp_helpgroup**	• **sp_showcontrolinfo**
• **sp_helpindex**	• **sp_showexeclass**
• **sp_helpjava**	• **sp_showplan**
• **sp_helpjoins**	• **sp_spaceused**
• **sp_helpkey**	• **sp_who**
• **sp_helplanguage**	• **sp_help_resource_limit**
• **sp_helplog**	

For complete information about the system procedures, see the *Reference Manual: Procedures*.

Keys in System Tables

Primary, foreign, and common keys for system tables are defined in the `master` and `model` databases. You can generate a report on defined keys by executing **sp_helpkey**. For a report on columns in two system tables that are likely join candidates, execute **sp_helpjoins**..

The *System Tables Diagram* shows the relationships between columns in the system tables.

Updating System Tables

The SAP ASE system tables contain information that is critical to the operation of your databases. Under ordinary circumstances, you need not perform direct data modifications to system tables.

Update system tables only when you are instructed to do so by SAP Technical Support, by an instruction in the *Error Messaging and Troubleshooting Guide*, or as required by some system administration tasks.

Before you update system tables, you must issue an **sp_configure** command that enables system table updates. While this command is in effect, any user with appropriate permission can modify a system table. Other requirements for direct changes to system tables are:

- Modify system tables only inside a transaction. Issue a **begin transaction** command before you issue the data modification command.
- Verify that only the rows you wanted changed have been affected by the command, and that the data has been changed correctly.
- If the command was incorrect, issue a **rollback transaction** command. If the command was correct, issue a **commit transaction** command.

Warning! Some system tables should not be altered by any user under any circumstances: these tables are dynamically built by system processes, contain encoded information, or display only a portion of their data when queried. Imprudent, ad hoc updates can prevent SAP ASE from running, make database objects inaccessible, scramble permissions on objects, or terminate a user session. Never attempt to alter system table definitions; for example, do not alter them to include constraints: triggers, defaults, and rules are not allowed in system tables.

System Procedures

The names of all system procedures begin with "sp_". They are located in the `sybsystemprocs` database, but you can run many of them in any database by issuing the stored procedure from the database or by qualifying the procedure name with the database name.

SAP-supplied system procedures (such as **sp_who**) are created using the `installmaster` installation script. Use **sp_version** to determine the version of `installmaster` that was most recently executed. See the *Reference Manual: System Procedures* for more information about **sp_version**.

If you execute a system procedure in a database other than `sybsystemprocs`, the procedure operates on the system tables in the database from which it was executed. For example, if the database owner of `pubs2` runs **sp_adduser** from `pubs2` or issues the command **pubs2..sp_adduser**, the new user is added to `pubs2..sysusers`. However, this does not apply to system procedures that update only tables in the master database.

Permissions on system procedures are discussed in the *Reference Manual: Procedures*.

Using System Procedures

A *parameter* is an argument to a stored or system procedure.

If a parameter value for a system procedure contains reserved words, punctuation, or embedded blanks, you must enclose it in single or double quotes. If the parameter is an object name, and the object name is qualified by a database name or owner name, enclose the entire name in single or double quotes.

System procedures can be invoked during a session using either chained or unchained transaction mode. Chained mode implicitly begins a transaction before any data retrieval or modification statement. Unchained mode requires explicit **begin transaction** statements paired with **commit transaction** or rollback transaction statements to complete the transaction. See *Transact-SQL Users Guide > Transactions: Maintaining Data Consistency and Recovery*.

You cannot execute the system procedures that modify data in system tables in the `master` database from within a transaction, since this may compromise recovery. You cannot run system procedures that create temporary worktables from transactions.

If no transaction is active when you execute a system procedure, SAP ASE turns off chained mode and sets **transaction isolation level 1** for the duration of the procedure. Before returning, the session's chained mode and isolation level are reset to their original settings. See *Transact-SQL Users Guide > Transactions: Maintaining Data Consistency and Recovery*.

All system procedures report a return status. For example, the following means that the procedure executed successfully:

```
return status = 0
```

If the system procedures do not execute successfully, the return status is a number other than 0.

System Procedure Tables

System procedures use several *system procedure tables* in the `master` and `sybsystemdb` databases to convert internal system values (for example, status bits) into human-readable format.

One of these tables, `spt_values`, is used by a variety of system procedures, including:

- **sp_configure**
- **sp_dboption**
- **sp_depends**
- **sp_help**
- **sp_helpdb**
- **sp_helpdevice**

- **sp_helpindex**
- **sp_helpkey**
- **sp_helprotect**
- **sp_lock**

The `spt_values` table can be updated only by an upgrade; you cannot modify it. To see how it is used, execute **sp_helptext** and look at the text for one of the system procedures that references it.

The other system procedure tables are `spt_monitor`, `spt_committab`, and tables needed by the catalog stored procedures. (The `spt_committab` table is located in the `sybsystemdb` database.)

In addition, several system procedures create, and then drop, temporary tables. For example, **sp_helpdb** creates `#spdbdesc`, **sp_helpdevice** creates `#spdevtab`, and **sp_helpindex** creates `#spindtab`.

Creating Stored Procedures

System administrators can write stored procedures that can be executed in any database.

The *System Administration Guides* include discussions about some system procedures that are relevant to system administration. Many system procedures are explained in this manual, in the sections where they are relevant. For detailed reference information about all procedures, see the *Reference Manual: Procedures*.

Create a stored procedure in `sybsystemprocs` and assign it a name that begins with "**sp_**". The `uid` of the stored procedure must be 1, the `uid` of the database owner.

Most system procedures that you create query the system tables. SAP recommends that you do not create stored procedures that modify the system tables.

1. Use **sp_configure** to set **allow updates to system tables** on.
2. Use **create procedure** to create the stored procedure.
3. Use **sp_configure** to set **allow updates to system tables** off.

Warning! Use caution when you modify system tables. Always test the procedures that modify system tables in development or test databases, rather than in your production database.

System Extended Stored Procedures

An extended stored procedure (ESP) lets you call external language functions from within SAP ASE.

SAP ASE includes a set or predefined ESPs; users can also create their own. The names of all system extended stored procedures begin with “xp_”, and are located in the `sybsystemprocs` database.

One very useful system ESP is **xp_cmdshell**, which executes an operating system command on the system that is running SAP ASE.

Invoke a system ESP just like a system procedure. The difference is that a system ESP executes procedural language code rather than Transact-SQL statements. All ESPs are implemented by an Open Server™ application called XP Server™, which runs on the same machine as SAP ASE. XP Server starts automatically on the first ESP invocation.

For information about the system ESPs provided with SAP ASE, see the *Reference Manual: Procedures*.

Creating System ESPs

Use **create procedure** to create a system ESP in the `sybsystemprocs` database.

System procedures are automatically included in the `sybsystemprocs` database. The name of the ESP, and its procedural language function, must begin with “xp_”. The `uid` of the stored procedure must be 1, the `uid` of the database owner.

For general information about creating ESPs, see *Transact-SQL Users Guide* > *Using Extended Stored Procedures*.

Logging Error Messages

Each time it starts, SAP ASE writes start-up information to a local error log file. The installation program automatically sets the error log location when you configure a new SAP ASE.

Many error messages from SAP ASE are written only to the user’s terminal. However, fatal error messages (severity levels 19 and above), kernel error messages, and informational messages from SAP ASE are recorded in the error log file.

SAP ASE keeps the error log file open until you stop the server process. Before deleting old messages to reduce the size of the error log, stop the SAP ASE process.

Note: On some platforms, such as Windows, SAP ASE also records error messages in the operating system event log. See the installation guide and configuration guide for your platform.

See the *Configuration Guide* for your platform to view the default location and file name of the error log for your platform.

Connecting to SAP ASE

SAP ASE can communicate with other SAP ASEs, Open Server applications, and client software on the network.

Clients can talk to one or more servers, and servers can communicate with other servers using remote procedure calls. For products to interact with one another, each must know where the others reside on the network. This network service information is stored in the `interfaces` file.

The interfaces File

The `interfaces` file, which lists the name and address of every known server, is usually named `interfaces`, `interface`, or `sql.ini`, depending on the operating system.

When you use a client program to connect to a server, the program looks up the server name in the `interfaces` file and then connects to the server using the address.

Figure 1: Connecting to SAP ASE

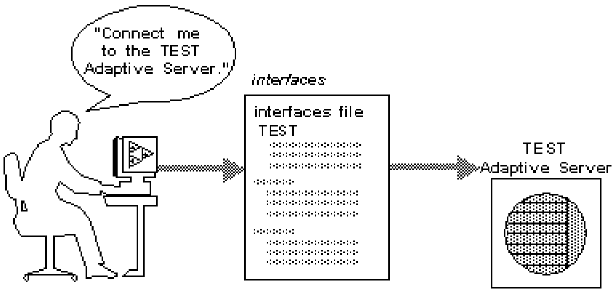

The name, location, and contents of the `interfaces` file differ between operating systems. Also, the format of the SAP ASE addresses in the `interfaces` file differs between network protocols.

The SAP ASE installation program creates a simple `interfaces` file that you can use for local connections to SAP ASE over one or more network protocols. As a system administrator, modify the `interfaces` file and distribute it to users so that they can connect to SAP ASE over the network. See the *Configuration Guide* for your platform for information about the `interfaces` file.

See *Performance and Tuning Series: Basics* > *Networks and Performance* for more information about the interfaces file and network listeners.

Directory Services

A directory service manages the creation, modification, and retrieval of network service information. Directory services are provided by platform or third-party vendors and must be purchased and installed separately from SAP ASE. Two examples of directory services are Registry, and Distributed Computing Environment (DCE).

The `$SYBASE/$SYBASE_OCS/config/libtcl.cfg` file is a SAP-supplied configuration file used by servers and clients to determine:

- Which directory service to use, and
- The location of the specified directory service driver.

If no directory services are installed or listed in the `libtcl.cfg` file, SAP ASE defaults to the `interfaces` file for obtaining network service information.

The system administrator must modify the `libtcl.cfg` file as appropriate for the operating environment.

Some directory services are specific to a given platform; others can be used on several different platforms. Because of the platform-specific nature of directory services, see the configuration documentation for your platform for detailed information about configuring for directory services.

LDAP as a Directory Service

Lightweight Directory Access Protocol (LDAP) is an industry standard for accessing directory services. Directory services allow components to look up information by a distinguished name (DN) from an LDAP server that stores and manages server, user, and software information that is used throughout the enterprise or over a network.

The LDAP server can be located on a different platform from the one on which SAP ASE or clients are running. LDAP defines the communication protocol and the contents of messages exchanged between clients and servers. Messages are operators, such as client requests for read, write and query, and server responses, including metadata (data about data).

The LDAP server can store and retrieve information about:

- SAP ASE, such as IP address, port number, and network protocol
- Security mechanisms and filters
- High availability companion server name
- Authentication information for user access to SAP ASE

  You can authenticate users logging in to SAP ASE through information stored in the `syslogins` directory or through a centralized LDAP server that enables a single login and password throughout the enterprise. See *Security Adiminstration Guide > Managing SAP ASE Logins and Database Users*.

You can configure the LDAP server to use these access restrictions:

- Anonymous authentication – all data is visible to any user.
- User name and password authentication – SAP ASE uses the default user name and password from the appropriate file:
  - UNIX, 32-bit – $SYBASE/$SYBASE_OCS/config/libtcl.cfg
  - UNIX, 64-bit – $SYBASE/$SYBASE_OCS/config/libtcl64.cfg
  - Windows – `%SYBASE%\%SYBASE_OCS%\ini\libtcl.cfg`

User name and password authentication properties establish and end a session connection to an LDAP server.

Note: The default user name and password stored in `libtcl.cfg` and passed to the LDAP server for authentication purposes are distinct and different from those used to access SAP ASE. The default user name and password allow access to the LDAP server for administrative tasks.

When an LDAP server is specified in the `libtcl.cfg` or `libtcl64.cfg` file (collectively called `libtcl*.cfg` file), the server information is then accessible only from the LDAP server; SAP ASE ignores the `interfaces` file.

If multiple directory services are supported in a server, the order in which they are searched is specified in `libtcl*.cfg`. You cannot use the **dataserver** command line option to specify the search order.

Multiple Directory Services

Any type of LDAP service, whether it is an actual server or a gateway to other LDAP services, is called an LDAP server.

You can specify multiple directory services for high-availability failover protection in `libtcl*.cfg`. Not every directory service in the list must be an LDAP server.

In the following example, if the connection to *test:389* fails, the connection fails over to the DCE driver with the specified directory information tree (DIT) base. If this also fails, a connection to the LDAP server on *huey:11389* is attempted. Different vendors employ different DIT base formats.

```
[DIRECTORY]
   ldap=libdldap.so ldap://test:389/dc=sap,dc=com
   dce=libddce.so ditbase=/.:/subsys/sap/dataservers
   ldap=libdldap.so ldap://huey:11389/dc=sap,dc=com
```

See the *Open Client Client-Library/C Programmer's Guide* and the *Open Client Client-Library/C Reference Manual*.

LDAP Directory Services Versus the SAP Interfaces File

The LDAP driver implements directory services for use with an LDAP server.

The LDAP infrastructure provides:

- A network-based alternative to the traditional `interfaces` file
- A single, hierarchical view of information, including users, software, resources, networks, files, and so on

Table 1. Interfaces File Versus LDAP Directory Services

`interfaces` File	Directory Services
Platform-specific	Platform-independent
Specific to each installation	Centralized and hierarchical
Contains separate master and query entries	One entry for each server that is accessed by both clients and servers
Cannot store metadata about the server	Stores metadata about the server

LDAP Performance

Performance when using an LDAP server may be slower than when using an `interfaces` file because the LDAP server requires time to make a network connection and retrieve data.

Since this connection is made when SAP ASE is started, changes in performance are seen at login time, if at all. During normal system load, the delay should not be noticeable. During high system load with many connections, especially repeated connections with short duration, the overall performance difference of using an LDAP server versus the traditional `interfaces` file might be noticeable.

Security Features Available in SAP ASE

SAP ASE includes a series of major security features, including identification and authentication control, discretionary access control, role division, accountability, and data confidentiality.

- Identification and authentication controls – ensures that only authorized users can log in to the system. In addition to password-based login authentication, SAP ASE supports external authentication using Kerberos, LDAP, or pluggable authentication modules (PAM).
- Discretionary access controls (DACs) – provides access controls that let object owners restrict access to objects, usually with the **grant** and **revoke** commands. This type of control is dependent upon an object owner’s discretion.

- Division of roles – allows an administrator to grant privileged roles to specified users so only designated users can perform certain tasks. SAP ASE has predefined roles, called “system roles,” such as system administrator and system security officer, and also SAP ASE allows system security officers to define additional roles, called “user-defined roles.”
- Accountability – provides the ability to audit events such as logins, logouts, server start operations, remote procedure calls, accesses to database objects, and all actions performed by a specific user or with a particular role active. SAP ASE also provides a single option to audit a set of server-wide, security-relevant events.
- Confidentiality of data – maintains a confidentiality of data using encryption for client/server communication, available with Kerberos or Secure Sockets Layer (SSL). Inactive data is kept confidential with password-protected database backup.

See *Security Administration Guide > Getting Started with Security Administration in SAP ASE*.

CHAPTER 2 System and Optional Databases

SAP ASE includes a number of system databases, optional SAP-supplied databases that you can install, and the `sybdiag` database, which SAP Product Support may install for diagnostic purposes.

Overview of System Databases

A default SAP ASE installation includes a number of system databases.

- The `master` database
- The `model` database
- The system procedure database, `sybsystemprocs`
- The two-phase commit transaction database, `sybsystemdb`
- The temporary database, `tempdb`

Optionally, you can install:

- The auditing database, `sybsecurity`
- The sample databases, `pubs2` and `pubs3`
- The **dbcc** database, `dbccdb`
- The Job Scheduler database, `sybmgmtdb`

For information about installing the `master`, `model`, `sybsystemprocs`, `tempdb`, and `sybmgmtdb` databases, see the installation guide for your platform. For information about installing `dbccdb`, see *System Administration Guide: Volume 2 > Checking Database Consistency*. For information about using Job Scheduler, see the *Job Scheduler Users Guide*.

The `master`, `model`, `sybsystemdb`, and temporary databases reside on the master device which is named during installation. The `master` database must be contained entirely on the master device and cannot be expanded onto any other device. Create all other databases and user objects on other devices.

Warning! Do not store user databases on the master device; doing so makes it difficult to recover both the system databases and any user databases stored on the master device.

Install the `sybsecurity` and `sybmgmtdb` databases on their own devices and segment. See the installation documentation for your platform.

Install the `sybsystemprocs` database on a device of your choice. You may want to modify the installation scripts for `pubs2` and `pubs3` to share the device you create for `sybsystemprocs`.

Use the `installjsdb` script (located in `$SYBASE/ASE-15_0/scripts`) to install the `sybmgmtdb` database. `installjsdb` looks for a device named `sybmgmtdev` on which to create the `sybmgmtdb` database and its accompanying tables and stored procedures. If the `sybmgmtdb` database already exists, `installjsdb` creates the Job Scheduler tables and stored procedures in the existing database. If `installjsdb` cannot find either a `sybmgmtdev` device or a `sybmgmtdb` database, it creates `sybmgmtdb` on the master device. However, SAP strongly recommends that you remove the `sybmgmtdb` database from the master device.

The `installpubs2` and the `installpubs3` scripts do not specify a device in their **create database** statement, so they are created on the default device. During installation, the master device is the default device. To change this, edit the scripts.

The master Database

The `master` database controls the operation of SAP ASE and stores information about all user databases and their associated database devices.

Table 2. Information About the master Database Stores

Information	System table
User accounts	`syslogins`
Remote user accounts	`sysremotelogins`
Remote servers that this server can interact with	`sysservers`
Ongoing processes	`sysprocesses`
Configurable environment variables	`sysconfigures`
System error messages	`sysmessages`
Databases on SAP ASE	`sysdatabases`
Storage space allocated to each database	`sysusages`
Tapes and disks mounted on the system	`sysdevices`
Active locks	`syslocks`
Character sets	`syscharsets`
Languages	`syslanguages`
Users who hold server-wide roles	`sysloginroles`
Server roles	`syssrvroles`

Information	System table
SAP ASE engines that are online	sysengines

Because the master database stores information about user databases and devices, you must be in the master database to issue the **create database**, **alter database**, **disk init**, **disk refit**, **disk reinit**, and disk mirroring commands.

The minimum size of your master database depends on your server's logical page size. The master database must contain at least 6656 logical pages, so its minimum physical size for each logical page size is:

- 2K page – 13MB
- 4K page – 26MB
- 8K page – 52MB
- 16K page – 104MB

Controlling Object Creation in the master Database

When you install SAP ASE, only a system administrator can create objects in the master database, because the system administrator implicitly becomes the database owner (“dbo”) of any database he or she uses.

Any objects created on the master database should be used only for system administration. Set permissions in master so that most users cannot create objects there.

Warning! Do not place user objects in master. Storing user objects in master causes the transaction log to fill quickly. If the transaction log runs out of space completely, you cannot use **dump transaction** commands to free space in master.

You may also want to use **alter login** to change the default database for users (the database to which a user is connected when he or she logs in). See *Security Administration Guide > Managing SAP ASE Logins and Database Users*.

Create any system procedures in the sybsystemprocs database rather than in master.

Backing Up Master and Keeping Copies of System Tables

To be prepared for hardware or software failure on SAP ASE, perform frequent backups of all databases, and keep copies of system tables.

- Perform frequent backups of the master database and all user databases. See *System Administration Guide: Volume 2 > Restoring the System Databases*.
- Keep a copy (preferably offline) of these system tables: sysusages, sysdatabases, sysdevices, sysloginroles, and syslogins. If you have copies of these scripts, and a hard-disk failure or other disaster makes your database unusable, you can use the recovery procedures described in *System Administration Guide: Volume 2 > Restoring*

the System Databases. If you do not have current copies of your scripts, it is much more difficult to recover SAP ASE when the `master` database is damaged.

model Database

SAP ASE includes the `model` database, which provides a template, or prototype, for new user databases.

Each time a user enters the **create database** command, SAP ASE makes a copy of the `model` database and extends the new database to the size specified by the **create database** command.

Note: New databases must be at least as large as the `model` database.

The `model` database contains the required system tables for each user database. You can modify `model` to customize the structure of newly created databases—everything you do to `model` is reflected in each new database. Some of the changes that system administrators commonly make to `model` are:

- Adding user-defined datatypes, rules, or defaults.
- Adding users who need access to all databases on SAP ASE.
- Granting default privileges, particularly for “guest” accounts.
- Setting database options such as **select into/bulkcopy/pllsort**. These settings are reflected in all new databases. The default settings for these options in `model` is **off**.

Typically, most users do not have permission to modify the `model` database. There is not much point in granting read permission either, since SAP ASE copies its entire contents into each new user database.

The `model` database cannot be larger than `tempdb`. By default, the size of `model` is six allocation units (an allocation unit is 256 logical pages.). You see error message if you increase the size of `model` without making `tempdb` at least as large.

Note: Keep a backup copy of the `model` database, and back up `model` with **dump database** each time you change it. In case of media failure, restore `model` as you would a user database.

sybsystemprocs Database

SAP system procedures are stored in the database `sybsystemprocs`.

When a user in any database executes a system stored procedure (that is, a procedure that has a name beginning with **sp_**), SAP ASE first looks for that procedure in the user’s current database. If there is no procedure there with that name, SAP ASE looks for it in `sybsystemprocs`. If there is no procedure in `sybsystemprocs`, SAP ASE looks for the procedure in `master`.

If the procedure modifies system tables (for example, **sp_adduser** modifies the `sysusers` table), the changes are made in the database from which the procedure was executed.

To change the default permissions on system procedures, modify those permissions in `sybsystemprocs`.

Note: Any time you make changes to `sybsystemprocs`, back up the database.

tempdb database

SAP ASE has a *temporary database*, `tempdb`, that provides a storage area for temporary tables and other temporary working storage needs. The space in `tempdb` is shared among all users of all databases on the server.

The default size of `tempdb` depends on the logical page size for your server, 2, 4, 8, or 16K. Certain activities may make it necessary for you to increase the size of `tempdb`:

- Large temporary tables.
- A lot of activity on temporary tables, which fills up the `tempdb` logs.
- Large or many simultaneous sorts. Subqueries and aggregates with **group by** also cause some `tempdb` activity.

Use **alter database** to increase the size of `tempdb`. `tempdb` is initially created on the master device. You can add space to `tempdb` from the master device or from any other database device.

If you run **update index statistics** against large tables, the command fails with error number 1105 if `tempdb` is not large enough.

You can create and manage multiple temporary databases in addition to the system temporary database, `tempdb`. Multiple temporary databases reduce contention on system catalogs and logs in `tempdb`.

Creating Temporary Tables

No special permissions are required to create temporary tables or to execute commands that may require storage space in the temporary database.

Create temporary tables either by preceding the table name in a **create table** statement with a pound sign (#), or by specifying the name prefix “tempdb..”.

Temporary tables created with a pound sign are accessible only by the current SAP ASE session: users on other sessions cannot access them. These nonsharable, temporary tables are destroyed at the end of each session. The first 13 bytes of the table’s name, including the pound sign (#), must be unique. SAP ASE assigns the names of such tables a 17-byte number suffix. (You can see the suffix by querying `tempdb..sysobjects`.)

Temporary tables created with the “tempdb..” prefix are stored in `tempdb` and can be shared among SAP ASE sessions. SAP ASE does not change the names of temporary tables created this way. The table exists either until you restart SAP ASE or until its owner drops it using **drop table**.

System procedures work on temporary tables, but only if you use them from `tempdb`.

If a stored procedure creates temporary tables, the tables are dropped when the procedure exits. You can also explicitly drop temporary tables before a session ends.

Warning! Do not create temporary tables with the “tempdb..” prefix from inside a stored procedure unless you intend to share those tables among other users and sessions.

Each time you restart SAP ASE, it copies `model` to `tempdb`, which clears the database. You cannot recover temporary tables.

sybsecurity Database

The `sybsecurity` database, which contains the auditing system for SAP ASE, includes these system tables, is installed when you configure SAP ASE for auditing.

`sybsecurity` includes these system tables:

- The `sysaudits_01`, `sysaudits_02`, ... `sysaudits_08` system tables, which contain the audit trail
- The `sysauditoptions` table, which contains rows that describe the global audit options
- All other default system tables that are derived from `model`

See *Security Administration Guide > Auditing*.

sybsystemdb Database

The `sybsystemdb` database stores information about distributed transactions.

SAP ASE versions 12.0 and later can provide transaction coordination services for transactions that are propagated to remote servers using remote procedure calls (RPCs) or Component Integration System (CIS). Information about remote servers participating in distributed transactions is stored in the `syscoordinations` table.

The `sybsystemdb` database also stores information about SYB2PC transactions that use the SAP two-phase commit protocol. The `spt_committab` table, which stores information about and tracks the completion status of each two-phase commit transaction, is stored in the `sybsystemdb` database.

See the *Configuration Guide* for your platform for information about two-phase commit transactions and how to create the `sybsystemdb` database.

sybmgmtdb Database

The `sybmgmtdb` database stores jobs, schedules, scheduled jobs information, and data needed by the internal Job Scheduler task for processing.

`sybmgmtdb` also maintains the output and results from these executed tasks. See the *Job Scheduler Users Guide*.

pubs2 and pubs3 Sample Databases

The `pubs2` and `pubs3` sample databases are provided as a learning tool for SAP ASE.

The `pubs2` sample database is used for most of the examples in the SAP ASE documentation, except for examples, where noted, that use the `pubs3` database. The sample databases are intended for training purposes only, and installing them is optional. Do not install them in an SAP ASE production environment.

For information about installing `pubs2` and `pubs3`, see the installation guide for your platform. For information about the contents of these sample databases, see the *Transact-SQL Users Guide*.

Maintaining the Sample Databases

The sample databases include a "guest" user login that allows access to the database by any authorized SAP ASE user.

The "guest" login has a wide range of privileges in `pubs2` and `pubs3`, including permissions to select, insert, update, and delete user tables. See *Security Administration Guide > Managing SAP ASE Logins and Database Users*.

The size of the `pubs2` and `pubs3` databases are determined by the size of the logical page size for your server; 2, 4, 8, and 16K. If possible, give each new user a clean copy of `pubs2` and `pubs3` so that she or he is not confused by other users' changes. To place `pubs2` or `pubs3` on a specific database device, edit the installation script before installing the database.

If space is a problem, instruct users to issue the **begin transaction** command before updating a sample database. After the user has finished updating one of the sample databases, he or she can issue the **rollback transaction** command to undo the changes.

pubs2 Image Data

SAP ASE includes a script for installing `image` data in the `pubs2` database (`pubs3` does not use the `image` data).

The `image` data consists of six pictures, two each in PICT, TIF, and Sun raster file formats. SAP does not provide any tools for displaying `image` data. You must use the appropriate screen graphics tools to display the images after you extract them from the database.

See the installation documentation for your platform for information about installing the `image` data in `pubs2`.

dbccdb Database

dbcc checkstorage records configuration information for the *target database*, operation activity, and the results of the operation in the `dbccdb` database.

Stored in the database are **dbcc** stored procedures for creating and maintaining `dbccdb` and for generating reports on the results of **dbcc checkstorage** operations. See *System Administration Guide: Volume 2 > Checking Database Consistency*.

sybdiag Database

SAP Product Support may create the `sybdiag` database on your system for debugging purposes. This database holds diagnostic configuration data, and should not be used by customers.

Determining the Version of the Installation Scripts

sp_version lets you determine the current version of the scripts (`installmaster`, `installdbccdb`, and so on) installed on SAP ASE, whether they ran successfully or not, and the length of time they took to complete.

The syntax for **sp_version** is:

```
sp_version [script_file [, “all”]]
```

where:

- *script_file* is the name of the installation script (the default value is NULL).
- **all** reports details about each script, such as the date executed, and the length of time for execution.

If you run **sp_version** without any parameters, it reports on all scripts.

This example describes what installation scripts were run, what time they were run, and what time they finished:

```
sp_version null, 'all'
```

```
Script           Version
Status
-----------      ---------------------------------------------------
---------
-----------
installmaster    15.0/EBF XXXXX/B/Sun_svr4/OS 5.8/asemain/1/32-bit/
OPT/Thu Sep 23 22:12:12 2004
Complete [Started=Sep 24 2004  3:39PM]-[Completed=Sep 24 2004
3:45PM]
```

CHAPTER 3 System Administration for Beginners

New system administrators should be aware of a number of concepts like logical page sizes, "test" servers, installation issues, physical resources, backup and recovery, maintenance, troubleshooting, and the importance of keeping accurate records.

Logical Page Sizes

Database objects are built with logical pages. A databases and any of its related objects must use the same logical page size. That is, you cannot create a server that uses more than one logical page size.

SAP ASE allows you to create master devices and databases with logical page sizes of 2K, 4K, 8K, or 16K, but a given server installation can use only one of these four logical page sizes.

All databases in a server—and all objects in every database—use the same logical page size. For example, all the pages on a server with a logical page size of 4K must be 4K, even though you may not use some pages beyond the initial 2K.

Select the page size when you create the master device with **dataserver** `-z`.

For more information about the **dataserver** command, see the *Utility Guide*. For more information about logical page sizes, see *System Administration Guide: Volume 2 > Configuring Memory*.

Using “Test” Servers

SAP suggests that you install and use a test or development server, then remove it before you create the production server.

Using a test server makes it easier to plan and test different configurations and less stressful to recover from mistakes. It is much easier to learn how to install and administer new features when there is no risk of having to restart a production server or re-create a production database.

SAP suggests that you use a test server from installation or upgrade through configuration. It is in these initial steps that you make some of the most important decisions about your final production system.

Planning Resources

Using a test server helps you plan the final resource requirements for your system and helps you discover resource deficiencies that you might not have anticipated.

In particular, disk resources can have a dramatic effect on the final design of the production system. For example, you may decide that, in the event of a media failure, a particular database requires nonstop recovery. This means you must configure one or more additional database devices to mirror the critical database. Discovering these resource requirements in a test server allows you to change the physical layout of databases and tables without affecting database users.

Use a test server to benchmark both SAP ASE and your applications using different hardware configurations. This allows you to determine the optimal setup for physical resources at both the SAP ASE level and the operating system level before bringing the entire system online for general use.

Achieving Performance Goals

Most performance objectives can be met only by carefully planning a database's design and configuration.

For example, you may discover that the insert and I/O performance of a particular table causes a bottleneck. In this case, the best course of action may be to re-create the table on a dedicated segment and partition the table. Changes of this nature are disruptive to a production system; even changing a configuration parameter may require you to restart SAP ASE.

Considerations When Installing SAP Products

The responsibility for installing SAP ASE and other SAP products is sometimes placed with the system administrator. If installation is one of your responsibilities, use the following pointers to help you in the process. there are several items you should consider, including product compatibility, installation and upgrade issues, third-party software installation, and configuring and testing client connections.

Check Product Compatibility

Before installing new products or upgrading existing products, always read the release bulletin included with the products to understand any compatibility issues that might affect your system.

Compatibility problems can occur between hardware and software, and between different release levels of the same software. Reading the release bulletin in advance can save the time and guesswork of troubleshooting known compatibility problems. Pay particular attention to the lists of known problems that are included in the release bulletin.

Install or Upgrade SAP ASE

Read through the installation guide for your platform before you begin a new installation or upgrade.

You may also want to consult with the operating system administrator to discuss operating system requirements for SAP ASE. These requirements can include the configuration of memory, raw devices, asynchronous I/O, and other features, depending on the platform you use. Many of these tasks must be completed before you begin the installation.

If you are upgrading a server, back up all data (including the `master` database, user databases, triggers, and system procedures) offline before you begin. After upgrading, immediately create a separate, full backup of your data, especially if there are incompatibilities between older dump files and the newer versions.

Install Additional Third-Party Software

SAP ASE generally includes support for the network protocols that are common to your hardware platform. If your network supports additional protocols, install the required protocol support.

As an alternative to the SAP `interfaces` file, you can use a directory service to obtain a server's address and other network information. Directory services are provided by platform or third-party vendors and must be purchased and installed separately from the installation of SAP ASE. See the *Configuration Guide* for your platform for a list of the directory services that SAP ASE currently supports.

Configure and Test Client Connections

A successful client connection depends on the coordination of SAP ASE, the client software, and network products.

If you are using one of the network protocols installed with SAP ASE, see the *Configuration Guide* for your platform for information about testing network connections. If you are using a different network protocol, follow the instructions that are included with the network product. You can also use "ping" utilities that are included with SAP connectivity products to test client connections with SAP ASE. For details about the name and contents of the `interfaces` file, see the *Configuration Guide* for your platform

Allocating Physical Resources

Allocating physical resources is providing SAP ASE the memory, disk space, worker processes, and CPU power required to achieve your performance and recovery goals.

When installing a new server, every system administrator must make decisions about resource utilization. If you upgrade your platform, or if the design of your database system changes, you must also reallocate SAP ASE resources by adding new memory, disk controllers, or

CPUs. Early benchmarking of SAP ASE and your applications can help you identify hardware resource deficiencies that create performance bottlenecks.

See *System Administration Guide: Volume 2 > Overview of Disk Resources* to understand the kinds of disk resources required by SAP ASE. See also *System Administration Guide: Volume 2 > Configuring Memory* and *System Administration Guide: Volume 2 > Managing Mulitprocessor Servers* for information about memory and CPU resources.

The following sections provide helpful pointers in determining physical resource requirements.

Dedicated Versus Shared Servers

Part of planning SAP ASE resources means learning about the resources that are required by other applications running on the same machine.

Generally, system administrators dedicate an entire machine for SAP ASE use, which means that only the operating system and network software consume resources that might otherwise be reserved for SAP ASE. On a shared system, other applications, such as SAP ASE client programs or print servers, run on the same machine as SAP ASE. It can be difficult to calculate the resources available to SAP ASE on a shared system, because the types of applications and their pattern of use may change over time.

It is the system administrator’s responsibility to take into account the resources used by operating systems, client programs, windowing systems, and so forth when configuring resources for SAP ASE. Configure SAP ASE to use only the resources that are available to it. Otherwise, the server may perform poorly or fail to start.

Decision-Support and OLTP Applications

SAP ASE contains many features that optimize performance for OLTP, decision-support, and mixed-workload environments. However, to make optimal use of these features, determine in advance the requirements of your system’s applications.

For mixed-workload systems, list the individual tables that you anticipate will be most heavily used for each type of application; this can help you achieve maximum performance for applications.

Advance Resource Planning

Plan resource usage before you make changes to SAP ASE configuration.

In the case of disk resources, for example, after you initialize and allocate a device to SAP ASE, that device cannot be used for any other purpose (even if SAP ASE never fills the device with data). Likewise, SAP ASE automatically reserves the memory for which it is configured, and this memory cannot be used by any other application.

When planning resource usage:

- For recovery purposes, always place a database’s transaction log on a separate physical device from its data. See *System Administration Guide: Volume 2 > Creating and Managing User Databases*.
- Consider mirroring devices that store mission-critical data. See *System Administration Guide: Volume 2 > Mirroring Database Devices*. If your operating system supports these features, consider using disk arrays and disk mirroring for SAP ASE data.
- If you are working with a test SAP ASE, for convenience, you may find it easier to initialize database devices as operating system files, rather than raw devices. SAP ASE supports either raw partitions or certified file systems for its devices.
- Changing configuration options can affect the way SAP ASE consumes physical resources, especially memory.

Operating System Configuration

After determining the resources that are available to SAP ASE and the resources you require, configure the physical resources at the operating system level.

- If you are using raw partitions, initialize the raw devices to the sizes required by SAP ASE. If you initialize a raw device for SAP ASE, you cannot use that device for any other purpose (for example, to store operating system files). Ask your operating system administrator for assistance in initializing and configuring raw devices to the required sizes.
- Configure the number of network connections. Make sure that the machine on which SAP ASE runs can actually support the number of connections you configure. See your operating system documentation.
- Additional configuration may be required for your operating system and the applications that you use. Read the installation guide for your platform. Also read your client software documentation or consult with your engineers to understand the operating system requirements for your applications.

Backup and Recovery

Making regular backups of your databases is crucial to the integrity of your database system. Although SAP ASE automatically recovers from system crashes (for example, power outages) or server failures, only *you* can recover from data loss caused by media failure.

See these topics in the *System Adminstration Guide: Volume 2* for information about how to develop and implement a backup and recovery plan:

- *Developing a Backup and Recovery Plan*
- *Backing Upa and Restoring User Databases*
- *Restoring the System Databases*
- *Managing Free Space with Thresholds*

Keep Up-To-Date Backups of the master Database

Backing up the `master` database is the most crucial element of any backup and recovery plan.

The `master` database contains details about the structure of your entire database system. It stores information about the SAP ASE databases, devices, and device fragments that make up those databases. Because SAP ASE needs this information for recovery, it is crucial that you maintain an up-to-date backup copy of the `master` database at all times.

To ensure that your backup of `master` is always up to date, back up the database after each command or procedure that affects disks, storage, databases, or segments, including:

- Creating or deleting databases
- Initializing new database devices
- Adding new dump devices
- Using any device mirroring command
- Creating or dropping system stored procedures, if they are stored in `master`
- Creating, dropping, or modifying a segment
- Adding new SAP ASE logins

To back up `master` to a tape device, start **isql** and enter the command:

```
dump database master to "tape_device"
```

where *tape_device* is the name of the tape device (for example, `/dev/rmt0`).

Keep Offline Copies of System Tables

In addition to backing up `master` regularly, keep offline copies of the `sysdatabases`, `sysdevices`, `sysusages`, `sysloginroles`, and `syslogins` system tables

Use the **bcp** utility described in the *Utility Guide* and store a printed copy of the contents of each system table. Create a printed copy by printing the output of:

```
select * from sysusages order by vstart
select * from sysdatabases
select * from sysdevices
select * from sysloginroles
select * from syslogins
```

If you have copies of these tables, and a hard-disk failure or other disaster makes your database unusable, you can use the recovery procedures described in *System Administration Guide: Volume 2 > Restoring the System Databases*.

Also keep copies of all data definition language (DDL) scripts for user objects.

Automating Backup Procedures

Automate your backup procedure using an operating system script or a utility (for example, the UNIX **cron** utility) to perform the necessary backup commands.

You can further automate the procedure by using thresholds, which are discussed in *System Administration Guide: Volume 2 > Managing Free Space with Thresholds*.

Although the commands required to create an automated script vary, depending on the operating system you use, all scripts should accomplish the same basic steps:

1. Start **isql** and dump the transaction log to a holding area (for example, a temporary file).
2. Rename the dump file to a name that contains the dump date, time, and database name.
3. In a history file, record information about the new backup.
4. In a separate file, record any errors that occurred during the dump.
5. Automatically send mail to the system administrator for any error conditions.

Verify Data Consistency Before Backing Up a Database

Your database backups must be consistent and accurate, especially for `master`. If you back up a database that contains internal errors, the errors persist in a restored version of the database.

Use the **dbcc** commands to check a database for errors before backing it up. Always use **dbcc** commands to verify the integrity of a database before dumping it. If **dbcc** detects errors, correct them before dumping the database.

Over time, if you discover few or no errors while running **dbcc**, you may decide that the risk of database corruption is small and that you need to run **dbcc** only occasionally. If the consequences of losing data are too high, continue to run **dbcc** commands each time you back up a database.

Note: For performance considerations, many sites choose to run **dbcc** checks outside of peak hours or on separate servers.

See *System Administration Guide: Volume 2 > Checking Database Consistency*.

Monitor the Log Size

When the transaction log becomes nearly full, it may be impossible to use standard procedures to dump transactions and reclaim space.

The system administrator should monitor the log size and perform regular transaction log dumps (in addition to regular database dumps) to avoid this situation. Set up a threshold stored procedure that notifies you (or dumps the log) when the log reaches a certain capacity. See *System Administration Guide: Volume 2 > Managing Free Space with Thresholds*. SAP also suggests that, to shorten the time required to dump and load the database, dump the transaction log immediately prior to performing a full database dump.

You can also monitor the space used in the log segment manually using **sp_helpsegment**, as described in *System Administration Guide: Volume 2 > Creating and Using Segments*.

Ongoing Maintenance and Troubleshooting

In addition to making regularly scheduled backups, the system administrator generally also performs other maintenance activities throughout the life of a server.

Starting and Stopping SAP ASE

Most system administrators automate the procedure for starting SAP ASE to coincide with the start-up of the server machine. Do this by editing operating system start-up scripts, or by using other operating system procedures.

See the *Configuration Guide* for your platform to determine how to start and stop SAP ASE.

Viewing and Pruning the Error Log

You can use operating system scripts to scan the error log for particular messages and to automatically notify the system administrator when specific errors occur. Checking the error log regularly may help determine whether there are continuing problems of the same nature, or whether a particular database device is likely to fail.

The error log file can grow large over time, since SAP ASE appends informational and status messages to it each time it starts. You can periodically “prune” the log file by opening the file and deleting old records. Keeping the log file to a manageable size saves disk space and makes it easier to locate current errors.

Keeping Records

Keeping records about your SAP ASE system is an important part of your job as a system administrator.

Accurate records of changes and problems that you have encountered can be a valuable reference when you are contacting SAP Technical Support or recovering databases. They can also provide vital information for administrators who manage the SAP ASE system in your absence.

Contact Information

Maintain a list of contact information for yourself as well as the system security officer, operator, and database owners on your system.

Also, record secondary contacts for each role. Make this information available to all SAP ASE users so that the appropriate contacts receive enhancement requests and problem reports.

Configuration Information

If you use script files to create databases and database objects, configure SAP ASE, and store the scripts in a safe place, you can then re-create your entire system in the event of a disaster.

You can also use script files to quickly re-create database systems for evaluation purposes on new hardware platforms. If you use a third-party tool to perform system administration, remember to generate equivalent scripts after performing administration tasks.

Consider recording the following kinds of information:

- Commands used to create databases and database objects (DDL scripts)
- Commands that add new SAP ASE logins and database users
- The current SAP ASE configuration file.
- The names, locations, and sizes of all files and raw devices initialized as database devices

Maintain a dated log of all changes to the SAP ASE configuration. Mark each change with a brief description of when and why you made the change, as well a summary of the end result.

Maintenance Schedules

SAP recommends that you keep a calendar of regularly scheduled maintenance activities.

List any of the procedures you perform at your site:

- Using **dbcc** to check database consistency
- Backing up user and system databases
- Monitoring the space left in transaction logs (if this is not done automatically)
- Dumping the transaction log
- Examining the error log contents for SAP ASE and Backup Server
- Running the **update statistics** command (see *Performance and Tuning Series: Improving Performance with Statistical Analysis > Using the set statistics Commands*)
- Examining auditing information, if the auditing option is installed
- Recompiling stored procedures
- Monitoring resource utilization of the server machine

System Information

Record information about the hardware and operating system on which you run SAP ASE.

This information should include:

- Copies of operating system configuration files or start-up files
- Copies of network configuration files (for example, the `hosts` and `services` files)
- Names and permissions for the SAP ASE executable files and database devices
- Names and locations of the tape devices used for backups

- Copies of operating system scripts or programs for automated backups, starting SAP ASE, or performing other administration activities

Disaster Recovery Plan

Consolidate the basic backup and recovery procedures, the guidelines in *Backup and Recovery*, and your personal experiences in recovering data into a concise list of recovery steps that are tailored to your system.

This list can be useful to both yourself and to other system administrators who may need to recover a production system in the event of an emergency.

Additional Resources

The amount of information for system administrators to learn may seem overwhelming. There are several software tools that can help you learn and facilitate basic administration tasks.

The SAP Control Center simplifies many administration tasks and provides monitoring information. There are also many third-party software packages available designed to help system administrators manage daily maintenance activities.

CHAPTER 4 Managing and Monitoring SAP ASE

SAP ASE includes managing and monitoring tools.

SAP Control Center for SAP ASE

SAP Control Center (SCC) is the preferred tool for managing and monitoring the status and availability of SAP ASE.

SAP Control Center is a Web-based tool for monitoring and administration. See the SCC installation Guide for information about downloading and installing SAP Control Center.

You can use a single instance of SAP Control Center to monitor and manage multiple SAP ASEs.

SAP Control Center supports SAP ASE versions 15.0.2 and later. See the SAP Control Center documentation for a list of supported features.

The SAP Control Center client/server architecture allows multiple clients to monitor all SAP ASEs in an enterprise through one or more SAP Control Center servers. SAP Control Center for SAP ASE provides availability monitoring, historical performance monitoring, and administration capabilities in a scalable Web application that is integrated with management modules for other SAP products. It offers shared, consolidated management of heterogeneous resources from any location, e-mail alerts that provide state- and threshold-based notifications about availability and performance in real time, and intelligent tools for spotting performance and usage trends, all via a thin-client, rich Internet application (RIA) delivered through your Web browser.

Use SAP Control Center for SAP ASE to gather statistics that over time will give you powerful insight into patterns of use and the behavior of databases, devices, caches, and processes on your servers. You can display collected data as tables or graphs. By plotting results over any period of time you choose, from a minute to a year, you can both see the big picture and focus on the particulars. Detailed knowledge of how your servers have performed in the past helps ensure that SAP ASE meets your needs in the future.

See the online help for more information about using SAP Control Center.

Configuration History Tracking

SAP ASE allows you to track the history of configuration changes made to your server. The **sp_confighistory** system procedure manages the history of configuration changes, and stores data about the changes in the `sybsecurity` database.

Configuration properties that are tracked include:

- Server-wide configuration parameters
- Database options
- Data cache and data cache pool properties
- Engine threads
- Changes to the server configuration file.

You must install the the `sybsecurity` auditing database to track these properties.

sp_confighistory displays SAP ASE configuration changes, including which configuration option has been changed, the old and new values, the user who made the change, and when the change was made. SAP ASE stores records of configuration changes in the `sybsecurity` database. Query the `ch_events` view or run **sp_confighistory** to access these records.

For example, the output below includes changes that include enabling auditing and changing the value of **max online engines** from 5 to 7:

```
area type target element oldvalue newvalue mode timestamp username
instanceid
---- ---- ------ ------- -------- -------- ---- --------- --------
----------
AUDIT global auditing  NULL NULL off on NULL Jul 15 2013 2:22PM sa
NULL
SERVER sp_configure NULL max online engines 5 7 static Jul 15 2013
2:23PM sa NULL
```

Configuring SAP ASE to Track Configuration Changes

To install **sp_confighistory**, run the `installsecurity` script.

Install and enable the audit system. See the *Security Administration Guide*.

1. Enable configuration history tracking (requires the sa_role, sso_role, or `manage auditing` if granular permission is enabled):

   ```
   sp_audit "config_history", "all", "all", "on"
   ```

   Note: Issuing **sp_audit** is recorded in the configuration history.

2. Enable auditing:

   ```
   sp_configure 'auditing', 1
   ```

3. Move to the `sybsecurity` database:

```
use sybsecurity
```

4. Create the `ch_events` view:

   ```
   sp_confighistory create_view
   ```

   `ch_events` collects information from all the audit tables, and becomes out of sync if you add or remove audit tables. If this occurs, `ch_events` may not include some configuration history record changes, or you may see error messages 208 (`table not found`) and 4413 (`view unusable`) when you query `ch_events`.

   Use **sp_confighistory create_view** to update `ch_events` when you add or remove audit tables. **sp_confighistory create_view** drops the view if it exists, and creates a new view that corresponds to the current audit tables.

Changes Captured

When configuration history auditing is enabled, SAP ASE captures a number of events.

The `ch_events` view does not record changes if the new value is the same as the old value.

Startup Configuration Changes

If you modify the SAP ASE configuration file while the SAP ASE is shut down, SAP ASE records any changes to the configuration in the `ch_events` table it starts (recording a value of NULL for the `mode` and `username` values for these changes).

Reading the Configuration File

`ch_events` records the event when you read, write, verify, and restore the configuration file, but does not record the configuration value changes. For example, if you change the value for **number of user connections** and then issue:

```
sp_configure "configuration file", 0, "read", "srv.config"
```

`ch_events` records that you read the configuration file, but does not record the configuration value changes.

sp_configure Changes

SAP ASE records all changes made by **sp_configure**, including:

- Name of the configuration parameter
- Old configuration value
- New configuration value
- Whether the parameter is dynamic or static
- Timestamp of the date and time the change was made
- Login of the user making the change

Configuration changes caused by reading from the configuration file are not recorded. That is, SAP ASE records the reading, writing, verifying, and restoring operations, but does not record the configuration changes caused by a reading operation. You can also change configuration values by reading a different or manually modified configuration file. Although SAP ASE records that it read the file, it does not record the individual parameter changes.

Changes to Server Options
`ch_events` records all changes made by **sp_serveroption**, including:

- Name of affected server
- Name of option that was changed
- Old option value
- New option value
- Timestamp of the date and time the change was made
- Login of the user making the change

Changes to Database Options
`ch_events` records these changes made by **sp_dboption**:

- Name of affected database
- Name of option that was changed
- Old option value
- New option value
- Timestamp of the date and time the change was made
- Login of the user making the change

Changes to Cache Configuration
`ch_events` records all changes made by **sp_cacheconfig** and **sp_poolconfig** to cache configurations.

Recorded changes from **sp_cacheconfig** include:

- Name of affected cache
- Old cache size
- New cache size
- Attribute (cache type, cache replacement policy, partition number), if applicable
- Timestamp of the date and time the change was made
- Login of the user making the change
- (Cluster Edition only) Instance to which this change applies

Recorded changes from **sp_poolconfig** include:

- Name of affected cache
- Configuration pool
- Old cache pool size
- New cache pool size

- Name of changed attribute (affected pool, wash size, asynchronous prefetch (APF) percentage)
- Timestamp of the date and time the change was made
- Login of the user making the change
- (Cluster Edition only) Instance to which this change applies

Changes to Trace flags and Switches

`ch_events` records changes to **dbcc traceon | off** and **set switch on | off**.

Recorded changes from **dbcc traceon | off** include:

- Trace flag affected
- Session ID
- Old trace flag state (**on** or **off**)
- New trace flag state (**on** or **off**)
- Timestamp of the date and time the change was made
- Login of the user making the change
- (Cluster Edition only) Instance to which this change applies

Recorded changes from **set switch on | off** include:

- Switch number affected
- Old switch state (**on** or **off**)
- New switch state (**on** or **off**)
- Name of changed attribute (server-wide or session-specific, with override, with no_info)
- Timestamp of the date and time the change was made
- Login of the user making the change

Changes to Number of Engines

The changes tracked by **create thread pool**, **alter thread pool**, and **drop thread pool** include:

- Name of pool name affected
- Old pool size
- New pool size
- Name of changed attribute (new pool name, idle timeout, and so on)
- Timestamp of the date and time the change was made
- Login of the user making the change
- (Cluster Edition only) Instance to which this change applies

SAP ASE Startup and Shutdown Events

`ch_events` records this information for startup, **shutdown**, **shutdown with nowait**, and abrupt (unscheduled) shutdown events for SAP ASE and instances from the Cluster Edition:

- Name of the action (startup, **shutdown**, **shutdown with nowait**, abrupt shutdown).

- Time spent waiting for a shutdown. Not applicable for **shutdown with no_wait**.
- Name of the host on which the server starts.
- Timestamp of the date and time the change was made.
- Login of the user making the change.
- (Cluster Edition only) Instance to which this change applies.

Note: Because `ch_events` records the shutdown when you issue the **shutdown** command, `ch_events` may record multiple shutdowns during a polite shutdown if you issue the command more than once.

Enabling or Disabling Auditing

`ch_events` records this information about tracking, global auditing, and configuration history auditing:

- Name of the action (**enable** or **disable**)
- Timestamp of the date and time the change was made
- Login of the user making the change

Querying ch_events

SAP ASE includes the `ch_events` view as part of the `sybsecurity` database.

`ch_events` presents configuration change history data in an easy to read format. You can query `ch_events` directly, or use the **sp_confighistory** system procedure to generate reports on configuration changes history. Either method provides the same information.

Using the **select** command provides the flexibility of the Transact-SQL™ language to qualify your result set (you must first move to the `sybsecurity` database before selecting from the `ch_events` view). **sp_confighistory** provides a more streamlined result set.

For example, if you make these configuration changes in SAP ASE:

```
sp_dboption sybsystemprocs, "delayed commit", false
sp_cacheconfig pub_cache, '10M'
sp_cacheconfig pub_log_cache, '2000K', logonly
```

Then shut down and restart the server, **sp_confighistory** returns:

```
sp_confighistory

 area     type            target         element             oldvalue
         newvalue    mode timestamp           username instanceid
 -------- --------------- -------------- --------------------
--------
         ----------- ---- ------------------- -------- ----------
 AUDIT    global auditing NULL           NULL                off
         on          NULL Aug 22 2013 11:56AM sa             NULL
 DATABASE sp_dboption     sybsystemprocs delayed commit      true
         false       NULL Aug 22 2013  3:16PM sa             NULL
 CACHE    sp_cacheconfig  pub_cache      NULL                10240
         not changed NULL Aug 22 2013  3:18PM sa             NULL
```

```
 CACHE    sp_cacheconfig  pub_log_cache  cache type: logonly 2000
        not changed NULL Aug 22 2013  3:19PM sa             NULL
 SUSD    shutdown        NULL           NULL                NULL
        NULL        NULL Aug 22 2013  3:49PM sa             NULL
 SUSD    startup         NULL           tigger              NULL
        NULL        NULL Aug 22 2013  3:50PM NULL           NULL
```

Include the date with **sp_confighistory** to select the changes over a period of time. This example shows all changes made after August 23, 2013:

```
sp_confighistory "Aug 23 2013"

 area type     target element oldvalue newvalue mode timestamp
        username instanceid
 ---- -------- ------ ------- -------- -------- ----
-------------------
        -------- ----------
 SUSD shutdown NULL   NULL    NULL     NULL     NULL Aug 23 2013
9:00AM
        sa            NULL
 SUSD startup  NULL   tigger  NULL     NULL     NULL Aug 23 2013
10:38AM
        NULL          NULL
```

Issuing **select** provides this result set:

```
use sybsecurity
go
select * from ch_events
go

area       type               target           element
oldvalue    newvalue    mode
                 timestamp               username       instanceid
----------  -------------------  ----------------
-----------------    ----------  ---------    --------------
                   ----------------------   -----------
---------------
 AUDIT     global auditing     NULL            NULL
off         on          NULL
                   Aug 22 2013 11:56AM       sa             NULL
 DATABASE  sp_dboption        sybsystemprocs   delayed commit
true        false       NULL
                   Aug 22 2013  3:16PM       sa             NULL
 CACHE     sp_cacheconfig     pub_cache        NULL
10240       not changed  NULL
                   Aug 22 2013  3:18PM       sa             NULL
 CACHE     sp_cacheconfig     pub_log_cache    cache type: logonly
2000        not changed  NULL
                   Aug 22 2013  3:19PM       sa             NULL
 SUSD      shutdown           NULL            NULL
NULL        NULL        NULL
                   Aug 22 2013  3:49PM       sa             NULL
 SUSD      startup            NULL            tiger
NULL        NULL        NULL
                   Aug 22 2013  3:50PM       NULL           NULL
```

Include the **last** parameter to see the last items changed:

```
sp_confighistory last

 area type    target element oldvalue newvalue mode timestamp
         username instanceid
 ---- ------- ------ ------- -------- -------- ----
--------------------
         -------- ----------
 SUSD startup NULL   tigger  NULL     NULL     NULL Aug 22 2013  3:50PM
         NULL           NULL
```

CHAPTER 5 Setting Configuration Parameters

A configuration parameter is a user-definable setting that you set with **sp_configure**. Configuration parameters are used for a wide range of services, from basic to specific server operations, and for performance tuning. SAP ASE supplies default values for all configuration parameters.

Use configuration parameters to tailor SAP ASE for an installation's particular needs. Determine which configuration parameters you should reset to optimize server performance.

Warning! Change configuration parameters with caution. Arbitrary changes in parameter values can adversely affect SAP ASE performance and other aspects of server operation.

SAP ASE stores the values of configuration parameters in a configuration file, which is an ASCII text file.

When you install a new SAP ASE, your parameters are set to the default configuration; the default name of the file is `server_name.cfg`, and the default location of the file is the SAP ASE home directory ($SYBASE_ASE). Each time you modify a configuration parameter, SAP ASE creates a copy of the outdated configuration file, using the naming convention `server_name.001`, `server_name.002`, `server_name.003`…`server_name.999`. SAP ASE writes the new values to the file `server_name.cfg` or to a file name you specify at start-up.

Modifying Configuration Parameters

You can modify configuration parameters using a variety of methods.

These methods include:

- Executing **sp_configure** with the appropriate parameters and values,
- Editing your configuration file and then invoking **sp_configure** with the **configuration file** option, or
- Specifying the name of a configuration file at start-up.

Configuration parameters are either dynamic or static. Dynamic parameters take effect as soon as you execute **sp_configure**. Static parameters require memory to be reallocated, so they take effect only after you have restarted SAP ASE. The description of each parameter in this chapter indicates whether it is static or dynamic.

SAP ASE writes the new value to the system table `sysconfigures` and to the configuration file when you change the value. The current configuration file and `sysconfigures` reflect configured values, not run values. The system table `syscurconfigs` reflects current run values of configuration parameters.

Required Roles for Modifying Configuration Parameters

The SAP ASE role you have determines how you execute **sp_configure**.

- Any user can execute **sp_configure** to display information about parameters and their current values.
- Only a system administrator or a system security officer can execute **sp_configure** to modify configuration parameters.
- Only a system security officer can execute **sp_configure** to modify values for:
  - **allow procedure grouping**
  - **allow remote access**
  - **allow sendmsg**
  - **allow updates to system tables**
  - **auditing**
  - **audit queue size**
  - **check password for digit**
  - **current audit table**
  - **enable ldap user auth**
  - **enable pam user auth**
  - **enable ssl**
  - **log audit logon failure**
  - **log audit logon success**
  - **maximum failed logins**
  - **minimum password length**
  - **msg confidentiality reqd**
  - **msg integrity reqd**
  - **secure default login**
  - **select on syscomments.text**
  - **SQL Perfmon Integration**
  - **syb_sendmsg port number**
  - **suspended audit when device full**
  - **systemwide password expiration**
  - **unified login required**
  - **use security services**

Unit Specification Using sp_configure

sp_configure allows you to specify the value for configuration parameters in unit specifiers.

The unit specifiers are **p** or **P** for pages, **m** or **M** for megabytes, and **g** or **G** for gigabytes. If you do not specify a unit, and you are configuring a parameter that controls memory, SAP ASE uses the logical page size for the basic unit.

Note: When you are configuring memory-related parameters, use only the **P** (page size) parameter for your unit specification. If you use any other parameter to configure memory related parameters, SAP ASE may issue an arithmetic overflow error message.

The syntax to indicate a particular unit specification is:

```
sp_configure "parameter name", 0, "p|P|k|K|m|M|g|G"
```

You must include the “0” as a placeholder.

You can use this unit specification to configure any parameter. For example, when setting **number of locks** to 1024 you can enter:

```
sp_configure "number of locks", 1024
```

or:

```
sp_configure "number of locks", 0, "1K"
```

This functionality does not change the way in which SAP ASE reports **sp_configure** output.

Global Versus Session Settings

Some SAP ASE global settings are similar to session-level settings. For example, both **sp_configure** and **sp_passwordpolicy** include a **minimum password length** parameter.

If a global and session-level parameters configure the same setting, only the session level setting applies to the login (login level settings override global settings).

For example, if you set this global configuration configuration parameter:

```
sp_configure 'minimum password length' 12
```

And you attempt to add user joe to the server with a insufficiently long password:

```
sp_adduser joe, joejoe
```

The command fails:

```
Msg 10317, Level 14, State 1:
Procedure 'sp_password', Line 118:
The specified password is too short. Passwords must be at least 12
character(s) long.
Msg 17720, Level 16, State 1:
```

```
Procedure 'sp_password', Line 128:
Error:  Unable to set the Password.
```

However, if you set **sp_passwordpolicy** to a shorter length, SAP ASE requires a shorter password length for this session:

```
sp_passwordpolicy 'set', 'minimum password length', 6
```

The command succeeds:

```
sp_addlogin joe, joejoe
```

```
Password correctly set.
Account unlocked.
New login created.
(return status = 0)
```

SAP ASE tests all configuration limits independently, and the login attempt must pass all the applicable limits to succeed.

Getting Help Information on Configuration Parameters

Use either **sp_helpconfig** or **sp_configure** to display information about a particular configuration parameter.

For example:

```
sp_helpconfig "number of open"
```

```
Configuration option is not unique.
option_name                    config_value  run_value
------------------------------ ------------ -----------
number of open databases                 12          12
number of open indexes                  500         500
number of open objects                  500         500
```

```
sp_helpconfig "number of open indexes"
```

```
number of open indexes sets the maximum number of indexes that can be
open at
one time on SQL Server. The default value is 500.
Minimum Value Maximum Value Default Value Current Value Memory Used
------------- ------------- ------------- ------------- -----------
          100    2147483647           500           500         208
```

```
sp_configure "number of open indexes"
```

```
Parameter Name           Default  Memory Used  Config Value  Run Value
-----------------------  -------  -----------  ------------  -------
--
number of open indexes       500          208           500        500
```

See *System Administration Guide: Volume 2 > Configuring Memory*.

Using sp_configure

sp_configure displays and resets configuration parameters.

You can restrict the number of parameters that **sp_configure** shows by using **sp_displaylevel** to set your display level to one of:

- Basic
- Intermediate
- Comprehensive

For information about **sp_displaylevel**, see the *Reference Manual: Procedures*.

In this table, the "Effect" column assumes that your display level is set to "comprehensive."

Command	Effect
sp_configure	Displays all configuration parameters by group, their current values, their default values, the value to which they have most recently been set, and the amount of memory used by this particular setting.
sp_configure "*parameter*"	Displays current value, default value, most recently changed value, and amount of memory used by the specified parameter.
sp_configure "*parameter*", *value*	Resets *parameter* to *value*.
sp_configure "*parameter*", 0, "default"	Resets the specified parameter to its default value.
sp_configure "*group_name*"	Displays all configuration parameters in *group_name*, their current values, their default values, the values to which they were recently set, and the amount of memory used by each.
sp_configure "configuration file", 0, "*sub_command*", "*file_name*"	Sets configuration parameters from the configuration file.

sp_configure Syntax Elements

sp_configure includes a variety of variables to help you configure SAP ASE.

- *parameter* – is any valid SAP ASE configuration parameter or parameter substring.
- *value* – is any integer within the valid range for that parameter. (See the descriptions of the individual parameters for valid range information.) Parameters that work as toggles have only two valid values: 1 (on) and 0 (off).
- *group_name* – is the name of any group in the parameter hierarchy.

sp_configure parses each parameter (and parameter name fragment) as "*%parameter%*". A string that does not uniquely identify a particular parameter returns values for all parameters matching the string.

The following example returns values for all configuration parameters that include "lock," such as **lock shared memory**, **number of locks**, **lock promotion HWM**, **server clock tick length**, **print deadlock information**, and **deadlock retries**:

```
sp_configure "lock"
```

Note: If you attempt to set a parameter value with a nonunique parameter name fragment, **sp_configure** returns the current values for all parameters matching the fragment and asks you to specify a unique parameter name.

Issue sp_configure with the Configuration File

Configure SAP ASE either interactively, by using **sp_configure**, or noninteractively, by instructing SAP ASE to read values from an edited or restored version of the configuration file.

By making your changes from the configuration file, you can:

- Replicate a specific configuration across multiple servers by using the same configuration file.
- Use a configuration file as a baseline for testing configuration values on your server.
- Use a configuration file to perform validation checking on parameter values before actually setting the values.
- Create multiple configuration files and switch between them as your resource needs change.

Configuration File Naming Recommendations

To work with a configuration file that has a name other than the default name, keep the *server_name* part of the file name, and include at least one alphabetic character in the extension (for example `my_server.A001`).

Alternatively, you can change the *server_name* part of the file name (for example, `A_my_server.001`). Doing this avoids confusion with the backup configuration files generated by SAP ASE when you modify a parameter.

Read or Write the Configuration File with sp_configure

sp_configure includes syntax options for using configuration files.

The syntax is:

```
sp_configure "configuration file", 0, "subcommand", "file_name"
```

where:

- **"configuration file"** – including quotes, specifies that this command uses the configuration file.
- **0** – required—for backward compatibility—after the **configuration file** parameter.
- **"*subcommand*"** – is one of:

  - **write** – creates a file named *file_name* with the current configuration. If *file_name* already exists, a message is written to the error log; the existing file is renamed using the convention *server_name*.`001`, *server_name.002*, and so on. If you have changed a static parameter, but you have not restarted your server, **write** displays the currently running value for that parameter. If you do not specify a directory with *file_name*, the file is written to the directory from which SAP ASE was started.
  - **read** – performs validation checking on values contained in *file_name* and reads those values that pass validation into the server. If any parameters are missing from *file_name*, the current values for those parameters are used.

    If the value of a static parameter in *file_name* is different from its current running value, **read** fails and a message is printed. However, validation is still performed on the values in *file_name*.
  - **verify** – performs validation checking on the values in *file_name*. This is useful if you have edited the configuration file, as it prevents you from attempting to configure your server with invalid configuration values.
  - **restore** – creates *file_name* with the most recently configured values. If you have configured static parameters to new values, this subcommand writes the configured, not the currently running, values to the file. This is useful if all copies of the configuration file have been lost and you must generate a new copy. If you do not specify a directory with *file_name*, the file is written to the directory from which SAP ASE was started.
- *file_name* – specifies the configuration file to use in conjunction with any *subcommand*. If you do not specify a directory as part of the file name, the directory where SAP ASE was started is used.

Examples

- Example 1 – performs validation checking on the values in the file `srv.config` and reads the parameters that pass validation into the server. Current run values are substituted for values that do not pass validation checking:

  ```
  sp_configure "configuration file", 0, "read", "srv.config"
  ```
- Example 2 – creates the file `my_server.config` and writes the current configuration values the server is using to that file:

  ```
  sp_configure "configuration file", 0, "write", "my_server.config"
  ```

Edit the Configuration File

The configuration file is an ASCII file that you can edit with any text editor that can save files in ASCII format.

The syntax for each parameter is:

```
parameter_name={value | DEFAULT}
```

where:

- *parameter_name* – is the name of the parameter you want to specify.

- *value* – is the numeric value for set *parameter_name*.
- “DEFAULT” – specifies that you want to use the default value for *parameter_name*.

Examples

- Example 1 – specifies that the transaction can retry its attempt to acquire a lock one time when deadlocking occurs during an index **page split** or shrink:
  ```
  deadlock retries = 1
  ```
- Example 2 – specifies that the default value for the parameter **cpu accounting flush interval** should be used:
  ```
  cpu accounting flush interval=DEFAULT
  ```

When you edit a configuration file, your edits are not validated until you check the file using the **verify** option, read the file with the **read** option, or restart SAP ASE with that configuration file.

If all your configuration files are lost or corrupted, you can re-create one from a running server by using the **restore** subcommand and specifying a name for the new file. The parameters in the new file are set to the values with which your server is currently running.

Permissions for Configuration Files

Configuration files are nonencrypted ASCII text files.

By default, configuration files are created with read and write permissions set for the file owner, and read permission set for all other users. If you created the configuration file at the operating system level, you are the file owner; if you created the configuration file from SAP ASE, using the **write** or **restore** parameter, the file owner is the user who started SAP ASE. Usually, this is the user “sybase.” To restrict access to configuration files, use your operating system’s file permission command to set read, write, and execute permissions as appropriate.

Note: You must set permissions accordingly on *each* configuration file created.

Backing Up Configuration Files

Configuration files are not automatically backed up when you back up the `master` database. They are operating system files—back them up in the same way you back up other operating system files.

Verify the Name of the Configuration File Currently in Use

Due to space limitations, **sp_configure** output truncates the name of the configuration file.

To see the full name of the configuration file, use:

```
select s1.value2
from syscurconfigs s1, sysconfigures s2
where s1.config = s2.config
and s2.name = "configuration file"
```

Start SAP ASE Using a Configuration File

By default, SAP ASE reads the configuration file `server_name.cfg` in the start-up directory when it starts. If this file does not exist, it creates a new file and uses SAP ASE defaults for all values.

You can start SAP ASE with a specified configuration file. For more information, see the *Utility Guide*.

If the configuration file you specify does not exist, SAP ASE prints an error message and does not start.

If the command is successful, the file `server_name.bak` is created. This file contains the configuration values stored in `sysconfigures` prior to the time `sysconfigures` was updated with the values read in from the configuration file you specified. This file is overwritten with each subsequent start-up.

Configuration File Errors

When there are errors in the configuration file, SAP ASE may not start, or may use default values.

SAP ASE uses default values if:

- There are illegal values. For example, if a parameter requires a numeric value, and the configuration file contains a character string, SAP ASE uses the default value.
- Values are below the minimum allowable value.

The Parameter Hierarchy

Configuration parameters are grouped according to the area of SAP ASE behavior they affect. This makes it easier to identify all parameters that you might need to tune to improve a particular area of SAP ASE performance.

Although each parameter has a primary group to which it belongs, many have secondary groups to which they also belong. For example, **number of remote connections** belongs primarily to the network communication group, but it also belongs secondarily to the memory use group. This reflects the fact that some parameters have implications for a number of areas of SAP ASE behavior. **sp_configure** displays parameters in all groups to which they belong.

Table 3. Configuration Groups

Parameter Group	Configures SAP ASE for
Backup/Recovery	Backing up and recovering data
Cache manager	Data and procedure caches
Component Integration Services administration	Component Integration Services

Parameter Group	Configures SAP ASE for
DTM administration	Distributed transaction management (DTM) facilities
Diagnostics	Diagnostic principles
Disk I/O	Disk I/O
Error log	Error log, and the logging of SAP ASE events to the Windows event log
Extended stored procedures	The behavior of extended stored procedures (ESPs).
General information	Basic system administration
Java services	Memory for Java in SAP ASE See *Java in SAP Adaptive Server Enterprise* for complete information about Java in the database. If you use method calls to JDBC, you may need to increase the size of the execution stack available to the user.
Languages	Languages, sort orders, and character sets
Lock manager	Locking
Memory use	Memory consumption
Metadata caches	Setting the metadata cache size for frequently used system catalog information. The metadata cache is a reserved area of memory used for tracking information on databases, indexes, or objects. The greater the number of open databases, indexes, or objects, the larger the metadata cache size. For a discussion of metadata caches in a memory-usage context, see *System Administration Guide: Volume 2 > Configuring Memory*.
Monitoring	Collecting monitoring information. By default, SAP ASE does not collect monitoring information. See *Performance and Tuning Guide: Monitoring and Analyzing > Monitoring Tables*.
Network communication	Communication between SAP ASE and remote servers, and between SAP ASE and client programs
O/S resources	Use of operating system resources
Physical memory	Your machine's physical memory resources
Processors	Processors in an SMP environment
Query Tuning	Query optimization
RepAgent thread administration	Replication via Replication Server

Parameter Group	Configures SAP ASE for
Shared Disk Cluster	Cluster Edition
SQL Server administration	General SAP ASE administration
Security related	Security-related features
Unicode	Unicode-related features
User environment	User environments

The syntax for displaying all groups and their associated parameters, and the current values for the parameters, is:

```
sp_configure
```

Note: The number of parameters returned by **sp_configure** depends on the value to which you have your display level set.

The following is the syntax for displaying a particular group and its associated parameter:

```
sp_configure "group_name"
```

For example, to display the disk I/O group, enter:

```
sp_configure "Disk I/O"
```

```
Group: Disk I/O
Parameter Name           Default Memory Used Config Value Run Value
unit         type
--------------           ------- ----------- ------------ ---------
------       -------------
allow sql server async i/o     1           0            1         1
switch       static
diable disk mirroring          1          0            1         1
switch       static
disk i/o structures          256          0          256       256
number       dynamic
number of devices             10          0           10        10
number       dynamic
number of large I/O buffers    6      12352            6         6
number       dynamic
page utilization percent      95          0           95        95
percent      dynamic
```

Note: If the server uses a case-insensitive sort order, **sp_configure** with no parameters returns a list of all configuration parameters and groups in alphabetical order with no grouping displayed.

User-Defined Subsets of the Parameter Hierarchy: Display Levels

Depending on how you use SAP ASE, you may need to adjust some parameters more frequently than others. It may be easier to work with a subset of parameters.

The default display level is comprehensive. When you set your display level, the setting persists across multiple sessions. However, you can reset it at any time.

- Basic – shows only the most basic parameters, and is appropriate for general server tuning.
- Intermediate – includes parameters that are somewhat more complex, in addition to the Basic parameters.
- Comprehensive – includes all the parameters, including the most complex ones. This level is appropriate for users doing highly detailed server tuning.

The syntax for showing your current display level is:

```
sp_displaylevel
```

To set the display level, use the following, where *user_name* is your SAP ASE login name:

```
sp_displaylevel user_name[, basic | intermediate | comprehensive]
```

The Effect of the Display Level on sp_configure Output

If your display level is set to either **basic** or **intermediate**, **sp_configure** returns only a subset of the parameters that are returned when your display level is set to **comprehensive**.

For example, if your display level is set to **intermediate**, and you want to see the parameters in the languages group, enter:

```
sp_configure "Languages"
```

The output looks like this:

```
sp_configure
Group: Languages

Parameter Name         Default Memory Used Config Value Run Value Un
it Type
---------------        ------- ----------- ------------ --------- --
-- ----
default character set    1          0            1          1     i
d   static
default language id     0        0           0        0     id   dyna
. . .
```

This represents only a subset of the parameters in the languages group; some language parameters appear only when your display level is **comprehensive**.

Performance Tuning with sp_configure and sp_sysmon

sp_sysmon monitors SAP ASE performance and generates statistical information that describes the behavior of your SAP ASE system.

You can run **sp_sysmon** before and after using **sp_configure** to adjust configuration parameters. The output gives you a basis for performance tuning and allows you to observe the results of configuration changes.

See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

Using Configuration Parameters in a Clustered Environment

In the Cluster Edition, SAP supports both cluster-wide and instance-specific configuration.

Cluster-wide configuration parameters are applied to all instances in the cluster. Local configuration parameters are applied only to a specified instance.

- Local configuration overrides cluster-wide configuration.
- If an instance-specific configuration has not been applied, the cluster-wide configuration applies.
- Some parameters, such as **default character set id**, cannot be applied to a specific instance. These parameters can only be used over an entire cluster.

The cluster configuration file includes an instance-specific configuration block. Parameter settings in the instance-specific block override cluster-wide settings. For example:

```
max online engines = DEFAULT

[Instance:ase1]
max online engines = 5
[Instance:ase2]
max online engines = 3
```

See the *Clusters Users Guide*.

sp_configure Output

The values sp_configure prints vary, depending on your platform and on what values you have already changed.

This output shows the type of information **sp_configure** prints if your display level is **comprehensive**, and you execute **sp_configure** with no parameters:

```
sp_configure
Group: Configuration Options

Group: Backup/Recovery

Parameter Name        Default Memory Used Config Value Run Value Uni
t Type
```

```
--------------        ------- ----------- ------------ --------- ---
- ----
allow remote access       1         0         1        1    switch dyn
print recovery info       0         0         0        0    switch dyn
recovery interval in m      5          0          5          5    min
utes dyn
...
```

Note: All configuration groups and parameters appears in output if your display level is set to "comprehensive."

where:

- The "Default" column displays the default value. If you do not explicitly reconfigure a parameter, it retains its default value.
- "Memory Used" shows the amount of memory, in kilobytes, used by the parameter at its current value. Some related parameters draw from the same memory pool. For instance, the memory used for **stack size** and **stack guard size** is already accounted for in the memory used for **number of user connections**. If you added the memory used by each of these parameters separately, the sume is more than the amount actually used. Parameters that "share" memory with other parameters are marked with a hash mark ("#").
- "Config Value" displays the most recent value to which the configuration parameter has been set. When you execute **sp_configure** to modify a dynamic parameter:
  - The configuration and run values are updated.
  - The configuration file is updated.
  - The change takes effect immediately.

  When you modify a static parameter:
  - The configuration value is updated.
  - The configuration file is updated.
  - The change takes effect only when you restart SAP ASE.
- "Run Value" displays the value SAP ASE is currently using. It changes when you modify a dynamic parameter's value and, for static parameters, after you restart SAP ASE.
- "Unit" displays the unit value of the configuration parameter. SAP ASE displays information in the following units:

Name of Unit	Unit Description
number	Number of items.
clock ticks	Number of clock ticks.
microseconds	Number of microseconds.
milliseconds	Number of millisecond.s
seconds	Number of seconds.
minutes	Number of minutes.

Name of Unit	Unit Description
hours	Number of hours.
bytes	Number of bytes.
days	Number of days.
kilobytes	Number of kilobytes.
megabytes	Number of megabytes.
memory pages (2K)	Number of 2K memory pages.
virtual pages (2K)	Number of 2K virtual pages.
logical pages	Number of logical pages. This value depends on the logical page size your server is using: 2, 4, 8, or 16K.
percent	Value of the configured parameter as a percentage.
ratio	Value of the configured parameter as a ratio.
switch	Value of the parameter is either TRUE (the parameter is turned on), or FALSE.
id	ID of the configured parameter you are investigating.
name	Character string name assigned to the run or configure value of the parameter. For example, “binary” appears under the “Run Value or “Config Value” column for the output of **sp_configure "lock scheme"**.
row	Number of rows

- “Type” displays whether the configuration option is static or dynamic. Changes to static parameters require that you restart SAP ASE for the changes to take effect. Changes to dynamic parameters take effect immediately without having to restart SAP ASE.

sysconfigures and syscurconfigs Tables

The report shown by **sp_configure** is constructed mainly from the `master..sysconfigures` and `master..syscurconfigs` system tables, with additional information provided from `sysattributes`, `sysdevices`, and other system tables.

The `value` column in the `sysconfigures` table records the last value set from **sp_configure** or the configuration file; the `value` column in `syscurconfigs` stores the value currently in use. For dynamic parameters, the two values match; for static parameters, which require a restart of the server to take effect, the two values are different if the values have been changed since SAP ASE was last started. The values may also be different when the default values are used. In this case, `sysconfigures` stores 0, and `syscurconfigs` stores the value that SAP ASE computes and uses.

sp_configure performs a join on `sysconfigures` and `syscurconfigs` to display the values reported by **sp_configure**.

Example syscurconfigs and sysconfigures Query

You can query `sysconfigures` and `syscurconfigs` to organize information the way you want.

For example, **sp_configure** without any arguments lists the memory used for configuration parameters, but does not list minimum and maximum values. Use this query to get a complete list of current memory usage, as well as minimum, maximum, and default values:

```
select b.name, memory_used, minimum_value, maximum_value, defvalue
from master.dbo.sysconfigures b,
master.dbo.syscurconfigs c
where b.config *= c.config and parent != 19
and b.config > 100
```

Named Cache Configuration Parameter Group

The Named Cache configuration parameter group provides details for named caches.

- **cache size** – size of the cache. By default SAP ASE creates 8MB caches. Change this parameter dynamically with **sp_cacheconfig**, or change the value in the server configuration file to have the change take place after the next server restart.
- **cache status** – status of the cache. One of `default data cache`, `log on ly`, `mixed`, or `in-memory storage`. The default data cache must have a cache status of `default data cache`, and cannot be changed. **cache status** for named caches can be `log only`, `mixed`, or, for in-memory databases, `in-memory storage` (you cannot change the **cache status** for in-memory databases).

  You cannot dynamically change the **cache status** in a clustered environment from `log only` on a local cache while other instance use a different **cache status**.
- **cache replacement** – describes the cache replacement policy. For named caches and default data caches, the replacement policy is `strict LRU` or `relaxed LRU`. Change this parameter dynamically with **sp_cacheconfig**, or change the value in the server configuration file to have the change take place after the next server restart. The cache replacement policy must be `none` for in-memory databases because they do not use buffer or page replacement.
- **local cache partition number** – number of cache partitions. You may partition a named cache into multiple cache partitions. The acceptable values are 0, 2, 4, 8, 16, 32, 64 or 128. You cannot change the number of cache partitions dynamically; you must restart SAP ASE for the change to take effect.

Dump Configuration Parameter Group

The dump configuration configuration parameter group represents these user-created dump configurations.

- **stripe directory** – is the directory in which files are archived during the dump operation. Archived files are typically named using this convention:

  ```
  database_name.nump_type.date-timestamp.stripeID
  ```
- **external api name** – is the name of the external API (byte stream device) to be used for the dump operation, and must conform to this format:

  ```
  External API Name::Options
  ```
- **number of stripes** – is the number of stripe devices to use during the dump operation. By default, a single stripe device is used.
- **number of retries** – is the number of times the server tries the dump operation for nonfatal errors up to a maximum of 5 times. The default is 0.
- **block size** – is the block size for the dump device and overrides the default block size for the device. **blocksize** must be at least 1 database page, and an exact multiple of the database page size.
- **compression level** – is the compression level for compressed dumps. By default, compression is disabled.
- **retain days** – is the number of days during which the dump cannot be overwritten. Backup Server requires confirmation to overwrite an unexpired volume. By default, **retaindays** is 0, and dumps can be overwritten.
- **init** – specifies whether the volume must be reinitialized. The default is "**noinit**".
- **verify** – specifies if Backup Server must perform a minimal page-header or full structural row check on the data pages as they are copied to archives. There is no structural check made to global allocation map (GAM), object allocation map (OAM), allocation pages, indexes, text, or log pages. By default, there is no verification of data pages durig archiving.
- **notify** – the default message destination to Backup Server. One of:
  - **client** - route messages to the terminal that initiated the **dump** command.
  - **operator_console** - route messages to the terminal on which Backup Server is running
- **remote backup server name** – specifies the remote Backup Server used for a dump operation. The default is SYB_BACKUP.

This example contains multiple dump configurations as created in the configuration file:

```
[dump configuration : dmp_cfg1]
   stripe_dir = /work/dmp_cfg1_dir
   ext_api = DEFAULT
   num_stripes = 5
   retry = 0
   blocksize = DEFAULT
   compression = 9
```

```
   retaindays = DEFAULT
   init = DEFAULT
   verify = DEFAULT
   backup_srv_name = DEFAULT

[dump configuration : dmp_cfg2]
   stripe_dir = /work/dmp_cfg2_dir
   ext_api = syb_tsm
   num_stripes = DEFAULT
   retry = 3
   blocksize = DEFAULT
   compression = DEFAULT
   retaindays = DEFAULT
   init = DEFAULT
   verify = DEFAULT
   backup_srv_name = SYB_REMOTE
```

Configuration Parameters

In many cases, the maximum allowable values for configuration parameters are usually limited by available memory, rather than by **sp_configure** limitations.

Note: To find the maximum supported values for your platform and version of SAP ASE, see *Installation Guide > SAP ASE Specifications* for your platform.

Alphabetical Listing of Configuration Parameters

Parameter descriptions include both summary and detailed information about each configuration parameter.

abstract plan cache

abstract plan cache enables caching of abstract plan hash keys.

See *Performance and Tuning Series: Query Processing and Abstract Plans > Creating and Using Abstract Plans*.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

abstract plan load must be enabled for plan caching to take effect.

abstract plan dump

abstract plan dump enables the saving of abstract plans to the **ap_stdout** abstract plans group.

See *Performance and Tuning Series: Query Processing and Abstract Plans > Creating and Using Abstract Plans*.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

abstract plan load

abstract plan load enables association of queries with abstract plans in the `ap_stdin` abstract plans group.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

See *Performance and Tuning Series: Query Processing and Abstract Plans > Creating and Using Abstract Plans*.

abstract plan replace

abstract plan replace enables plan replacement for abstract plans in the `ap_stdout` abstract plans group.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic

Summary Information	
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

abstract plan load must be enabled for replace mode to take effect.

See *Performance and Tuning Series: Query Processing and Abstract Plans > Creating and Using Abstract Plans*.

abstract plan sharing

abstract plan sharing enables abstract plan sharing between different users.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Tables of a query using a shared abstract plan must be explicitly prefixed with the owner name or belong to the DBO user.

additional network memory

additional network memory sets the amount of additional memory allocated to the network memory pool.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, Network Communication, Physical Memory

You must increase the size of the network memory pool if you configure the **max network packet size** greater than the default network packet size. SAP ASE rounds down the value you enter to the nearest 2K value.

When a login requests a large packet size, SAP ASE verifies it has sufficient memory available to satisfy the request. If it does not, SAP ASE finds the largest available block of memory and tries the appropriate size (which is a multiple of **default network packet size**) less than the largest memory block. If that fails, SAP ASE decreases the value of the request by the number of bytes equal to **default network packet size**, if this is available. SAP ASE continues for 10 iterations, or until the size equals the value of **default network packet size**, whichever comes first. On the tenth iteration, SAP ASE uses the value of the **default network packet size** for the packet size.

SAP ASE guarantees that every user connection can log in at the default packet size. However, if you increase **max network packet size**, you must also increase the value for **additional network memory**.

Note: Users may be able to connect to SAP ASE with a larger packet size with **additional network memory** set to zero because of system-configured overhead in the network memory pool.

If you increase **max network packet size** but do not increase **additional network memory**, SAP ASE does not guarantee that clients who request network packet sizes larger than the default size can login at the requested packet size.

Increasing **additional network memory** may improve performance for applications that transfer large amounts of data. To determine the value for **additional network memory** when your applications use larger packet sizes:

1. Estimate the number of simultaneous users who will request the large packet sizes, and the sizes their applications will request,
2. Multiply this sum by three, since each connection needs three buffers,
3. Add two percent for overhead for 32-bit servers, or four percent for 64-bit servers, and
4. Round the value to the next highest multiple of 2048.

For example, if you estimate these simultaneous needs for larger packet sizes:

Application	Packet size	Overhead
bcp	8192	
Client-Library	8192	
Client-Library	4096	
Client-Library	4096	
Total	24576	
Multiply by 3 buffers/user	* 3=73728	

Application	Packet size	Overhead
Compute 2% overhead		* .02=1474
Add overhead	+ 1474	
Additional network memory	75202	
Round up to multiple of 2048	75776	

Set **additional network memory** to 75,776 bytes.

aggresive task stealing

aggressive task stealing sets the SAP ASE scheduler task stealing policy to aggressive.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System Administrator
Configuration groups	SQL Server Administration

allocate max shared memory

allocate max shared memory determines whether SAP ASE allocates all the memory specified by **max memory** at start-up, or only the amount of memory the configuration parameter requires.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Physical Memory

By setting **allocate max shared memory** to 0, you ensure that SAP ASE uses only the amount of shared memory required by the current configuration, and allocates only the amount of memory required by the configuration parameters at start-up, which is a smaller value than **max memory**.

If you set **allocate max shared memory** to 1, SAP ASE allocates all the memory specified by **max memory** at start-up. If you set **allocate max shared memory** to 1, and if you increase **max**

memory, SAP ASE attempts to allocate the memory immediately. If the memory allocation fails, SAP ASE writes messages to the error log. Check the error log to verify that no errors have occurred.

A successful memory allocation means that SAP ASE always has the memory required for any memory configuration changes you make and there is no performance degradation while the server readjusts for additional memory. However, if you do not predict memory growth accurately, and **max memory** is set to a large value, you may waste total physical memory.

allow backward scans

allow backward scans controls how the optimizer performs **select** queries that contain the **order by...desc** command.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Query Tuning

- When the value is set to 1, the optimizer can access the index or table rows by following the page chain in descending index order.
- When the value is set to 0, the optimizer selects the rows into a worktable by following the index page pointers in ascending order, and then sorts the worktable in descending order.

The first method—performing backward scans—can speed access to tables that need results ordered by descending column values. Some applications, however, may experience deadlocks due to backward scans. In particular, look for increased deadlocking if you have **delete** or **update** queries that scan forward using the same index. There may also be deadlocks due to page splits in the index.

Use **print deadlock information** to send messages about deadlocks to the error log. Alternatively, you can use **sp_sysmon** to check for deadlocking. See the *Performance and Tuning Series: Locking and Concurrency Control*.

allow nested triggers

allow nested triggers controls the use of nested triggers.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)

Summary Information	
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

When the value is set to 1, data modifications made by triggers can fire other triggers. Set **allow nested triggers** to 0 to disable nested triggers. A **set** option, **self_recursion**, controls whether the modifications made by a trigger can cause that trigger to fire again.

Note: The multiple triggers feature does not change the behavior of the **allow nested triggers** configuration parameter.

allow procedure grouping

allow procedure grouping controls the ability to group stored procedures of the same name so that they can be dropped with a single **drop procedure** statement.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System security officer
Configuration group	Security Related

allow remote access

allow remote access controls logins from remote SAP ASEs. The default value of 1 allows SAP ASE to communicate with Backup Server.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration groups	Backup/Recovery, Network Communication

Note: Setting the value to 0 disables server-to-server RPCs. Since SAP ASE communicates with Backup Server via RPCs, setting this parameter to 0 makes it impossible to back up a database.

Since other system administration actions are required to enable remote servers other than Backup Server to execute RPCs, leaving this option set to 1 does not constitute a security risk.

allow resource limits

allow resource limits controls the use of resource limits.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, SQL Server Administration

When the value is set to 1, the server allocates internal memory for time ranges, resource limits, and internal server alarms. The server also internally assigns applicable ranges and limits to user sessions. The output of **showplan** and **statistics io** displays the optimizer's cost estimate for a query. Set **allow resource limits** to 0 to disable all resource limits.

allow sendmsg

allow sendmsg enables or disables sending messages from SAP ASE to a User Datagram Protocol (UDP) port.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System security officer
Configuration group	Network Communication

When **allow sendmsg** is set to 1, any user can send messages using **sp_sendmsg** or **syb_sendmsg**.

Note: Sending messages to UDP ports is not supported on Windows.

allow sql server async i/o

allow sql server async i/o enables SAP ASE to run with asynchronous disk I/O.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Disk I/O

To use asynchronous disk I/O, enable it on both SAP ASE and your operating system. See your operating system documentation for information on enabling asynchronous I/O at the operating system level.

Disk I/O always runs faster asynchronously than synchronously. This is because when SAP ASE issues an asynchronous I/O, it does not have to wait for a response before issuing further I/Os.

allow updates to system tables

allow updates to system tables enables users with the system administrator role to make changes to the system tables and to create stored procedures that can modify system tables.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

A database administrator can update system tables in any tables that he or she owns if **allow updates to system tables** is enabled.

System tables include:

- All SAP-supplied tables in the `master` database
- All tables in user databases that begin with “sys” and that have an ID value in the `sysobjects` table of less than or equal to 100

Warning! Incorrect alteration of a system table can result in database corruption and loss of data. To protect against errors that might corrupt your databases, always use **begin transaction** when modifying a system table. Immediately after finishing your modifications, disable **allow updates to system tables**.

Stored procedures and triggers you create while **allow updates to system tables** is set on can update the system tables, even after the parameter has been set off. When you set **allow updates to system tables** to on, you create a "window of vulnerability," a period of time during which users can alter system tables or create a stored procedure with which the system tables can be altered in the future.

Because the system tables are so critical, SAP suggests that you set this parameter to on only in highly controlled situations. To guarantee that no other users can access SAP ASE while the system tables can be directly updated, restart SAP ASE in single-user mode. For details, see **startserver** and **dataserver** in the *Utility Guide*.

Note: The server-wide configuration option **allow updates to system tables** takes precedence over the stored procedure settings for **allow updates to system tables**. If you do not **enable allow updates to system tables** at the server level, individual stored procedure settings determine whether you can modify system catalogs.

average cap size

average cap size is reserved for future use.

Summary Information	
Default value	200
Range of values	100 – 10000
Status	Static
Display level	
Required role	
Configuration group	Diagnostics

audit queue size

audit queue size determines the size of an audit queue. The in-memory audit queue holds audit records generated by user processes until the records can be processed and written to the audit trail.

Summary Information	
Default value	100
Range of values	1 – 65535
Status	Dynamic

Summary Information	
Display level	Intermediate
Required role	System security officer
Configuration groups	Memory Use, Security Related

When you configure the queue suze, there is a trade-off between performance and risk. If the queue is too large, records can remain in it for some time. As long as records are in the queue, they are at risk of being lost if the system fails. However, if the queue is too small, it can repeatedly become full, which affects overall system performance; user processes that generate audit records sleep if the audit queue is full.

Following are some guidelines for determining how big your audit queue should be. You must also take into account the amount of auditing to be performed at your site.

- The memory requirement for a single audit record is 424 bytes; however, a record can be as small as 22 bytes when it is written to a data page.
- The maximum number of audit records that can be lost in a system failure is the size of the audit queue (in records), plus 20. After records leave the audit queue, they remain on a buffer page until they are written to the current audit table on the disk. The pages are flushed to disk every 20 records, less if the audit process is not constantly busy.
- In the system audit tables, the `extrainfo` field and fields containing names are of variable length, so audit records that contain complete name information are generally larger.

The number of audit records that can fit on a page varies from 4 to as many as 80 or more. The memory requirement for the default audit queue size of 100 is approximately 42K.

auditing

auditing enables or disables auditing for SAP ASE.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

automatic cluster takeover

automatic cluster takeover determines whether an instance automatically recovers from a cluster failure. Setting **automatic cluster takeover** to 1 allows an instance that is starting to automatically recover from an abrupt total cluster failure.

Summary Information	
Default value	1
Valid values	1 (enabled), 0 (disabled)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

Setting **automatic cluster takeover** to:

- 1 – allows an instance that is starting to automatically recover from an abrupt total cluster failure
- 0 – the cluster may not be able to recover from an abrupt cluster failover unless you include the `--cluster_takeover` parameter.

The Cluster Edition uses quorum heartbeats and a cluster takeover algorithm to determine when cluster takeover should be performed. This algorithm allows an instance that is starting to distinguish between an inability to join the cluster because the cluster has crashed (in which case takeover is appropriate) and an inability to join the cluster because the instance that is starting does not have network connectivity (in which case takeover is not appropriate).

If you disable **automatic cluster takeover** (set it to 0), the Cluster Edition writes the results of the algorithm to the error log as an advisory message and then exits.

If you enable **auotomatic cluster takeover** (set it to 1), the Cluster Edition starts as the cluster coordinator and recovers the databases. This is guaranteed to be a safe operation in environments that have I/O fencing enabled.

In environments without I/O fencing, a malfunction of the algorithm could introduce data corruption, so you can set the configuration parameter to 0 to disable this algorithm. However, environments without I/O fencing have a risk of data corruption, and disabling automatic cluster takeover does not mitigate all of those risks.

builtin date strings

builtin date strings determines if the server interprets strings given to chronological builtins as bigdatetimes. If a string is given as an argument in place of the chronological value, the server interprets it as a `datetime` value regardless of its apparent precision.

Summary Information	
Default value	0
Range of values	0 – 1
Status	Dynamic
Display level	
Required role	
Configuration group	Query tuning

This default behavior may be changed by setting the configuration parameter **builtin date strings** or the set option **builtin_date_strings**.

caps per ccb

caps per ccb is reserved for future use.

Summary Information	
Default value	50
Range of values	5 – 50
Status	Static
Display level	
Required role	
Configuration group	Diagnostics

capture compression statistics

capture compression statistics enables the `monTableCompression` monitoring table to begin capturing compression statistics.

Summary Information	
Default value	0 (off)
Range of values	1 (on), 0 (off)
Status	Dynamic
Display level	10

Summary Information	
Required role	System administrator
Configuration group	SQL Server Administration

capture missing statistics

capture missing statistics enables or disables SAP ASE to capture information about columns that have missing statistics.

Summary Information	
Default value	0 (off)
Range of values	1 (on), 0 (off)
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	Query Tuning

check password for digit

check password for digit enables or disables checking for at least one character or digit in a password. If set, this parameter does not affect existing passwords.

Summary Information	
Default value	0 (off)
Range of values	1 (on), 0 (off)
Status	Dynamic
Display level	10
Required role	System security officer
Configuration group	Security Related

CIPC large message pool size

CIPC large message pool size specifies the number of large message buffers allocated by CIPC at start-up time.

Summary Information	
Default value	512
Valid values	512 – 2147483647
Status	Static

Summary Information	
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

CIPC regular message pool size

CIPC regular message pool size specifies the number of regular message buffers allocated by CIPC at start-up time.

Summary Information	
Default value	8192
Valid values	2048 – 2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

cis bulk insert array size

cis bulk insert array size determines the size of the array CIS internally buffers (and asks the Open Client bulk library to transfer as a block) when performing a bulk transfer of data from one SAP ASE to another SAP ASE.

Summary Information	
Default value	50
Range of values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

cis bulk insert batch size

cis bulk insert batch size determines how many rows from the source tables are to be bulk copied into the target table as a single batch using **select into**.

Summary Information	
Default value	0

Summary Information	
Range of values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

If you leave **cis bulk insert batch size** at 0, all rows are copied as a single batch. Otherwise, after the count of rows specified by this parameter has been copied to the target table, the server issues a bulk commit to the target server, causing the batch to be committed.

If a normal client-generated bulk copy operation (such as that produced by the **bcp** utility) is received, the client is expected to control the size of the bulk batch, and the server ignores the value of this configuration parameter.

cis connect timeout

cis connect timeout determines the wait time, in seconds, for a successful Client-Library connection.

Summary Information	
Default value	0
Range of values	0–32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

cis cursor rows

cis cursor rows specifies the cursor row count for **cursor open** and **cursor fetch** operations. Increasing this value means more rows are fetched in one operation, which increases speed but requires more memory.

Summary information	
Default value	50
Range of values	1 – 2147483647
Status	Dynamic
Display level	Comprehensive

Summary information	
Required role	System administrator
Configuration group	Component Integration Services

cis idle connection timeout

cis idle connection timeout configures SAP ASE to check for CIS connections to any remote server that have been unused longer than the specified number of seconds. SAP ASE deletes the unused connections and reallocates their resources.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

Although the number you specify is in seconds, the housekeeper task wakes up, at most , once a minute, so idle connections may be idle for much longer than the configured value. SAP ASE does not drop idle connections if a transaction is active on the connection, and reestablishes the connection automatically if the user executes any command that accesses the connection.

cis packet size

cis packet size specifies the size of Tabular Data Stream™ (TDS) packets that are exchanged between the server and a remote server when a connection is initiated.

Summary Information	
Default value	512
Range of values	512–32768
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

The default packet size on most systems is 512 bytes, and this may be adequate for most applications. However, larger packet sizes may result in significantly improved query performance, especially when `text`, `unitext`, and `image` or bulk data is involved.

If you specify a packet size larger than the default, then the target server must be configured to allow variable-length packet sizes, using:

- **additional netmem**
- **maximum network packet size**

If you specify a packet size on the target server that is smaller than the default, it uses the default packet size instead of the **cis packet size**.

cis rpc handling

cis rpc handling specifies the default method for remote procedural call (RPC) handling.

Summary Information	
Default value	0 (off), default value of 1 for the Cluster Edition
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

Setting **cis rpc handling** to 0 sets the SAP ASE site handler as the default RPC handling mechanism. Setting the parameter to 1 forces RPC handling to use Component Integration Service access methods. See **set cis rpc handling** in the *Component Integration Services Users Guide*.

cluster heartbeat interval

cluster heartbeat interval controls the interval that cluster instances use to send and check the heartbeat status.

Summary Information	
Default value	10
Valid values	1– 127
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

Using a lower value for **cluster heartbeat interval** reduces the failure detection time but increases the risk of a false failure because of a transient problem (such as an overloaded CPU). Tuning **cluster heartbeat interval** to a larger value reduces the risk of a false failure but increases the time needed to detect a failure.

cluster heartbeat retries

cluster heartbeat retries controls the number of times an instance retries a failed cluster heartbeat before entering failure mode.

Summary Information	
Default value	1
Valid values	1– 127
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

Tuning **cluster heartbeat retries** to a lower value reduces the time to detect failure but increases the risk of a false failure because of a transient problem (such as an overloaded CPU). Tuning **cluster heartbeat retries** to a larger value reduces the risk of a false failure but increases the time needed to detect a failure.

cluster redundancy level

cluster redundancy level (CRL) controls the number of recoverable failed instances in a shared-disk cluster. It is the maximum number of instances that can fail simultaneously while allowing recovery to proceed concurrently with other activity.

Summary Information	
Default value	1
Valid values	1 to n – 1, where n is the maximum number of instances specified in `cluster.cfg` or the quorum file.
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

If the number of failed instances exceeds the maximum number specified by the CRL, the cluster shuts down.

For the cluster to start, the value for **cluster redundancy level** must be at least one less than the value of maximum number of instances specified in `cluster.cfg` or the quorum file. Thus, the cluster cannot start if you set either of these:

- The value of maximum number of instances to a value that is equal to or less than the value of **cluster redundancy level**
- The value of **cluster redundancy level** to a value equal to or greater than the value of maximum number of instances

When the **cluster redundancy level** value increases, so does the messaging traffic, as there are multiple copies of each lock in the cluster, and an increase in overhead required to maintain this redundancy level.

Other configuration parameters, such as **number of locks** and **cache size** need more resources for **cluster redundancy level** values greater than 1, which means you must increase **max memory** for the same **number of locks** value.

SAP recommends that you set **cluster redundancy level** to 1 if the failure of multiple instances is expected to be uncommon.

cluster vote timeout

cluster vote timeout controls the maximum amount of time an instance waits for other instances to vote during the voting period. An instance waits only for those instances which it believes are running.

Summary Information	
Default value	60
Valid values	1– 127
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

Tuning **cluster vote timeout** to a lower value can reduce failover time, but increases the risk that an instance that is running is excluded from the new cluster view. Tuning **cluster vote timeout** to a larger value reduces the risk that an running instance is excluded from the new cluster view, but may increase failover time.

column default cache size

column defult cache size determines the size of the cache that SAP ASE must keep in memory to provide defaults for nonmaterialized columns.

Summary Information	
Default value	1024 pages (2MB)
Range of values	128 – 8192 (pages)

Summary Information	
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

compression info pool size

compression info pool size determines the size of the memory pool used for compression.

Summary Information	
Default value	4096
Range of values	0 – 2147483647
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	SQL Server Administration

compression memory size

compression memory size determines the size (in 2KB pages) of the memory pool SAP ASE uses to decompress a compressed dump. SAP ASE uses the value for **compression memory size** while loading a compressed dump into an archive database.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Physical Memory

When you set **compression memory size** to 0, no pool is created and a compressed dump cannot be loaded.

See *System Administration Guide: Volume 2 > Creating a Compression Memory Pool*.

configuration file

configuration file specifies the location of the configuration file currently in use.

Summary Information	
Default value	0 (off)
Range of values	One of: 0, **verify**, **read**, **write**, or **restore**
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	General Information

In **sp_configure** output, the "Run Value" column displays only 10 characters, so the output may not display the entire path and name of your configuration file.

cost of a logical io

cost of a logical io specifies the cost of a single logical I/O.

Summary Information	
Default value	2
Range of values	0 – 254
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

cost of a physical io

cost of a phsyical io specifies the cost of a single physical I/O.

Summary Information	
Default value	25
Range of values	0 – 254
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

cost of a cpu unit

cost of a cpu unit specifies the cost of a single CPU operation.

Summary Information	
Default value	1000
Range of values	1 – 65534
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

The cost of a serial plan in the optimizer is described by this formula:

```
Cost = PIO X estimated_pio + LIO X estimated_lio + 100 X CPU/
estimated_cpu
```

Where the default values are:

- *estimated_pio* = 25
- *estimated_lio* = 2
- *estimated_cpu* = 1000

If your SAP ASE has sufficient memory, then all tables exist in memory, and a value of 0 for **cost of a physical io** is appropriate.

If your CPU is fast enough so the value for **cost of a cpu unit** is not a issue, use this formula to determine the cost of CPU, which combines 2 LIO and 25 PIO (the default values):

```
CPU X 100/configuration_value
```

The default value for *configuration_value* is 1000.

As you increase the value for **cost of a cpu unit**, this formula reduces the impact of CPU on cost.

cpu accounting flush interval

cpu accounting flush interval specifies the amount of time, in machine clock ticks (non-SAP ASE clock ticks), that SAP ASE waits before flushing CPU usage statistics for each user from `sysprocesses` to `syslogins`, a procedure used in charge-back accounting.

Summary Information	
Default value	200
Range of values	1–2147483647
Status	Dynamic

Summary Information	
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

When a user logs in to SAP ASE, the server begins accumulating figures for CPU usage for that user process in `sysprocesses`. When a user logs off SAP ASE, or when the value of **cpu accounting flush interval** is exceeded, the accumulated CPU usage statistics are flushed from `sysprocesses` to `syslogins`. These statistics continue accumulating in `syslogins` until you clear the totals. Display the current totals from `syslogins` using **sp_reportstats**.

The value to which you set **cpu accounting flush interval** depends on the type of reporting you intend to do. If you run reports on a monthly basis, set **cpu accounting flush interval** to a relatively high value. With infrequent reporting, it is less critical that the data in `syslogins` be updated frequently.

However, if you perform periodic ad hoc selects on the `totcpu` column in `syslogins` to determine CPU usage by process, set **cpu accounting flush interval** to a lower value to increase the likelihood of the data in `syslogins` being up-to-date when you execute your selects.

Setting **cpu accounting flush interval** to a low value may cause the lock manager to mistakenly identify processes as potential deadlock victims. When the lock manager detects a deadlock, it checks the amount of CPU time accumulated by each competing processes. The process with the lesser amount is chosen as the deadlock victim and is terminated by the lock manager. Additionally, when **cpu accounting flush interval** is set to a low value, the task handlers that store CPU usage information for processes are initialized more frequently, thus making processes appear as if they have accumulated less CPU time than they actually have. Because of this, the lock manager may select a process as the deadlock victim when, in fact, that process has more accumulated CPU time than the competing process.

If you do not intend to report on CPU usage at all, set **cpu accounting flush interval** to its maximum value. This reduces the number of times `syslogins` is updated, and reduces the number of times its pages must be written to disk.

cpu grace time

cpu grace time, together with **time slice**, specifies the maximum amount of time that a user process can run without yielding the CPU before SAP ASE preempts it and terminates it with a timeslice error.

Summary Information	
Default value	500

Summary Information	
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

The units for **cpu grace time** are time ticks, as defined by **sql server clock tick length**.

When a process exceeds **cpu grace time** SAP ASE "infects" it by removing the process from the internal queues. The process is killed, but SAP ASE is not affected. This prevents runaway processes from monopolizing the CPU. If any of your user processes become infected, you may be able to temporarily fix the problem by increasing the value of **cpu grace time**. However, be sure that the problem really is a process that takes more than the current value of **cpu grace time** to complete, rather than a runaway process.

Temporarily increasing the **cpu grace time** value is a workaround, not a permanent fix, since it may cause other complications. See *Performance and Tuning Series: Basics > Using Engines and CPUs* for a more detailed discussion of task scheduling.

current audit table

current audit table establishes the table where SAP ASE writes audit rows.

Summary Information	
Default value	1
Range of values	0–8
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

A system security officer can change the current audit table, using:

```
sp_configure "current audit table", n
  [, "with truncate"]
```

where `n` is an integer that determines the new current audit table, as follows:

- 1 means `sysaudits_01`, 2 means `sysaudits_02`, and so forth, up to 8.
- 0 tells SAP ASE to set the current audit table to the next table. For example, if your installation has three audit tables, `sysaudits_01`, `sysaudits_02`, and `sysaudits_03`, SAP ASE sets the current audit table to:

- 2 if the current audit table is `sysaudits_01`
- 3 if the current audit table is `sysaudits_02`
- 1 if the current audit table is `sysaudits_03`

"with truncate" specifies that SAP ASE should truncate the new table if it is not already empty. **sp_configure** fails if this option is not specified and the table is not empty.

Note: If SAP ASE truncates the current audit table, and you have not archived the data, the table's audit records are lost. Be sure that the audit data is archived before using the **with truncate** option.

To execute **sp_configure** to change the current audit table, you must have the **sso_role** active. You can write a threshold procedure to change the current audit table automatically.

deadlock checking period

deadlock checking period specifies the minimum amount of time (in milliseconds) before SAP ASE initiates a deadlock check for a process that is waiting on a lock to be released.

Summary Information	
Default value	500
Range of values	0–2147483
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Lock Manager

Deadlock checking is time-consuming overhead for applications that experience no or very few deadlocks, and the overhead grows as the percentage of lock requests that must wait for a lock also increases.

If you set **deadlock checking period** to a nonzero value (*n*), SAP ASE initiates a deadlock check after a process waits at least *n* milliseconds. For example, you can make a process wait at least 700 milliseconds for a lock before each deadlock check by entering:

```
sp_configure "deadlock checking period", 700
```

If you set **deadlock checking period** to 0, SAP ASE initiates deadlock checking when each process begins to wait for a lock. Any value less than the number of milliseconds in a clock tick is treated as 0.

Configuring **deadlock checking period** to a higher value produces longer delays before deadlocks are detected. However, since SAP ASE grants most lock requests before this time elapses, the deadlock checking overhead is avoided for those lock requests. If your applications deadlock infrequently, set **deadlock checking period** to a higher value. Otherwise, the default value of 500 should suffice.

Use **sp_sysmon** to determine the frequency of deadlocks in your system and the best setting for **deadlock checking period**. See the *Performance and Tuning Series: Monitoring SAP ASE with sp_sysmon*.

deadlock pipe active

deadlock pipe active controls whether SAP ASE collects deadlock messages.

Summary Information	
Default value	0
Range of values	0–1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Monitoring

If both **deadlock pipe active** and **deadlock pipe max messages** are enabled, SAP ASE collects the text for each deadlock. Use `monDeadLock` to retrieve these deadlock messages.

deadlock pipe max messages

deadlock pipe max messages determines the number of deadlock messages SAP ASE stores per engine.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

The total number of messages in the `monSQLText` table is the value of **sql text pipe max messages** times the number of engines running.

deadlock retries

deadlock retries specifies the number of times a transaction can attempt to acquire a lock when deadlocking occurs during an index **page split** or shrink.

Summary Information	
Default value	5
Range of values	0–2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Lock Manager, SQL Server Administration

For example, in the figure below:

- Transaction A locks page 1007 and needs to acquire a lock on page 1009 to update the page pointers for a page split.
- Transaction B is also inserting an index row that causes a page split, holds a lock on page 1009, and needs to acquire a lock on page 1007.

Figure 2: Deadlocks during page splitting in a clustered index

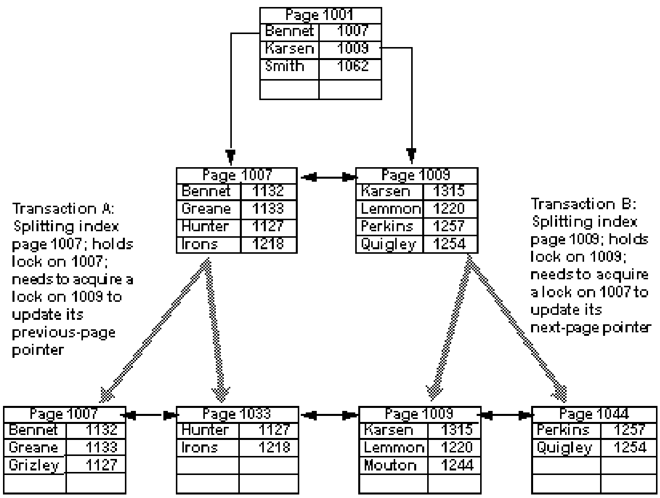

In this situation, rather than immediately choosing a process as a deadlock victim, SAP ASE relinquishes the index locks for one of the transactions. This often allows the other transaction to complete and release its locks.

For the transaction that surrendered its locking attempt, the index is rescanned from the root page, and the page split operation is attempted again, up to the number of times specified by **deadlock retries**.

sp_sysmon reports on deadlocks and retries. See the *Performance and Tuning Series: Locking and Concurrency Control*.

default character set id

default character set id specifies the number of the default character set used by the server.

Summary Information	
Default value	1
Range of values	0–255
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	Languages

The default is set at installation, and can be changed later with the SAP installation utilities.

default database size

default database size sets the default number of megabytes allocated to a new user database if **create database** is issued without any size parameters.

Summary Information	
Default value	3MB
Range of values	2[a] –10000 (the minimum determined by server's logical page size)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

A database size given in a **create database** statement takes precedence over the value set by this configuration parameter.

If most of the new databases on your SAP ASE require more than one logical page size, you may want to increase the default.

Note: If you alter the `model` database, you must also increase the **default database size**, because the **create database** command copies `model` to create a new user database.

default exp_row_size percent

default exp_row_size percent reserves space for expanding updates in data-only-locked tables, to reduce row forwarding.

Summary Information	
Default value	5
Range of values	0–100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

An "expanding update" is any update to a data row that increases the length of the row. Data rows that allow null values or that have variable-length columns may be subject to expanding updates. In data-only-locked tables, expanding updates can require row forwarding if the data row increases in size so that it no longer fits on the page.

The default value sets aside 5 percent of the available data page size for use by expanding updates. Since 2002 bytes are available for data storage on pages in data-only-locked tables, this leaves 100 bytes for expansion. This value is applied only to pages for tables that have variable-length columns.

Setting **default exp_row_size percent** to 0 means that all pages are completely filled and no space is left for expanding updates.

default exp_row_size percent is applied to data-only-locked tables with variable-length columns when **exp_row_size** is not explicitly provided with **create table** or set with **sp_chgattribute**. If a value is provided with **create table**, that value takes precedence over the configuration parameter setting. See the *Performance and Tuning Series: Locking and Concurrency Control*.

default fill factor percent

default fill factor percent determines how full SAP ASE makes each index page when it is creating a new index on existing data, unless the fill factor is specified in the **create index** statement.

Summary Information	
Default value	0
Range of values	0–100
Status	Dynamic

Summary Information	
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

The **fillfactor** percentage is relevant only when the index is created. As data changes, pages are not maintained at any particular level of fullness.

default fill factor percent affects:

- The amount of storage space used by your data – SAP ASE redistributes the data as it creates the clustered index.
- Performance – splitting up pages uses SAP ASE resources.

There is seldom a reason to change **default fill factor percent**, especially since you can override it in the **create index** command. See **create index** in the *Reference Manual: Commands*.

default language id

default language id is the number of the language that is used to display system messages unless a user has chosen another language from those available on the server. us_english always has an ID of NULL.

Summary Information	
Default value	0
Range of values	0–32767
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Languages

Additional languages are assigned unique numbers as they are added.

default network packet size

default network packet size configures the default packet size for all SAP ASE users.

Summary Information	
Default value	2048
Range of values	512– 65024

Summary Information	
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, Network Communication, User Environment

You can set **default network packet size** to any multiple of 512 bytes; values that are not even multiples of 512 are rounded down.

Memory for all users who log in with the default packet size is allocated from the SAP ASE memory pool, as set with **total logical memory**. This memory is allocated for network packets when SAP ASE is started.

Each SAP ASE user connection uses:

- One read buffer
- One buffer for messages
- One write buffer

Each of these buffers requires **default network packet size** bytes. The total amount of memory allocated for network packets is:

```
(number of user connections + number of worker processes) * 3 *
default network packet size
```

For example, if you set **default network packet size** to 1024 bytes, and you have 50 user connections and 20 worker processes, the amount of network memory required is:

(50 + 20) * 3 * 1024 = 215040 bytes

If you increase **default network packet size**, you must also increase **max network packet size** to at least the same size. If the value of **max network packet size** is greater than the value of **default network packet size**, increase the value of **additional network memory**.

Use **sp_sysmon** to see how changing the **default network packet size** parameter affects network I/O management and task switching. For example, try increasing **default network packet size** and then checking **sp_sysmon** output to see how this affects **bcp** for large batches. See the *Performance and Tuning Series: Monitoring SAP ASE with sp_sysmon*.

Requesting a Larger Packet Size at Login

Clients automatically use the SAP ASE default packet size. Use the `-A` flag to SAP ASE client programs to request a large packet size.

For example:

```
isql -A2048
```

default sortorder id

default sortorder id is the number of the sort order that is installed as the default on the server.

Summary Information	
Default value	50
Range of values	0–255
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Languages

default unicode sortorder

default unicode sortorder is a string parameter that uniquely defines the default Unicode sort order installed on the server.

Summary Information	
Default value	binary
Range of values	Not currently used
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Unicode

default XML sortorder

default XML sortorder is a string parameter that defines the sort order used by the XML engine. A string parameter is used rather than a numeric parameter to guarantee a unique ID.

Summary Information	
Default value	binary
Range of values	(not currently used)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Unicode

See *XML Services in SAP Adaptive Server Enterprise > XML Support for I18N*.

deferred name resolution

deferred name resolution determines whether deferred name resolution is applied globally to server connections.

Summary Information	
Default value	0 (disabled)
Range of values	0 to 1
Status	dynamic
Required role	System administrator
Configuration group	Query tuning

When **deferred name resolution** is active (set to 1), deferred name resolution is applied globally to all server connections; all procedures you create in the server are created using deferred name resolution.

Therefore, the stored procedures are created without resolving the objects referenced inside the stored procedure, postponing object resolution processing to the execution time. See *Transact-SQL Users Guide > Using Stored Procedures*.

disable character set conversions

disable character set conversions enables or disables character set conversion for data moving between clients and SAP ASE.

Summary Information	
Default value	0 (enabled)
Valid values	0 (enabled), 1 (disabled)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Languages

By default, SAP ASE performs conversion on data moving to and from clients that use character sets that are different than the server's. For example, if some clients use Latin-1 (iso_1) and SAP ASE uses Roman-8 (roman8) as its default character set, data from the clients is converted to Roman-8 when being loaded into SAP ASE. For clients using Latin-1, the data is reconverted when it is sent to the client; for clients using the same character set as SAP ASE, the data is not converted.

By setting **disable character set conversions**, you can request that no conversion take place. For example, if all clients are using a given character set, and you want SAP ASE to store all data in that character set, set **disable character set conversions** to 1, and no conversion takes place.

disable disk mirroring

disable disk mirroring enables or disables disk mirroring for SAP ASE.

Summary Information	
Default value	1
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Disk I/O

This is a global variable; SAP ASE does not perform any disk mirroring after this configuration parameter is set to 1 and SAP ASE is restarted. Setting **disable disk mirroring** to 0 enables disk mirroring.

Note: You must disable disk mirroring if your SAP ASE is configured for failover.

disable jsagent core dump

disable jsagent core dump enables or disables JS Agent core dumps.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	SQL Server Administration

Disables JS Agent core dump for all platforms. When off (0), the core dump for JS Agent is enabled during signal handling. Setting **disable jsagent coredump** to on (1) disables core dumps and is not recommended.

Note: Having JS Agent core dumps enabled allows you to diagnose JS Agent crash issues. Disabling core dumps for JS Agent is not recommended.

disable varbinary truncation

disable varbinary truncation controls whether SAP ASE includes trailing zeros at the end of `varbinary` or `binary` null data.

Summary Information	
Default value	0 (on)
Range of values	0 (on), 1 (off)
Status	Dynamic
Display level	Comprehensive
Required role	System Administrator
Configuration group	SQL Server Administration

disk i/o structures

disk i/o structures specifies the initial number of disk I/O control blocks SAP ASE allocates at start-up.

Summary Information	
Default value	256
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Disk I/O, Memory Use

User processes require a disk I/O control block before SAP ASE can initiate an I/O request for the process. The memory for disk I/O control blocks is preallocated when SAP ASE starts. To minimize the chance of running out of disk I/O structures, you should configure **disk i/o structures** to as high a value as your operating system allows. See your operating system documentation for information on concurrent disk I/Os.

Use **sp_sysmon** to determine whether to allocate more disk I/O structures. See *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*. You can set the **max async i/os per server** configuration parameter to the same value as **disk i/o structures**.

DMA object pool size

DMA object pool size specifies the number of DMA (direct memory access) objects allocated by CIPC at start-up time.

Summary Information	
Default value	4096
Valid values	2048 – 2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared disk cluster

dtm detach timeout period

dtm detach timeout period sets the amount of time, in minutes, that a distributed transaction branch can remain in the detached state.

Summary Information	
Default value	0 (minutes)
Valid values	0 – 2147483647 (minutes)
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	DTM Administration

In some X/Open XA environments, a transaction may become detached from its thread of control (usually to become attached to a different thread of control). SAP ASE permits transactions to remain in a detached state for the length of time specified by **dtm detach timeout period**. After this time has passed, SAP ASE rolls back the detached transaction.

dtm lock timeout period

dtm lock timeout period sets the maximum amount of time, in seconds, that a distributed transaction branch waits for lock resources to become available.

Summary Information	
Default value	300 (seconds)
Valid values	1 – 2147483647 (seconds)

Summary Information	
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	DTM Administration

After this time has passed, SAP ASE considers the transaction to be in a deadlock situation, and rolls back the transaction branch that triggered the deadlock. This ultimately rolls back the entire distributed transaction.

Distributed transactions may potentially deadlock themselves if they propagate a transaction to a remote server, and in turn, the remote server propagates a transaction back to the originating server (described in the figure below). The work of distributed transaction “dxact1” is propagated to SAP ASE 2 via “rpc1.” SAP ASE 2 then propagates the transaction back to the coordinating server via “rpc2.” “rpc2” and “dxact1” share the same **gtrid** but have different branch qualifiers, so they cannot share the same transaction resources. If “rpc2” is awaiting a lock held by “dxact1,” a deadlock situation exists.

Figure 3: Distributed transaction deadlock

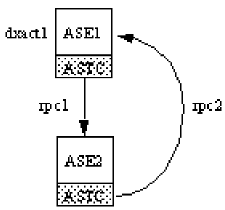

SAP ASE cannot detect interserver deadlocks. Instead, it relies on **dtm lock timeout period**. In the figure above, after **dtm lock timeout period** has expired, the transaction created for “rpc2” is aborted. This causes SAP ASE 2 to report a failure in its work, and “dxact1” is ultimately aborted as well.

The value of **dtm lock timeout period** applies only to distributed transactions. Local transactions may use a lock timeout period with the server-wide **lock wait period** parameter.

Note: SAP ASE does not use **dtm lock timeout period** to detect deadlocks on system tables.

dump history filename

dump history filename specifies the path and name of your dump history file.

Summary Information	
Default value	dumphist

Summary Information	
Valid values	
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Backup/Recovery

dump on conditions

dump on conditions determines whether SAP ASE generates a dump of data in shared memory when it encounters the conditions specified in **maximum dump conditions**.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Group Diagnostics

Note: The **dump on conditions** parameter is included for use only by SAP Technical Support. Do not modify it unless you are instructed to do so by SAP Technical Support.

dynamic allocation on demand

dynamic allocation on demand determines when memory is allocated for changes to dynamic memory configuration parameters.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Physical Memory

If you set **dynamic allocation on demand** to 1, memory is allocated only as it is needed. That is, if you change the configuration for **number of user connections** from 100 to 200, the memory for each user is added only when the user connects to the server. SAP ASE continues to add memory until it reaches the new maximum for user connections.

If **dynamic allocation on demand** is set to 0, all the memory required for any dynamic configuration changes is allocated immediately. That is, when you change the number of user connections from 100 to 200, the memory required for the extra 100 user connections is immediately allocated.

dynamic SQL plan pinning

dynamic SQL plan pinning improves performance by reducing the time spent by server connections waiting for access to the query plan manager.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Unicode

Once enabled, dynamic SQL plan pinning will only take effect if one of the following options is also enabled:

```
sp_configure 'streamlined dynamic SQL', 1
```

```
sp_configure 'enable functionality group', 1
```

When a program sends a dynamic prepared SQL statement to SAP ASE, a stored procedure containing the prepared SQL statement is created. This stored procedure is similar to a user-created stored procedure, except that it has no system catalog entries, that is, it exists in memory only. The first time that the prepared statement is executed, a query plan is compiled and executed. At the end of execution, the query plan is released to the query plan manager for re-use. When the same statement is executed again, the query plan manager is called to see if a query plan is available, and if so, the query plan manager returns the query plan to the server connection to execute. At the end of execution, the query plan is returned to the query plan manager.

All SAP ASE server connections can access the query plan manager to ask for available query plans as well as to store new query plans that they have compiled and finished executing. However, only one server connection can access the query plan manager at a time, to avoid multiple connections getting the same query plan at the same time (only a single connection can execute a given query plan at a time). Each connection will access the query plan manager twice for each dynamic prepared SQL statement that it executes: Once to acquire the query plan and once to release it for re-use.

In a highly concurrent environment (many server connections running dynamic prepared SQL statements at the same time), performance may be degraded because each connection must wait its turn to access the query plan manager when retrieving or storing a query plan. **dynamic SQL plan pinning** can improve performance by reducing the time spent by server connections waiting for access to the query plan manager. When query plan pinning is enabled, each server connection compiles a query plan for each dynamic prepared statement that it executes and does not release it to the query plan manager for re-use. Each connection keeps all query plans that it compiles for its own exclusive re-use, thus, avoiding the need to access to the query plan manager on the second and subsequent executions of the same dynamic prepared SQL statement.

Because each server connection keeps its own copy of each query plan and query plans are created from the procedure cache memory pool, this pool may need to be configured to a larger size when **dynamic SQL plan pinning** is enabled. Exactly how much larger the procedure cache needs to be depends upon the number of concurrent server connections executing dynamic SQL prepared statements: In extreme environments with small procedure cache sizes, 2-to-3 times larger may be required.

early row send increment

early row send increment configures the additional number of rows sent in the second and subsequent packets of a result set (subject to the maximum packet size).

Summary Information	
Default value	2147483647
Valid values	1 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Network Communication

See also

- *number of early send rows* on page 186

enable async database init

enable async database init ensures that all **create database** and **alter database** commands initialize databases asynchronously by default

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)

Summary Information	
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

enable backupserver HA

enable backupserver HA enables or disables the high availability Backup Server for the cluster.

Summary Information	
Default value	1
Valid values	1 (enabled), 0 (disabled)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared disk cluster

enable cis

enable cis enables or disables Component Integration Service.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

enable compression

enable compression enables or disables data compression.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)

Summary Information	
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

See the *Compression Users Guide*.

enable concurrent dump tran

enable concurrent dump tran enables or disables SAP ASE to use concurrent dumps.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

enable concurrent dump tran is part of the **enable functionality configuration** parameter group. The default value for the parameters in this group depends on the value to which **enable functionality group** is set. A value of DEFAULT for the individual configuration parameters in this group—other than **enable functionality group**—means they are set to the same value as **enable functionality group**. That is, if you set **enable functionality group** to 1, a value of DEFAULT for any other configuration parameter in the group is 1. Aside from the value for **enable functionality group**, you can ignore values of DEFAULT in the output from **sp_configure** and **sp_helpconfig** for individual configuration parameters in the Application Functionality group

enable console logging

enable console logging determines whether SAP ASE sends messages to the console.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	10

Summary Information	
Required role	System administrator
Configuration groups	SQL Server Administration

By default, **enable console logging** is disabled, and after startup, SAP ASE sends no messages to the console (the messages are still sent to the errorlog). Once enabled, **enable console logging** allows SAP ASE to send messages to the console and error log at all times.

enable DTM

enable DTM enables or disables the SAP ASE distributed transaction management (DTM) feature.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Static
Display level	10
Required role	System administrator
Configuration groups	DTM Administration, SQL Server Administration

When DTM is enabled, you can use SAP ASE as a resource manager in X/Open XA and MSDTC systems. You must restart the server for this parameter to take effect. See the *XA Interface Integration Guide for CICS, Encina, and TUXEDO* for more information about using SAP ASE in an X/Open XA environment. See *System Administration Guide: Volume 2 > Distributed Transaction Management* for information about transactions in MSDTC environments, and for information about SAP ASE native transaction coordination services.

enforce dump configuration

enforce dump configuration enables dump operations to use a dump configuration.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Backup/Recovery

enable dump history

dump history update determines whether there are updates to the dump history file at the end of the database dump operation.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Backup/Recovery

By default, SAP ASE updates the dump history file after every database dump.

enable encrypted columns

enable encrypted columns enables and disables encrypted columns.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Security Related

You cannot set **enable encrypted columns** unless you have purchased, installed, and registered the ASE_ENCRYPTION license on your server. Any attempt to set it without such licensing results in Msg. 10834:

```
Configuration parameter 'enable encrypted columns'
cannot be enabled without license 'ASE_ENCRYPTION'
```

Note: Using encrypted columns increases the logical memory used by 8198 kilobytes.

enable enterprise java beans

enable enterprise java beans enables and disables EJB Server in the SAP ASE database.

Summary Information	
Default value	0 (off)

Summary Information	
Range of values	0 (off), 1(on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Java Services

You cannot use EJB Server until the SAP ASE is enabled for EJB Server.

enable file access

enable file access enables access through proxy tables to the external file system.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

enable file access requires a license for ASE_XFS.

enable full-text search

enable full-text search enables Enhanced Full-Text Search services.

Summary Information	
Default value	1
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

enable full-text search requires a license for ASE_EFTS.

enable functionality group

enable functionality group enables or disables the changes available for specific SAP ASE version 15.7 and later features.

Summary Information	
Default value	0
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

The features include:

- Shareable inline defaults – when **enable functionality group** is set to 0, SAP ASE does not create shareable inline defaults, and does not reuse existing shareable inline defaults. The columns that share inline defaults (before you change this configuration parameter) continue to share the defaults until the defaults are removed.
- **select for update** acquiring exclusive locks
- Quoted identifiers
- Unicode noncharacters
- Monitoring cursor statements
- Reducing query processing latency
- Suppressing Job Scheduler **max task** messages

The default value for the parameters in this group depends on the value to which **enable functionality group** is set. A value of `DEFAULT` for the individual configuration parameters in this group—other than **enable functionality group**—means they are set to the same value as **enable functionality group**. That is, if you set **enable functionality group** to 1, a value of `DEFAULT` for any other configuration parameter in the group is 1.

Aside from the value for **enable functionality group**, you can ignore values of `DEFAULT` in the output from **sp_configure** and **sp_helpconfig** for individual configuration parameters in the Application Functionality group.

enable inline default sharing

enable inline default sharing enables SAP ASE to share inline defaults.

Summary Information	
Default value	0 (off)

Summary Information	
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

That is, once this configuration is enabled, SAP ASE looks for existing shareable inline defaults having the same value in the database belonging to the same user. If it finds an existing shareable default, SAP ASE binds this object to the column instead of creating a new default. However, if SAP ASE does not find an existing shareable inline default, it creates a new default

See *Transact-SQL Users Guide > Defining Defaults and Rules for Data*.

Note: The default value for **enable inline default sharing** depends on the value to which **enable functionality group** is set. If you set **enable functionality group** to:

- 0 – the default value for **enable inline default sharing** is 0
- 1 – the default value for **enable inline default sharing** is 1

However, if you set **enable inline default sharing** to 1, it uses a value of 1 regardless of what you set **enable functionality group** to.

See **enable functionality group**.

enable HA

enable HA enables or disables SAP ASE as a companion server in an active-active high availability subsystem.

Summary Information	
Default value	0 (off)
Range of values	0 – 2
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Set **enable HA** to 2 to configure SAP ASE as a companion server in an active-passive high availability subsystem.

SAP ASE uses SAP Failover to interact with the high availability subsystem. You must set **enable HA** to 1 before you run the `installhasvss` script (`insthasv` on Windows), which installs the system procedures for SAP Failover.

Note: The license information and the run value for **enable HA** are independent of each other. Whether or not you have a license for SAP Failover, the run value and the config value are set to 1 when you restart SAP ASE. Until you have a license, you cannot run SAP Failover. If you have not installed a valid license, SAP ASE logs an error message and does not activate the feature. See the installation guide for your platform for information about installing license keys.

Setting enable HA to 1 or 2 does not mean that SAP ASE is configured to work in a high availability system. You must perform the steps described in *Using SAP Failover in a High Availability System* to configure SAP ASE to be a companion server in a high availability system.

When enable HA is set to 0, you cannot configure for SAP Failover, and you cannot run `installhasvss` (`insthasv` on Windows).

enable housekeeper GC

enable housekeeper GC configures the housekeeper task.

Summary Information	
Default value	1 (on)
Range of values	0 – 5
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

The housekeeper garbage collection task performs space reclamation on data-only-locked tables. When a user task deletes a row from a data-only-locked table, a task is queued to the housekeeper to check the data and index pages for committed deletes.

See *Performance and Tuning Series: Basics > Using Engines and CPUs*.

These are valid values for **enable housekeeper GC**:

- 0 – disables the housekeeper garbage collection task, but enables the **delete** command's lazy garbage collection. You must use **reorg reclaim_space** to deallocate empty pages. This is the cheapest option with the lowest performance impact, but it may cause performance problems if many empty pages accumulate. SAP recommends that you do not use this value.

- 1 – enables lazy garbage collection for the housekeeper garbage collection task and the **delete** command. If more empty pages accumulate than your application allows, consider options 4 or 5. You can use the **optdiag** utility to obtain statistics of empty pages.
- 2 – reserved for future.
- 3 – reserved for future.
- 4 – enables aggressive garbage collection for the housekeeper garbage collection task and the **delete** command. This option is the most effective, but the **delete** command is expensive. This option is ideal if the deletes on your DOL tables are in a batch.
- 5 – enables aggressive garbage collection for the housekeeper, and lazy garbage collection for the **delete** command. This option is less expensive for deletes than option 4. This option is suitable when deletes are caused by concurrent transactions

sp_sysmon reports on how often the housekeeper garbage collection task performed space reclamation and how many pages were reclaimed. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

enable hp posix async i/o

enable hp posix async i/o enables or disables asynchronous I/O on database devices created on on HP-UX 11.31 and later file systems.

Summary Information	
Default value	0
Valid values	1 (enabled), 0 (disabled)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Disk I/O

You must first enable SAP ASE to use asynchronous I/O with the **allow sql server async i/o** configuration parameter.

The combination of these configuration parameters determines whether SAP ASE uses asynchronous or synchronous I/O for raw partitions and file systems:

enable hp posix async i/o	allow sql server async i/o = 0	allow sql server async i/o = 1
Set to 0	Synchronous I/O for file systems Synchronous I/O for raw devices	Synchronous I/O for file systems Uses /dev/async asynchronous I/O for raw devices

enable hp posix async i/o	allow sql server async i/o = 0	allow sql server async i/o = 1
Set to 1	Synchronous I/O for file systems Synchronous I/O for raw devices	POSIX asynchronous I/O for file systems POSIX asynchronous I/O for raw devices

Note: enable hp posix async i/o improves performance when you allocate database devices on file systems, but may decrease performance on database devices that are allocated on raw devices.

enable HugePages

enable HugePages enables and disables the use of huge pages on Linux platforms that support huge pages.

Summary Information	
Default value	0
Valid values	0–2
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Physical Memory

Valid values are:

- 0 – use huge pages if possible; if huge pages are not available, use regular pages.
- 1 – use only regular pages.
- 2 – use only huge pages.

enable i/o fencing

enable i/o fencing enables or disables I/O fencing for each database device that supports the SCSI-3 Persistent Group Reservation (PGR) standard.

Summary Information	
Default value	0
Valid values	1 (enabled), 0 (disabled)
Status	Static
Display level	Comprehensive

Summary Information	
Required role	System administrator
Configuration group	Shared Disk Cluster

enable ISM

enable ISM enables and disables SAP ASE to use integrated service management (ISM) on the Solaris platform.

Summary information	
Default value	0
Valid values	0–2
Status	Static
Display level	
Required role	System administrator
Configuration group	Physical Memory

Valid values are:

- 0 – use ISM if possible; if ISM is not available, use regular shared memory.
- 1 – use only regular shared memory.
- 2 – use only ISM.

enable java

enable java enables and disables Java in the SAP ASE database. You cannot install Java classes or perform any Java operations until the server is enabled for Java.

Summary Information	
Default value	0 (disabled)
Range of values	0 (disabled), 1 (enabled)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Java Services

Note: Decreasing the value for **runnable process search count** improves Java's performance when it is run on machines with a single-CPU, or when the number of SAP ASE engines is close to or equal to the number of CPUs. **runnable process search count** controls how quickly SAP ASE voluntarily yields CPU when there are no runnable processes (yielding the CPU allows the operating system to schedule the JVM to execute the Java code).

enable job scheduler

enable job scheduler determines whether Job Scheduler starts when Adaptive Server starts.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

enable js restart logging

enable js restart logging enables or disables diagnostics logging after the restart of Job Scheduler.

Summary Information	
Default value	0 (Job Scheduler logging disabled)
Valid values	1– 1
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	SQL Server Administration

enable large chunk elc

enable large chunk elc enables or disables large allocation in the engine local cache.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator

Summary Information	
Configuration group	Meta-Data Caches

enable large pool for load

enable large pool for load configures the use of large buffer pools during the recovery phase for load database and load transaction commands.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Note: The default value for **enable large pool for load** depends on the value to which **enable functionality group** is set. If you set **enable functionality group** to:

- 0 – the default value for **enable large pool for load** is 0
- 1 – the default value for **enable large pool for load** is 1

However, if you set **enable large pool for load** to 1, it uses a value of 1 regardless of what you set **enable functionality group** to.

See **enable functionality group**.

enable ldap user auth

enable ldap user auth enables or disables SAP ASE to authenticate each user on the LDAP server.

Summary Information	
Default value	0 (off)
Valid values	0 (off) – allows only `syslogins` authentication. 1 (on) – allows both LDAP and `syslogins` authentication. 2 (on) – allows only LDAP authentication.
Status	Dynamic
Display level	Comprehensive
Required role	System security officer

Summary Information	
Configuration group	Security Related

If the LDAP authentication fails, SAP ASE searches `syslogins` to authenticate the user. Use level 1 when you are migrating users from SAP ASE authentication to LDAP authentication.

enable literal autoparam

enable literal autoparam enables and disables literal server-wide parameterization.

Summary Information	
Default value	0
Range of values	1 (enabled), 0 (disabled)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Query Tuning

engine local cache percent

engine local cache percent enables or disables modification of the Engine Local Cache (ELC) as a percentage of procedure cache.

Summary Information	
Default value	50
Range of values	0–100
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Meta-Data Caches

Based on this configuration option, SAP ASE will configure procedure cache ELC size as given percentage of procedure cache size. The default value is 50, which means the ELC size will be 50% of procedure cache size.

This configurations option can be used only when large ELC is enabled (using boot time traceflag 758).

For optimal performance, a value no larger than 80 is recommended.

Example:

```
1> sp_configure "engine local cache percent",70
2> go
Parameter Name Default Memory Used Config Value Run Value Unit Type
-------------- ------- ----------- ------------ --------- ---- ----
    engine   local cache percent  50   0   70  50  percent static
(1 row affected)
Configuration option changed. Since the option is static,
Adaptive Server must be rebooted in order for the change to take
effect.
Changing the value of 'engine local cache percent' does not
increase the amount of memory Adaptive Server
```

enable logins during recovery

enable logins during recovery determines whether non-system administrator logins are allowed during database recovery.

Summary information	
Default value	1
Range of values	0 (enabled), 1 (disabled)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Security Related

A value of 1 indicates that logins are allowed during recovery, and a value of 0 indicates that logins are not allowed during recovery, that is, only the system administrator can log in to SAP ASE.

enable merge join

enable merge join enables or disables merge join at the server level.

Summary Information	
Default value	2
Range of values	0 – disables merge joins at the server level. 1 – enables merge joins at the server level. 2 – sets merge joins to their default values at the server level.
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

The default value for merge join depends on current value of the **optimization goal** configuration parameter:

Value for optimization goal	Default value for merge join
allrows_mix	on
allrows_dss	on
allrows_oltp	off

enable metrics capture

enable metrics capture enables SAP ASE to capture metrics at the server level.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

Metrics for ad hoc statements are captured in the system catalogs; metrics for statements in a stored procedure are saved in the procedure cache.

enable monitoring

enable monitoring controls whether SAP ASE collects the monitoring table data.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

enable monitoring acts as a master switch that determines whether any Monitoring configuration parameters are enabled.

Note: Any data that populates the monitoring tables without enabling **enable monitoring** should not be considered valid.

enable pam user auth

enable pam user auth controls the ability to authenticate users using pluggable authentication modules (PAM).

Summary Information	
Default value	0 (off)
Range of values	0 (off) – allows only `syslogins` authentication. 1 (on) – allows both PAM and `syslogins` authentication. 2 (on) – allows only PAM authentication.
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

When **enable pam user auth** is set to 1, SAP ASE uses the PAM provider to authenticate each user. If the PAM authentication fails, SAP ASE searches `syslogins` to authenticate the user. Use level 1 when you are migrating users from SAP ASE authentication to PAM authentication.

enable pci

enable pci enables or disables the Java PCI Bridge for SAP ASE.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on), 2 (on with operating system override)
Status	Static
Display level	Intermediate
Required role	System Administrator
Configuration group	User Environment

Note: Do not use setting “2” (on with operating system override) unless instructed to do so by SAP Technical Support. This setting enables the PCI Bridge on operating system versions that may not fully or correctly support PCI functionality.

enable permissive unicode

enable permissive unicode enables or disables SAP ASE to ignore Unicode noncharacters.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

enable permissive unicode configuration parameter allows you to include random binary data when enabled (set to 1). However, once you enable **enable permissive unicode**, SAP ASE correctly sorts only valid UTF-8 data.

Note: The default value for **enable permissive unicode** depends on the value to which **enable functionality group** is set. If you set **enable functionality group** to:

- 0 – the default value for **enable permissive unicode** is 0
- 1 – the default value for **enable permissive unicode** is 1

However, if you set **enable permissive unicode** to 1, it uses a value of 1 regardless of what you set **enable functionality group** to.

See **enable functionality group**.

enable plan sharing

enable plan sharing enables or disables SAP ASE to use shared query plans.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

The **enable plan sharing** configuration option is not automatically enabled when the **enable functionality group** configuration option is enabled: you must explicitly **enable plan sharing**.

enable predicated privileges

enable predicated privileges enables or disables predicated privileges.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Security Related

enable query tuning mem limit

enable query tuning mem limit enables or disables the query tuning memory limit.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Query Tuning

enable query tuning time limit

enable query tuning time limit enables or disables the query tuning time limit.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Intermediate
Display level	Intermediate
Required role	System administrator
Configuration group	Query Tuning

enable rapidlog

enable rapidlog enables the capture of diagnostic for 'Proc Cache Header' memory pool messages.

Note: Use diagnostic for 'Proc Cache Header' memory pool only under the guidance of the SAP support team.

Summary Information	
Default value	0 (off)
Valid values	0–255
Status	Dynamic
Display level	
Required role	System administrator
Configuration groups	Error Log

See *Performance and Turning Series: Basics > Memory Use Performance > Procedure Cache > Diagnostic for 'Proc Cache Header' Memory Pool*.

enable real time messaging

enable real time messaging enables or disables the real time messaging services.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

enable rep agent threads

enable rep agent threads enables or disables the RepAgent thread within SAP ASE.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic

Summary Information	
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Rep Agent Thread Administration

Other steps are also required to enable replication. For more information, see the Replication Server documentation.

enable row level access control

enable row level access control enables or disables row level access control.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System security officer
Configuration group	Security Related

You must have the security services license key enabled before you can configure **enable row level access control**.

enable semantic partitioning

enable semantic partitioning enables or disables partitioning other than round-robin (for example list, hash, and range partitioning) in SAP ASE.

Summary Information	
Default value	0
Range of values	1 (enabled), 0 (disabled)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Before you use any of these partitioning schemes, you must first have the appropriate license.

enable sort-merge join and JTC

enable sort-merge join and JTC enables or disables the query processor to select a sort merge or a nested loop join when SAP ASE compiles a query in compatibility mode.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

enable sort-merge join and JTC enables join transitive closure (JTC), which allows the query processor for releases earlier than 15.0 to use additional **join** clauses.

For more information about compatibility mode, see the *Migration Technology Guide*.

enable sql debugger

enable sql debugger enables or disables the SAP ASE SQL debugger, which allows you to step through your T-SQL code.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

enable ssl

enable ssl enables or disables Secure Sockets Layer session-based security.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static

Summary Information	
Display level	Comprehensive
Required role	System security officer
Configuration group	Security Related

enable stmt cache monitoring

enable stmt cache monitoring enables or disables SAP ASE to collect monitoring information about the statement cache.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administration
Configuration group	Monitoring

Once enabled, `monStatementCache` and `monCachedStatement` display valid data.

enable surrogate processing

enable surrogate processing enables or disables the processing and maintains the integrity of surrogate pairs in Unicode data.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Unicode

If **enable surrogate processing** is disabled, the server ignores the presence of surrogate pairs in the Unicode data, and all code that maintains the integrity of surrogate pairs is skipped. This enhances performance, but restricts the range of Unicode characters that can appear in the data.

enable unicode conversion

enable unicode conversion enables or disables character conversion using Unilib for the `char`, `varchar`, and `text` datatypes.

Summary Information	
Default value	1
Range of values	0 – uses only the built-in character-set conversion. 1 – uses the built-in conversion. If it cannot find a built-in conversion, SAP ASE uses the Unilib character conversion 2 – uses the appropriate Unilib conversion
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Languages, Unicode

enable unicode normalization

enable unicode normalization enables or disables Unilib character normalization.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Unicode

The normalization process modifies the data so there is only a single representation in the database for a given sequence of abstract characters. Often, characters followed by combined diacritics are replaced by precombined forms.

Set **enable unicode normalization** to 1 to use the built-in process that enforces normalization on all incoming Unicode data. If this parameter is disabled (set to 0), the normalization step is bypassed and the client code is responsible for normalization rather than the server. If normalization is disabled, performance is improved—but only if *all* clients present Unicode data to the server using the same representation.

Note: Once disabled, you cannot turn normalization on again. This limitation prevents non-normalized data from entering the database.

enable utility lvl 0 scan wait

enable utility lvl 0 scan wait enables or disables running **alter table ... add | drop partition** commands while Adaptive Server runs isolation level 0 scans.

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

The default value for **enable utility lvl 0 scan wait** depends on the value to which **enable functionality group** is set. If you set enable functionality group to:

- 0 – the default value for **enable utility lvl 0 scan wait** is 0.
- 1 – the default value for **enable utility lvl 0 scan wait** is 1.

However, if you set **enable utility lvl 0 scan wait** to 1, it uses a value of 1 regardless of what you set **enable functionality group** to. See **enable functionality group**.

enable webservices

enable webservices enables or disables Webservices.

Summary Information	
Default value	0
Range of values	1 (enabled), 0 (disabled)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

enable xact coordination

enable xact coordination enables or disables SAP ASE transaction coordination services.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1(on)

Summary Information	
Status	Static
Display level	10
Required role	System administrator
Configuration group	DTM Administration

When this parameter is set to 1 (on), coordination services are enabled, and the server can propagate transactions to other SAP ASEs. This may occur when a transaction executes a remote procedure call (RPC) to update data in another server, or updates data in another server using Component Integration Services (CIS). Transaction coordination services ensure that updates to remote SAP ASE data commit or roll back with the original transaction.

If this parameter is set to 0 (off), SAP ASE does not coordinate the work of remote servers. Transactions can still execute RPCs and update data using CIS, but SAP ASE cannot ensure that remote transactions are rolled back with the original transaction or that remote work is committed along with an original transaction, if remote servers experience a system failure. This corresponds to the behavior of SAP ASE versions earlier than version 12.x.

enable xml

enable xml enables or disables XML services.

Summary information	
Default value	0
Range of values	1 (enabled), 0 (disabled)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

engine memory log size

engine memory log size is for diagnostic use only and has no relevance in a production environment.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Dynamic
Display level	

Summary Information	
Required role	
Configuration group	Physical Memory

engine memory log size should be left at the default setting unless otherwise requested by SAP Technical Support.

errorlog pipe active

errorlog pipe active controls whether SAP ASE collects error log messages.

Summary Information	
Default value	0
Range of values	0–1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

If both **errorlog pipe active** and **errorlog pipe max messages** are enabled, SAP ASE collects all the messages sent to the error log. Use `monErrorLog` to retrieve these error log messages.

errorlog pipe max messages

errorlog pipe max messages determines the number of error log messages SAP ASE stores per engine.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Monitoring

The total number of messages in the `monSQLText` table is the value of **sql text pipe max messages** times the number of engines running.

esp execution priority

esp execution priority sets the priority of the XP Server thread for ESP execution.

Summary Information	
Default value	8
Range of values	0–15
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Extended Stored Procedure

Over long periods of time ESPs can be CPU-intensive. Also, since XP Server resides on the same machine as SAP ASE, XP Server can impact SAP ASE performance.

See the *Open Server Server-Library/C Reference Manual* for information about scheduling Open Server threads.

esp execution stacksize

esp execution stacksize sets the size of the stack, in bytes, to be allocated for ESP execution.

Summary Information	
Default value	Platform dependent: • 196608 – Linux AMD-64, IBM PLinux, HP IA64, Sun x86 64 • 139264 – SunSparc64 • 65536 – Sun Sparc32, Windows 32, IBM AIX 64, Linux IA 32, Windows 64 • 67584 – HP 64
Range of values	Platform default—2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Extended Stored Procedure

Use this parameter if you have your own ESP functions that require a larger stack size than the default, 34816.

esp unload dll

esp unload dll specifies whether DLLs that support ESPs should be automatically unloaded from XP Server memory after the ESP call has completed.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Extended Stored Procedure

If **esp unload dll** is set to 0, DLLs are not automatically unloaded. If it is set to 1, they are automatically unloaded.

If **esp unload dll** is set to 0, you can still unload individual DLLs explicitly at runtime, using **sp_freedll**.

event buffers per engine

Controls the number of monitor event buffers. Not used in the current version of SAP ASE.

event log computer name (Windows only)

event log computer name specifies the name of the Windows machine that logs SAP ASE messages in its Windows Event Log. This feature is available on Windows servers only.

Summary Information	
Default value	LocalSystem
Valid values	• Name of an Windows machine on the network configured to record SAP ASE messages • LocalSystem • NULL
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Error Log

A value of LocalSystem or NULL specifies the default local system.

You can also use the Server Config utility to set the **event log computer name** parameter by specifying the Event Log Computer Name under Event Logging.

Setting the **event log computer name** parameter with **sp_configure** or specifying the Event Log Computer Name under Event Logging overwrites the effects of the command line `-G` option, if it was specified. If SAP ASE was started with the `-G` option, you can change the destination remote machine by setting **event log computer name**.

For more information about logging SAP ASE messages to a remote site, see the *Configuration Guide for Windows*.

event logging

(Windows servers only)**event logging** enables and disables the logging of SAP ASE messages in the Windows Event Log.

Summary Information	
Default value	1
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Error Log

The default value of 1 enables SAP ASE message logging in the Windows Event Log; a value of 0 disables it.

Use the Server Config utility to set the **event logging** parameter by selecting Use Windows Event Logging under Event Logging.

Setting the **event logging** parameter or selecting Use Windows Event Logging overwrites the effects of the command line `-G` option, if it was specified.

executable codesize + overhead

executable codesize + overhead reports the combined size, in kilobytes, of the SAP ASE executable and overhead.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Calculated
Display level	Basic

Summary Information	
Required role	System administrator
Configuration group	Memory Use

This a calculated value that is not user-configurable.

extended cache size

extended cache size specifies the size of the secondary cache.

Summary Information	
Default value	0
Range of values	0 – 31457280
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	Cache Manager

FIPS login password encryption

FIPS login password encryption provides FIPS 140-2 cryptographic module support for encrypting passwords in transmission, in memory, and on disk.

Summary Information	
Default value	0
Range of values	0 – 1
Status	Static
Display level	Comprehensive
Required role	System security officer
Configuration group	Security related

Enabling **FIPS login password encryption** requires a Security and Directory Services license.

SAP ASE uses the FIPS 140-2 certified Certicom security provider for login encryption. If this configuration is not enabled, SAP ASE uses the OpenSSL security provider to perform login password encryption.

global async prefetch limit

global async prefetch limit specifies the percentage of a buffer pool that can hold the pages brought in by asynchronous prefetch that have not yet been read.

Summary Information	
Default value	10
Range of values	0–100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Cache Manager

global async prefetch limit sets the limit for all pools in all caches for which the limit has not been set explicitly with **sp_poolconfig**.

If the limit for a pool is exceeded, asynchronous prefetch is temporarily disabled until the percentage of unread pages falls below the limit. See *Performance and Tuning Series: Basics > Tuning Asynchronous Prefetch*.

global cache partition number

global cache partition number sets the default number of cache partitions for all data caches.

Summary Information	
Default value	1
Range of values	1 – 64, as powers of 2
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	Cache Manager

The number of partitions for a particular cache can be set using **sp_cacheconfig**; the local value takes precedence over the global value.

Use cache partitioning to reduce cache spinlock contention; in general, if spinlock contention exceeds 10 percent, partitioning the cache should improve performance. Doubling the number of partitions cuts spinlock contention by about one-half.

See *System Administration Guide: Volume 2 > Configuring Data Caches* for information on configuring cache partitions. Also see *Performance and Tuning Series: Basics > Tuning Asynchronous Prefetch*.

heap memory per user

heap memory per user configures the amount of heap memory per user. A heap memory pool is an internal memory created at start-up that tasks use to dynamically allocate memory as needed.

Summary Information	
Default value	4K
Valid values	0 – 2147483647 bytes
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Physical Memory

The heap memory pool is important if you are running tasks that use wide columns, which require a lot of memory from the stack. The heap memory allocates a temporary buffer that enables these wide column tasks to finish. The heap memory the task uses is returned to the heap memory pool when the task is finished.

The size of the memory pool depends on the number of user connections. SAP recommends that you set **heap memory per user** to three times the size of your logical page.

histogram tuning factor

histogram tuning factor controls the number of steps SAP ASE analyzes per histogram for **update statistics**, **update index statistics**, **update all statistics**, and **create index**.

Summary Information	
Default value	20
Range of values	1 – 100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

A value of 1 disables **histogram tuning factor**.

Note: For SAP ASE versions 15.0.2 ESD #2 and later, if you set **histogram tuning factor** to the default value of 20 and a large number of steps are requested for the histogram, the actual step count used for the histogram is limited to the value that reduces the procedure cache usage:

```
min (max (400, requested_steps), histogram_tuning_factor X
requested_steps)
```

In the following example, SAP ASE generates an intermediate 20-step histogram with 30 values:

```
sp_configure 'histogram tuning factor',20
update statistics tab using 30 values
```

SAP ASE analyzes the histogram and compresses it into the resulting histogram according to the following parameters:

- The first step is copied unchanged.
- The high-frequency steps are copied unchanged.
- The consecutive range steps are collapsed into the resulting step, so the total weight of the collapsed step is no bigger than one-thirtieth of the value.

The final histogram in `sysstatistics`:

- Has range steps generated in a way similar for a 30-step **update statistics**, and high frequency ranges are isolated as if the histogram were created with 600 steps.
- The total number of steps in the resulting histogram may differ between 30 and 600 values.
- For equally distributed data, the value should be very close to 30.
- More “frequent” values in the table means more steps in the histogram.
- If a column has few different values, all those values may appear as high-frequency cells.

You can achieve the same result by increasing the **number of histogram steps** to 600, but this uses more resources in the buffer and procedure cache

histogram tuning factor minimizes the resources histograms consume, and increases resource usage only when it is in the best interest for optimization, for example, when there is uneven distribution of data in a column, or highly duplicated values within a column. In this situation, up to 600 histogram steps are used. However, in most cases, **histogram tuning factor** uses the default value (30 in the example above).

housekeeper free write percent

housekeeper free write percent specifies the maximum percentage by which the housekeeper wash task can increase database writes.

Summary Information	
Default value	1
Range of values	0–100
Status	Dynamic
Display level	Intermediate
Required role	System administrator

Summary Information	
Configuration group	SQL Server Administration

For example, to stop the housekeeper task from working when the frequency of database writes reaches 5 percent above normal, set **housekeeper free write percent** to 5.

When SAP ASE has no user tasks to process, the housekeeper wash task automatically begins writing changed pages from cache to disk. These writes result in improved CPU utilization, decreased need for buffer washing during transaction processing, and shorter checkpoints.

In applications that repeatedly update the same database page, the housekeeper wash may initiate some unnecessary database writes. Although these writes occur only during the server's idle cycles, they may be unacceptable on systems with overloaded disks.

The table and index statistics that are used to optimize queries are maintained in memory structures during query processing. When these statistics change, the changes are not written to the `systabstats` table immediately, to reduce I/O contention and improve performance. Instead, the housekeeper chores task periodically flushes statistics to disk.

The default value allows the housekeeper wash task to increase disk I/O by a maximum of 1 percent. This results in improved performance and recovery speed on most systems.

To disable the housekeeper wash task, set the value of **housekeeper free write percent** to 0.

Set this value to 0 only if disk contention on your system is high, and it cannot tolerate the extra I/O generated by the housekeeper wash task.

If you disable the housekeeper tasks, keep statistics current. Commands that write statistics to disk are:

- **update statistics**
- **dbcc checkdb** (for all tables in a database) or **dbcc checktable** (for a single table)
- **sp_flushstats**

Run one of these commands on any tables that have been updated since the last time statistics were written to disk, at the following times:

- Before dumping a database
- Before an orderly shutdown
- After restarting, following a failure or orderly shutdown; in these cases, you cannot use **sp_flushstats**—you must use **update statistics** or **dbcc** commands
- After any significant changes to a table, such as a large bulk copy operation, altering the locking scheme, deleting or inserting large numbers of rows, or performing a **truncate table** command

To allow the housekeeper wash task to work continuously, regardless of the percentage of additional database writes, set **housekeeper free write percent** to 100.

Use **sp_sysmon** to monitor housekeeper performance. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

You might also want to look at the number of free checkpoints initiated by the housekeeper task. The *Performance and Tuning Series: Basics* describes this output.

i/o accounting flush interval

i/o accounting flush interval specifies the amount of time, in machine clock ticks, that SAP ASE waits before flushing I/O statistics for each user from `sysprocesses` to `syslogins`. This is used for charge-back accounting.

Summary Information	
Default value	1000
Range of values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

When a user logs in to SAP ASE, the server begins accumulating I/O statistics for that user process in `sysprocesses`. When the value of **i/o accounting statistics interval** is exceeded, or a user logs off SAP ASE, the accumulated I/O statistics for that user are flushed from `sysprocesses` to `syslogins`. These statistics continue accumulating in `syslogins` until you clear the totals by using **sp_clearstats**. You can display the current totals from `syslogins` by using **sp_reportstats**.

The value to which you set **i/o accounting flush interval** depends on the type of reporting you intend to do. If you run reports on a monthly basis, set **i/o accounting flush interval** to a relatively high value. With infrequent reporting, it is less critical that the data in `syslogins` be updated frequently.

If you perform periodic ad hoc selects on the `totio` column `syslogins` to determine I/O volume by process, set **i/o accounting flush interval** to a lower value. Doing so increases the likelihood of the data in `syslogins` being current when you execute your selects.

If you do not report on I/O statistics at all, set **i/o accounting flush interval** to its maximum value. This reduces the number of times `syslogins` is updated and the number of times its pages must be written to disk.

i/o batch size

i/o batch size sets the number of writes issued in a batch before the task goes to sleep. Once this batch is completed, the task is woken up, and the next batch of writes are issued, ensuring that the I/O subsystem is not flooded with many simultaneous writes.

Summary Information	
Default value	100
Range of values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Setting **i/o batch size** to the appropriate value can improve the performance of operations like **checkpoint**, **dump database**, **select into**, and so on.

i/o polling process count

i/o polling process count specifies the maximum number of processes that SAP ASE can run before the scheduler checks for disk and network I/O completions. Tuning **i/o polling process count** affects both the response time and throughput of SAP ASE.

Summary Information	
Default value	10
Range of values	1 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Note: i/o polling process count functions only when you configure SAP ASE for process kernel mode; it is nonfunctional for threaded kernel mode.

SAP ASE checks for disk or network I/O completions:

- If the number of tasks run since the last time SAP ASE checked for I/O completions equals the value for **i/o polling process count**, and
- At every SAP ASE clock tick.

As a general rule, increasing the value of **i/o polling process count** increases throughput for applications that generate a lot of disk and network I/O. Conversely, decreasing the value improves process response time in these applications, possibly at the risk of lowering throughput.

If your applications create both I/O and CPU-bound tasks, tuning **i/o polling process count** to a low value (1 – 2) ensures that I/O-bound tasks get access to CPU cycles.

For OLTP applications (or any I/O-bound application with user connections and short transactions), tuning **i/o polling process count** to a value in the range of 20 – 30 may increase throughput, but may also increase response time.

When tuning **i/o polling process count**, consider these other parameters:

- **sql server clock tick length**, which specifies the duration of the SAP ASE clock tick in microseconds.
- **time slice**, which specifies the number of clock ticks the the SAP ASE scheduler allows a user process to run.
- **cpu grace time**, which specifies the maximum amount of time, in clock ticks, a user process can run without yielding the CPU before SAP ASE preempts it and terminates it with a timeslice error.

Use **sp_sysmon** to determine the effect of changing **i/o polling process count**. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

identity burning set factor

identity burning set factor changes the percentage of potential column values that is made available in a block of column values.

Summary Information	
Default value	5000
Range of values	1–9999999
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

IDENTITY columns are of type `numeric` and scale zero whose values are generated by SAP ASE. Column values can range from a low of 1 to a high determined by the column precision.

For each table with an IDENTITY column, SAP ASE divides the set of possible column values into blocks of consecutive numbers, and makes one block at a time available in memory. Each time you insert a row into a table, SAP ASE assigns the IDENTITY column the next available value from the block. When all the numbers in a block have been used, the next block becomes available.

This method of choosing IDENTITY column values improves server performance. When SAP ASE assigns a new column value, it reads the current maximum value from memory and adds 1. Disk access becomes necessary only after all values within the block have been used. Because all remaining numbers in a block are discarded in the event of server failure (or **shutdown with nowait**), this method can lead to gaps in IDENTITY column values.

The value for **identity burning set factor** should be high enough for good performance, but not so high that gaps in column values are unacceptably large. The default value, 5000, releases .05 percent of the potential IDENTITY column values for use at one time.

To get the correct value for **sp_configure**, express the percentage in decimal form, and then multiply it by 10^{7} (10,000,000). For example, to release 15 percent (.15) of the potential IDENTITY column values at a time, specify a value of .15 times 10^7 (or 1,500,000) in **sp_configure**.

identity grab size

identity grab size allows each SAP ASE process to reserve a block of IDENTITY column values for inserts into tables that have an IDENTITY column.

Summary Information	
Default value	1
Range of values	1–2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

This is useful if you are performing inserts, and you want all the inserted data to have contiguous IDENTITY numbers. For instance, if you are entering payroll data, and you want all records associated with a particular department to be located within the same block of rows, set **identity grab size** to the number of records for that department.

identity grab size applies to all users on SAP ASE. Large **identity grab size** values result in large gaps in the IDENTITY column when many users insert data into tables with IDENTITY columns.

SAP recommends that you set **identity grab size** to a value large enough to accommodate the largest group of records you want to insert into contiguous rows.

identity reservation size

identity reservation size sets a limit for the number of identity values.

Summary Information	
Default value	1
Range of values	1–2147483647

Summary Information	
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

idle migration timeout

idle migration timeout specifies the amount of time after which an idle connection is closed without invalidating the migration request sent to the client, allowing you to stop an instance after a specified period of time without waiting for idle client connections to migrate.

Summary Information	
Default value	60
Valid values	0 – 32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

Setting **idle migration timeout** to a high value slows down a graceful shutdown because the instance must wait the specified period of time for all idle connections that issued a migration request without the client having initiated migration.

job scheduler interval

job scheduler interval sets the interval when the Job Scheduler checks which scheduled jobs are due to be executed.

Summary Information	
Default value	1 (in minutes)
Range of values	1 – 600
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

job scheduler tasks

job scheduler tasks sets the maximum number of jobs that can run simultaneously through Job Scheduler.

Summary Information	
Default value	4
Range of values	1 – 640
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

If you increase the value of **job scheduler tasks** to a higher value, you must increase the value for **number of user connections** by at least twice the value you incremented **job scheduler tasks** before starting the Job Scheduler.

However, if the SAP ASE running the scheduled jobs is the same SAP ASE that is hosting the Job Scheduler, you must increase the value for **number of user connections** by three times the value you incremented **job scheduler tasks** before starting the Job Scheduler.

Increasing the **number of user connections** may require that you increase the value for **max memory**.

For compatibility with RAP - The Trading Edition R4, you must set **job scheduler tasks** to 32.

If you set the value of **job scheduler tasks** to "default" before you upgrade SAP ASE, the server automatically sets the new default to 4.

js heartbeat interval

Specifies the intervals between two JS Agent heartbeat checks, in minutes

Summary Information	
Default value	1
Valid values	1– 1440
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	SQL Server Administration

js job output width

js job output width determines the line width the output uses for jobs stored in the `js_output` table.

Summary Information	
Default value	80
Range of values	1 – 32768
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

kernel mode

kernel mode determines the mode the SAP ASE kernel uses, threaded or process.

Summary Information	
Default value	**threaded**
Range of values	**threaded** or **process**
Status	Static
Display level	Basic
Required role	System administrator
Configuration group	SQL Server Administration

In threaded mode, SAP ASE uses operating system threads to support SAP ASE engines and specialized tasks. In this mode, SAP ASE is a single process running on the operating system. In process mode, SAP ASE uses individual processes to support SAP ASE engines. In this mode, each engine is a distinct process running on the operating system. The process mode is the same kernel used by SAP ASE versions earlier than 15.7.

On Windows, SAP ASE supports only threaded kernel mode.

The values for **kernel mode** are character data, so you must use 0 as a placeholder for the second **sp_configure** parameter, which must be numeric, and specify **threaded** or **process** as the third parameter.

SAP assumes you use threaded mode in your production server, and includes process mode for backward compatibilty. Process mode may not support features for SAP ASE 15.7 and later.

See *System Administration Guide: Volume 2 > Managing Multiprocessor Servers*.

kernel resource memory

kernel resource memory determines the size, in 2K pages, of the kernel resource memory pool from which all thread pools and other kernel resources are allocated memory.

Summary Information	
Default value	• Solaris – 6184 • All other platforms – 4096
Range of values	75 – 2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Phsyical Memory

The amount of memory available for **kernel resource memory** depends on the value of **max memory**. If **max memory** has an insufficient amount of memory to allocate to **kernel resource memory**, SAP ASE issues an error message that includes the value to which you must raise **max memory**.

The amount of **kernel resource memory** SAP ASE requires is determined by a combination of the **max online engines** and **number of user connections** configuration parameters, and the amount of memory required by the monitoring system

The monitoring system is the largest consumer of **kernel resource memory**. For example, even when configured for the default number of users and engines, SAP ASE requires additional **kernel resource memory** to manage the large number of spinlocks caused by the number of open objects, locks, and so on.

For configurations of 8 engines or fewer, SAP recommends that you add one page of kernel resource memory for every two user connections above 100. For configurations of 9 engines or more, add 1 page of kernel resource memory for every user connection.

For example, a server with 6 engines and 200 user connections should add a value of 100 to the default value. A server with 16 engines and 10000 user connections would add a value of 10000 to the default value.

large allocation auto tune

large allocation auto tune configures SAP ASE preallocate large amounts of memory for query execution, which reduces procedure cache contention.

Summary Information	
Default value	0 (off)

Summary Information	
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Meta-Data Caches

license information

license information allows SAP system administrators to monitor the number of user licenses used in SAP ASE. Enabling this parameter only monitors the number of licenses issued; it does not enforce the license agreement.

Summary Information	
Default value	25
Valid values	$0–2^{31}$
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

If **license information** is set to 0, SAP ASE does not monitor license use. If **license information** is set to a number greater than 0, the housekeeper chores task monitors the number of licenses used during the idle cycles in SAP ASE. Set **license information** to the number of licenses specified in your license agreement.

If the number of licenses used is greater than the number to which **license information** is set, SAP ASE writes this message to the error log:

```
WARNING: Exceeded configured number of user licenses
```

At the end of each 24-hour period, the maximum number of licenses used during that time is added to the `syblicenseslog` table. The 24-hour period restarts if SAP ASE is restarted.

See *Security Administration Guide > Managing SAP ASE Logins and Database Users*.

lock address spinlock ratio

address lock spinlock ratio sets the number of rows in the internal address locks hash table that are protected by one spinlock for SAP ASEs running with multiple engines.

Summary Information	
Default value	100
Range of values	1–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Lock Manager

SAP ASE manages the acquiring and releasing of address locks using an internal hash table with 1031 rows (known as hash buckets). This table can use one or more spinlocks to serialize access between processes running on different engines.

The default value for **address lock spinlock ratio** defines 11 spinlocks for the address locks hash table. The first 10 spinlocks protect 100 rows each, and the eleventh spinlock protects the remaining 31 rows. If you specify a value of 1031 or greater for **address lock spinlock ratio**, SAP ASE uses only 1 spinlock for the entire table.

lock hashtable size

lock hashtable size specifies the number of hash buckets in the lock hash table.

Summary Information	
Default value	2048
Range of values	1–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration groups	Lock Manager, Memory Use

The lock hash table manages all row, page, and table locks, and all lock requests. Each time a task acquires a lock, the lock is assigned to a hash bucket, and each lock request for that lock checks the same hash bucket. Setting this value too low results in large numbers of locks in each hash bucket and slows the searches. On SAP ASEs with multiple engines, setting this value too low can also lead to increased spinlock contention. Do not set the value to less than the default value, 2048.

lock hashtable size must be a power of 2. If the value you specify is not a power of 2, **sp_configure** rounds the value to the next highest power of 2 and prints an informational message.

The optimal hash table size is a function of the number of distinct objects (pages, tables, and rows) that can be locked concurrently. The optimal hash table size is at least 20 percent of the number of distinct objects that need to be locked concurrently. See the *Performance and Tuning Series: Locking and Concurrency Control*.

However, if you have a large number of users and have had to increase the **number of locks** parameter to avoid running out of locks, use **sp_sysmon** to check the average hash chain length at peak periods. If the average length of the hash chains exceeds 4 or 5, consider increasing the value of **lock hashtable size** from its current setting to the next power of 2.

The hash chain length may be high during large insert batches, such as bulk copy operations. This is expected behavior, and does not require you to reset **lock hashtable size**.

lock scheme

lock scheme sets the default locking scheme to be used by **create table** and **select into** commands when a lock scheme is not specified in the command.

Summary Information	
Default value	allpages
Range of values	allpages, datapages, datarows
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Lock Manager

The values for lock scheme are character data, so you must use 0 as a placeholder for the second parameter, which must be numeric, and specify **allpages**, **datapages**, or **datarows** as the third parameter:

```
sp_configure "lock scheme", 0, datapages
```

lock shared memory

lock shared memory disallows swapping of SAP ASE pages to disk and allows the operating system kernel to avoid the server's internal page locking code. This can reduce disk reads, which are expensive.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)

Summary Information	
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Physical Memory

Not all platforms support shared memory locking. Even if your platform does, **lock shared memory** may fail due to incorrectly set permissions, insufficient physical memory, or for other reasons. See operating system documentation for your platform for information on shared memory locking.

lock spinlock ratio

For SAP ASEs running with multiple engines, **lock spinlock ratio** sets a ratio that determines the number of lock hash buckets that are protected by one spinlock. If you increase the value for **lock hashtable size**, the number of spinlocks increases, so the number of hash buckets protected by one spinlock remains the same.

Summary Information	
Default value	85
Range of values	1–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration groups	Lock Manager, Memory Use

SAP ASE manages the acquiring and releasing of locks using an internal hash table with a configurable number of hash buckets. On SMP systems, this hash table can use one or more spinlocks to serialize access between processes running on different engines. Use **lock hashtable size** to set the number of hash buckets.

The SAP ASE default value for **lock spinlock ratio** is 85. With **lock hashtable size** set to the default value of 2048, the default spinlock ratio defines 26 spinlocks for the lock hash table. See *System Administration Guide: Volume 2 > Managing Mulitprocessor Servers*.

sp_sysmon reports on the average length of the hash chains in the lock hash table. See the *Performance and Tuning Series: Monitoring Adaptive Server with sp_sysmon*.

lock table spinlock ratio

For SAP ASEs running with multiple engines, **table lock spinlock ratio** sets the number of rows in the internal table locks hash table that are protected by one spinlock.

Summary Information	
Default value	20
Range of values	1–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Lock Manager

SAP ASE manages the acquiring and releasing of table locks using an internal hash table with 101 rows (known as hash buckets). This table can use one or more spinlocks to serialize access between processes running on different engines.

The SAP ASE default value for **table lock spinlock ratio** is 20, which defines 6 spinlocks for the table locks hash table. The first 5 spinlocks protect 20 rows each; the sixth spinlock protects the last row. If you specify a value of 101 or greater for **table lock spinlock ratio**, SAP ASE uses only 1 spinlock for the entire table.

lock timeout pipe active

lock timeout pipe active controls whether SAP ASE collects lock timeout messages.

Summary Information	
Default value	0 (off)
Range of values	1 (on), 0 (off)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

If **lock timeout pipe active** and **lock timeout pipe max messages** are enabled, SAP ASE collects the data for each lock timeout that occurs.

Retrieve the lock timeout messages from the `monLockTimeout` monitor table.

lock timeout pipe max messages

lock timeout pipe max messages controls the maximum number of rows per engine in the lock timeout pipe, which determines the maximum number of rows that can be returned by the `monLockTimeout` monitoring table.

Summary Information	
Default value	0
Range of values	0 – 2147483648
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

lock wait period

lock wait period limits the number of seconds that tasks wait to acquire a lock on a table, data page, or data row. If the task does not acquire the lock within the specified time period, SAP ASE returns error message 12205 to the user and rolls back the transaction.

Summary Information	
Default value	2147483647
Range of values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Lock Manager

The **lock wait** option of the **set** command sets a session-level number of seconds that a task waits for a lock. It overrides the server-level setting for the session.

lock wait period, used with the session-level setting **set lock wait nnn**, is applicable only to user-defined tables. These settings have no influence on system tables.

At the default value, all processes wait indefinitely for locks. To restore the default value, reset the value to 2147483647 or enter:

```
sp_configure "lock wait period", 0, "default"
```

log audit logon failure

log audit logon failure specifies whether to log unsuccessful SAP ASE logins to the SAP ASE error log and, on Windows servers, to the Windows Event Log, if event logging is enabled.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Error Log

log audit logon success

log audit logon success specifies whether to log successful SAP ASE logins to the SAP ASE error log and, on Windows servers, to the Windows Event Log, if event logging is enabled.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Error Log

max async i/os per engine

max async i/os per engine specifies the maximum number of outstanding asynchronous disk I/O requests for a single engine at one time.

Summary Information	
Default value	Platform dependent
Range of values	1– platform-dependent value
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	O/S Resources

On the Linux Platform

On the Linux platform, **max async i/os per engine** controls the number of asynchronous IOs reserved from the operating system.

In threaded mode, SAP ASE reserves the value for **max async i/os per engine** for the engine server. In process mode, SAP ASE reserves the value for **max async i/os per engine** for each engine. For example, if you set **max async i/os per engine** to 4096 and there are 4 engines, in threaded mode SAP ASE reserves 4096 I/Os, but in process mode SAP ASE reserves 16384 I/Os.

Your system may benefit from using a number greater than the default value.

You can use **sp_sysmon** to help tune **max async i/os per engine**. **sp_sysmon**'s `disk i/o section` contains information about the maximum number of outstanding IOs for each engine during the sample period and the number of I/Os that were delayed because of engine or operating system limits. Generally, any I/Os delayed by engine limits indicate that you should increase the value of **max async i/os per engine**.

Whether SAP ASE can perform asynchronous IO on a device depends on whether or not this device support kernel asynchronous I/O (KAIO). The Linux kernel requires that you implement kernel asynchronous I/O support at the file system level. Most major file systems provide support for kernel asynchronous I/O, including ext3, xfs, jfs, and raw devices. The tmpfs file system does not support kernel asynchronous I/O. If the device does not support kernel asynchronous I/O, SAP ASE cannot perform asyncronous IO on that device, and instead reverts to standard synchronous IO for all reads and writes to that device. SAP ASE prints a message similar to the following in the error log indicating that the device has switched to synchronous IO:

```
00:00000:00001:2006/12/15 11:47:17.98 kernel  Virtual device
'/dev/shm/tempdb.dat' does not support kernel asynchronous i/o.
Synchronous i/o
will be used for this device.
```

max async i/os per server

max async i/os per server specifies the maximum number of asynchronous disk I/O requests that can be outstanding for SAP ASE at one time.

Summary Information	
Default value	Platform dependent
Range of values	1– platform dependent value
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	O/S Resources

This limit is not affected by the number of online engines per SAP ASE. **max async i/os per engine** limits the number of outstanding I/Os per engine.

Most operating systems limit the number of asynchronous disk I/Os that can be processed at any one time; some operating systems limit the number per operating system process, some limit the number per system, and some do both. If an application exceeds these limits, the operating system returns an error message. Because operating system calls are relatively expensive, it is inefficient for SAP ASE to attempt to perform asynchronous I/Os that get rejected by the operating system.

To avoid this, SAP ASE maintains a count of the outstanding asynchronous I/Os per engine and per server; if an engine issues an asynchronous I/O that would exceed either **max async i/os per engine** or **max async i/os per server**, SAP ASE delays the I/O until enough outstanding I/Os have completed to fall below the exceeded limit.

For example, assume an operating system limit of 200 asynchronous I/Os per system and 75 per process and an SAP ASE with three online engines. The engines currently have a total of 200 asynchronous I/Os pending, distributed according to the following table:

Engine	Number of I/Os pending	Outcome
0	60	Engine 0 delays any further asynchronous I/Os until the total for the server is under the operating system per-system limit and then continues issuing asynchronous I/Os.
1	75	Engine 1 delays any further asynchronous I/Os until the per-engine total is under the operating system per-process limit and then continues issuing asynchronous I/Os.
2	65	Engine 2 delays any further asynchronous I/Os until the total for server is under the operating system per-system limit and then continues issuing asynchronous I/Os.

All I/Os (both asynchronous and synchronous) require a disk I/O structure, so the total number of outstanding disk I/Os is limited by the value of **disk i/o structures**. It is slightly more efficient for SAP ASE to delay the I/O because it cannot get a disk I/O structure than because the I/O request exceeds **max i/os per server**. Set **max async i/os per server** equal to the value of **disk i/o structures.**

If the limits for asynchronous I/O can be tuned on your operating system, make sure they are set high enough for SAP ASE. There is no penalty for setting them as high as needed.

Use **sp_sysmon** to see if the per server or per engine limits are delaying I/O on your system. If **sp_sysmon** shows that SAP ASE exceeded the limit for outstanding requests per engine or per server, raise the value of the corresponding parameter. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

max buffers per lava operator

max buffers per lava operator sets an upper limit for the number of buffers used by lava operators that perform sorting or hashing (which are “expensive” in terms of processing). Lava operators use buffers from the session’s `tempdb` data cache pool as a work area for processing rows.

Summary Information	
Default value	2048
Range of values	500 – 65535
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

Lava operators often recurse through their input streams. Sorting requires subsequent merge passes until there are enough buffers available to merge all of the remaining runs. Hashing requires subsequent passes to build hash tables on any spilled sets until all of the remaining data can fit into an in-memory hash table. Some queries require less I/O if you increase **max buffers per lava operator**. This is particularly true for queries that use the `HASH DISTINCT`, `HASH VECTOR AGGREGATE`, and `HASH UNION` operators.

Be careful when you increase the default value of **max buffers per lava operator** for servers with many concurrent users: SAP ASE may allocate more buffers solely for expensive operators, reducing the number of buffers available for caching user’s tables and other session’s worktables. Use **sp_sysmon** to analyze **tempdb**'s data caching effectiveness.

max buffers per lava operator works with **max resource granularity** to limit the number of buffers used. The limit is set to the minimum of:

- The value of **max buffers per lava operator**, or,
- (**max resource granularity**) X (the number of data buffers in `tempdb`’s pagesize pool)

Differences Between number of sort buffers and max data buffers per lava sort

There are a number of differences between **number of sort buffers** and **max data buffers per lava sort**.

- **number of sort buffers** – affects parallel sorts only from the **create index**, and **update statistics** commands:
  - The buffers are sourced from single-page default or named data caches.
  - SAP ASE must have a sufficient number of free procedure cache buffers available for storing the metadata per row of data being sorted, otherwise it may abort the query processing.

- The SAP ASE sort manager verifies the amount of procedure cache space required with this formula:

  (**number of sort buffers**) X (*rows_per_page*) X 100

  Where *rows_per_page* is the number of rows of the sorted columns that can fit into a single-page data buffer.

- **max data buffers per lava sort** – affects only regular lava `sort` or `hash` operations:
  - The buffers are sourced from the single-page pool from the `tempdb` data cache assigned to the current session.
  - The query processor uses the smaller of the following values as the upper limit for buffer space per lava operation:
    - The value of **max data buffers per lava operator**
    - The product of:

      max resource granularity X (the number of cache buffers in the tempdb's single-page pool)

Note: The differences betwen **max data buffers per lava sort** and **number of sort buffers** influences the lava operator during optimization. However, the query processor cannot predict the number of available cache buffers at runtime. If the Execute engine cannot find a sufficient number of `tempdb` single-page cache buffers, it may use system worktable for the sorting or hashing instead.

---

max cis remote connections

max cis remote connections specifies the maximum number of concurrent Client-Library connections that can be made to remote servers by Component Integration Services.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Component Integration Services

By default, Component Integration Services allows up to four connections per user to be made simultaneously to remote servers. If you set the maximum number of users to 25, as many as 100 simultaneous Client-Library connections are allowed by Component Integration Services.

If this number does not meet the needs of your installation, you can override the setting by specifying exactly how many outgoing Client-Library connections you want the server to be able to make at one time.

max concurrently recovered db

max concurrently recovered db determines the degree of parallelism.

Summary Information	
Default value	0
Valid values	1– number of engines at start-up minus 1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Backup/Recovery

The minimum value is 1, which uses serial recovery, but you can also use the default value of 0, to use a self-tuning approach.

The maximum value is the number of engines at start-up minus 1. max concurrently recovered db is also limited by the value of the configuration parameter number of open databases.

max js restart attempts

Restricts the number of restart attempts and prevents the Job Scheduler restart feature from going into an infinite loop.

Summary Information	
Default value	3
Valid values	1– 10
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	SQL Server Administration

The value 0 indicates that the Job Scheduler Auto restart feature is disabled.

max memory

max memory specifies the maximum amount of total physical memory that you can configure SAP ASE to allocate. **max memory** must be greater than the total logical memory consumed by the current configuration of SAP ASE.

Summary Information	
Default value	Platform-dependent

Summary Information	
Range of values	Platform-dependent minimum – 2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Physical Memory

There is no performance penalty for configuring SAP ASE to use the maximum memory available to it on your computer. However, assess the other memory needs on your system, or SAP ASE may not be able to acquire enough memory to start.

See *System Administration Guide: Volume 2 > Configuring Memory*.

If SAP ASE Cannot Start

When **allocate max shared memory** is set to 1, SAP ASE must have the amount of memory available that is specified by **max memory**.

If the memory is not available, SAP ASE does not start. If this occurs, reduce the memory requirements for SAP ASE by manually changing the value of **max memory** in the server's configuration file. You can also change the value of **allocate max shared memory** to 0 so that not all memory required by **max memory** is required at start-up.

You may also want to reduce the values for other configuration parameters that require large amounts of memory. Then restart SAP ASE to use the memory specified by the new values. If SAP ASE fails to start because the total of other configuration parameter values is higher than the **max memory** value, see *System Administration Guide: Volume 2 > Configuring Memory* for information about configuration parameters that use memory.

max native threads per engine

max native threads per engine defines the maximum number of native threads the server spawns per engine.

Summary Information	
Default value	50
Maximum values	50 – 1000
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	User Environment

When the limit for the native threads is reached, SAP ASE sessions that require a native thread sleep until another session releases a native thread.

Note: max native threads per engine is ignored in threaded mode.

max nesting level

maximum nesting level sets the maximum nesting level for stored procedures and triggers. Each increased nesting level requires about 160 bytes of additional memory. For example, if you increase the nesting level from 16 to 26, SAP ASE requires an additional 1600 bytes of memory.

Summary Information	
Default value	50
Range of values	16 – 100
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	User environment

In SAP ASE 15.0.3 and later, the maximum nesting level has been increased to 100, and the default value to 50.

max network packet size

max network packet size specifies the maximum network packet size that can be requested by clients communicating with SAP ASE.

Summary Information	
Default value	2048
Range of values	512–65024
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	Network Communication

If some of your applications send or receive large amounts of data across the network, these applications can achieve significant performance improvement by using larger packet sizes. Two examples are large bulk-copy operations and applications that read or write large `text`, `unitext`, and `image` values.

Generally, you want:

- **default network packet size** to be small for users who perform short queries, and
- **max network packet size** to be large enough to allow users who send or receive large volumes of data to request larger packet sizes.

max network packet size must always be as large as, or larger than, the **default network packet size**. Values that are not even multiples of 512 are rounded down.

For client applications that explicitly request a larger network packet size to receive, you must also configure **additional network memory**.

Open Client Server cannot accept a network packet size greater than 64K.

See **bcp** and **isql** in the *Utility Guide* for information on using larger packet sizes from these programs. Open Client Client-Library documentation includes information on using variable packet sizes.

Choosing Packet Sizes

For best performance, choose a server packet size that works efficiently with the underlying packet size on your network.

The goals are:

- Reducing the number of server reads and writes to the network
- Reducing unused space in network packets (increasing network throughput)

For example, if your network packet size carries 1500 bytes of data, setting the SAP ASE packet size to 1024 (512*2) will probably achieve better performance than setting it to 1536 (512*3). The figure below shows how four different packet size configurations would perform in such a scenario.

Figure 4: Factors in determining packet size

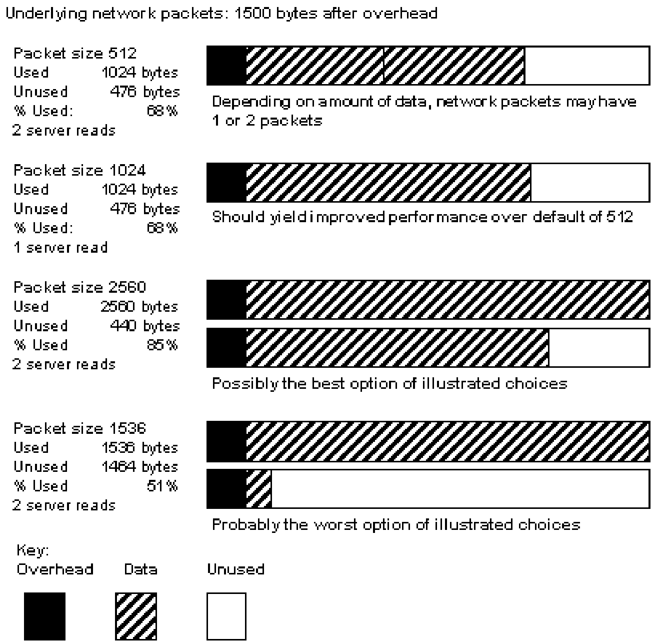

After you determine the available data space of the underlying packets on your network, perform your own benchmark tests to determine the optimum size for your configuration.

Use **sp_sysmon** to see how changing **max network packet size** affects network I/O management and task switching. For example, try increasing **max network packet size** and then checking **sp_sysmon** output to see how this affects **bcp** for large batches. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

max network peek depth

(UNIX only) **max network peek depth** specifies how many levels deep SAP ASE peeks into a connections operating system receive buffer for a pending cancel.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Network Communications

For example, if a client sends a new command followed by a cancel before SAP ASE finishes processing the current command, SAP ASE peeks into the operating system's receive buffer to the depth specified by **max network peek depth**. If the cancel occurs within the specified depth, both the current command and the command preceding the cancel are discarded, and SAP ASE waits for the next command.

max number network listeners

max number network listeners specifies the maximum number of network listeners allowed by SAP ASE at one time.

Summary Information	
Default value	5
Range of values	0–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration groups	Network Communication

Each master port has one network listener. Generally, there is no need to have multiple master ports, unless your SAP ASE must communicate over more than one network type. Some platforms support both socket and TLI (Transport Layer Interface) network interfaces (SAP ASE does not support the TLI interface in threaded mode). See the *Configuration Guide* for your platform for information on supported network types.

max online engines

max online engines places an upper limit of the number of engine threads that can be brought online. It does not take into account the number of CPUs available at start-up, and allows users to add CPUs at a later date.

Summary Information	
Default value	1
Range of values	1 – 1024 (threaded mode) 1 – 128 (process mode)
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, Processors

See *System Administration Guide: Volume 2 > Managing Mulitprocessor Servers* for a detailed discussion of how to set this parameter for your SMP environment.

The following is the sequence of events at SAP ASE start-up:

1. System databases are recovered on engine 0.
2. SAP ASE accepts user connections.
3. All engines that are configured to be online during start-up are brought online.
4. User databases are recovered in parallel by a "self-tuned" number of recovery tasks using the default data cache tuned for optimal recovery performance.

When tuning the **max engines online** parameter:

- Never have more engine threads online than there are CPUs.
- **max engines online** must be large enough to allow SAP ASE to bring all the engine thread pools online simultaneously.
- Depending on overall system load (including applications other than SAP ASE), you may achieve optimal throughput by leaving some CPUs free to run non-SAP ASE processes.
- You can achieve better throughput by running fewer engine threads with high CPU use, rather than by running more engine threads with low CPU use.
- Scalability is application-dependent. Conduct extensive benchmarks on your application to determine the best configuration of online engine threads.
- In process mode, use **sp_engine** to take engine threads offline or to bring them online. You can take all engine threads offline except engine zero.

See *Performance and Tuning Series: Basics > Using Engines and CPUs*.

max online Q engines

(Process mode only) **max online Q engines** specifies the maximum number of Q engines you can have online, and is required for MQ (see the *Active Messaging Users Guide*).

Summary Information	
Default value	0
Range of values	0 – 127
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Processors

The maximum number of Q engines depends on the value of **max online engines**:

- **max online Q engines** cannot be greater than **max online engines** minus **number of engines at startup**. That is, if the value for **max online engines** is 57, and the value for **number of engines at startup** is 28, the value for **max online Q engines** cannot be greater than 29.

- Setting **max online Q engines** reserves the high range of **max online engines** for Q engines. Once you set **max online Q engines**, SAP ASE engines can no longer use the engines in the range that is reserved for Q engines. For example, if you set **max online engines** to 10 and set **max online Q engines** to 4:
  - SAP ASE issues an error if you attempt to bind an engine group to engines 6, 7, 8, or 9
  - You cannot change the value of **number of engines at startup** to 7, 8, or 9
  - You cannot add engines 6, 7, 8, or 9 to an engine group

Because setting **max online Q engines** may affect existing production environments, SAP recommends that you increase the value of **max online engines** by the same amount that you increase **max online Q engines**. For example, if you increase **max online Q engines** by 4, also increase **max online engines** by 4.

max parallel degree

max parallel degree specifies the server-wide maximum number of worker processes allowed per query. This is called the "maximum degree of parallelism."

Summary Information	
Default value	1
Range of values	1–255
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Query Tuning

If **max parallel degree** is too low, the performance gain for a given query may not be as significant as possible; if **max parallel degree** is too high, the server may compile plans that require more processes than are actually available at execution time, or the system may become saturated, resulting in decreased throughput. To enable parallel partition scans, set this parameter to be equal to or greater than the number of partitions in the table you are querying.

The value of this parameter must be less than or equal to the current value of **number of worker processes**.

If you set **max parallel degree** to 1:

- SAP ASE scans all tables or indexes serially.
- SAP ASE forces serial query execution and the optimizer may select plans with a higher parallel degree than if it is disabled.

Changing **max parallel degree** causes all query plans in the procedure cache to be invalidated, and new plans are compiled the next time you execute a stored procedure or trigger.

See *Performance and Tuning Series: Query Processing and Abstract Plans > Parallel Sorting*

max pci slots

max pci slots sets the maximum number of PCI slots SAP ASE allows.

Summary Information	
Default value	0
Range of values	0 – 30
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

The values are:

- 0, 1 – default bridge with one PCA.

  Note: JVM support requires a single slot. Do not increase the number of slots.
- 2 – 30 – allocated for future releases.

For more information about PCI slots, see *Java in SAP Adaptive Server Enterprise*.

max query parallel degree

(Used when SAP ASE is in compatibility mode) Defines the number of worker processes to use for a given query.

Summary Information	
Default value	1
Range of values	1 – 255
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

This parameter is relevant only if you do not want to enable parallelism globally. The value for **number of worker process** cannot be less than the value for **max query parallel degree**.

See *Performance and Tuning Series: Query Processing and Abstract Plans > Parallel Query Processing*.

For more information about compatibility mode, see the *Migration Technology Guide*.

max repartition degree

max repartition degree configures the amount of dynamic repartitioning SAP ASE requires, which enables SAP ASE to use horizontal parallelism. However, if the number of partitions is too large, the system is flooded with worker processes that compete for resources, which degrades performance.

Summary Information	
Default value	1
Range of values	1 – value of **max parallel degree**
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

The value for **max repartition degree** enforces the maximum number of partitions created for these resources. If all of the tables and indexes are unpartitioned, SAP ASE uses the value for **max repartition degree** to provide the number of partitions to create as a result of repartitioning the data.

max resource granularity

max resource granularity indicates the maximum percentage of the system's resources a query can use.

Summary Information	
Default value	10
Range of values	1 – 100
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

max resource granularity is not enforced at execution time, but is only a guide for the query optimizer, and does not prevent the query processor from running queries in parallel. The query engine can avoid some memory-intensive strategies by using **max resource granularity** as a guide.

max scan parallel degree

max scan parallel degree specifies the server-wide maximum degree of parallelism for hash-based scans.

Summary Information	
Default value	1
Range of values	1–255
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Query Tuning

The degree of parallelism for hash-based scans may be used for the following access methods:

- Parallel index scans for partitioned and nonpartitioned tables
- Parallel table scans for nonpartitioned tables

max scan parallel degree applies per table or index; that is, if **max scan parallel degree** is 3, and one table in a join query is scanned using a hash-based table scan and the second can best be accessed by a hash-based index scan, the query can use 9 worker processes (as long as **max parallel degree** is set to 9 or higher).

The optimizer uses **max scan parallel degree** as a guideline when it selects the number of processes to use for parallel, nonpartition-based scan operations. It does not apply to parallel sort. Because there is no partitioning to spread the data across devices, parallel processes can be accessing the same device during the scan. This can cause additional disk contention and head movement, which may degrade performance. To prevent multiple disk accesses from becoming a problem, use **max scan parallel degree** to reduce the maximum number of processes that can access the table in parallel.

If this number is too low, the performance gain for a given query is not as significant as possible; if the number is too large, the server may compile plans that use enough processes to make disk access less efficient. A general rule is to set this parameter to no more than 2 or 3, because it takes only 2 to 3 worker processes to fully utilize the I/O of a given physical device.

Set the value of **max scan parallel degree** to less than or equal to the current value of **max parallel degree**. SAP ASE returns an error if you specify a number larger than the **max parallel degree** value.

If you set **max scan parallel degree** to 1, SAP ASE does not perform hash-based scans.

Changing **max scan parallel degree** causes all query plans in the procedure cache to be invalidated, and new plans are compiled the next time you execute a stored procedure or trigger.

max SQL text monitored

max SQL text monitored specifies the amount of memory allocated per user connection for saving SQL text to memory shared by Adaptive Server Monitor.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Monitoring

If you do not allocate enough memory for the batch statements, the text you want to view may be truncated. SAP recommends that you use an initial value of 1024 bytes of memory per user connection.

The total memory allocated from shared memory for the SQL text is the product of **max SQL text monitored** multiplied by the currently configured number of user connections.

max transfer history

max transfer history controls how many transfer history entries SAP ASE retains in the `spt_TableTransfer` table in each database.

Summary Information	
Default value	10
Range of values	1 – 255
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Adaptive Server Administration

For each table tracked, `spt_TableTransfer` retains:

(*N* successful entries) + (*N* unsuccessful entries)

Where *N* is the value for **max transfer history**.

Lowering this parameter does not automatically remove any entries from `spt_TableTransfer`. Entries are removed for a given transferred table the next time you initiate a transfer for that table. The table's successful transfer entries are cleared if the transfer succeeds. If the transfer is unsuccessful, its failed transfer entries are cleared.

For example, if a table has 12 successful and 9 unsuccessful history entries in `spt_configure`, and you change **max transfer history** to 5, the next successful transfer of that table places 5 successful entries in `spt_configure`, but `spt_configure` retains the previous 9 failed entries.

max utility parallel degree

max util parallel degree specifies the server-wide maximum number of worker processes allowed per query used by the **create index with consumers** and **update stats with consumers** commands.

Summary Information	
Default value	1
Valid values	1– 255
Status	Dynamic
Display level	Basic
Required role	Query Tuning
Configuration group	Application Functionality

maximum dump conditions

maximum dump conditions sets the maximum number of conditions you can specify under which SAP ASE generates a dump of data in shared memory.

Summary Information	
Default value	10
Range of values	10–100
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	Group Diagnostics

Note: This parameter is included for use only by SAP Technical Support. Do not modify it unless you are instructed to do so by SAP Technical Support.

maximum failed logins

maximum failed logins allows you to set the server-wide maximum number of failed login attempts for logins and roles.

Summary Information	
Default value	0
Range of values	-1 – 32767
Status	Dynamic
Display level	10
Required role	System security officer
Configuration group	Security Related

A value of -1 indicates that the failed login count in the `syslogins` column `logincount` is updated whenever an authentication failure occurs, but that the account is not locked. Compare with a 0 (zero) value, which avoids incrementing the column for every failed authentication and avoids locking the account due to authentication failures.

See the *Password Complexity Checks* in the *Security Administration Guide* for information about password checks and policies for authentication.

maximum job output

maximum job output sets limit, in bytes, on the maximum output a single job can produce.

Summary Information	
Default value	32768
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

If a job produces more output than specified in **maximum job output**, all the data returned above the value you enter is discarded.

memory alignment boundary

memory alignment boundary determines the memory address boundary on which data caches are aligned.

Summary Information	
Default value	Logical page size
Range of values	2048[a] – 16384 (the minimum determined by server's logical page size)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Cache Manager

Some machines perform I/O more efficiently when structures are aligned on a particular memory address boundary. To preserve this alignment, values for **memory alignment boundary** should always be powers of two between the logical page size and 2048K.

Note: The **memory alignment boundary** parameter is included for support of certain hardware platforms. Do not modify it unless you are instructed to do so by SAP Technical Support.

memory dump compression level

memory dump compression level controls the compression level for shared memory dumps.

Summary Information	
Default value	0
Valid values	0–9
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Diagnostics

The compression levels range from 0 (no compression) to 9 (highest compression). The speed of the compression is inversely proportional to the amount the dump is compressed. The lower the compression level, the faster SAP ASE compresses the dump, but the size of the compressed file may be greater.

memory per worker process

memory per worker process specifies the amount of memory, in bytes, used by worker processes.

Summary Information	
Default value	1024
Range of values	1024–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Memory Use

Each worker process requires memory for messaging during query processing. This memory is allocated from a shared memory pool; the size of this pool is **memory per worker process** multiplied by **number of worker processes**. For most query processing, the default size is more than adequate. If you use **dbcc checkstorage**, and have set **number of worker processes** to 1, you may need to increase **memory per worker process** to 1792 bytes.

See *System Administration Guide: Volume 2 > Configuring Memory*.

messaging memory

messaging memory configures the amount of memory available for SAP messaging.

Summary Information	
Default value	400
Range of values	60 – 2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Memory Use, Physical Memory

metrics elap max

metrics elap max configures maximum elapsed time and thresholds for QP metrics

Summary Information	
Default value	0
Range of values	0 – 2147483647

Summary Information	
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

metrics exec max

metrics exec max configures maximum execution time and thresholds for QP metrics.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

metrics lio max

metrics lio max configures maximum logical I/O and thresholds for QP metrics.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

metrics pio max

metrics pio max configures maximum physical I/O and thresholds for QP metrics.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Dynamic

Summary Information	
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

min pages for parallel scan

min pages for parallel scan controls the number of tables and indexes that SAP ASE can access in parallel.

Summary Information	
Default value	200
Range of values	20 - 2147483647
Status	dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

If the number of pages in a table is below the value you set, the table is accessed serially. **min pages for parallel scan** does not consider page size. If SAP ASE accesses the indexes and tables, SAP ASE attempts to repartition the data, if that is appropriate, and to use parallelism above the scans, if that is appropriate.

minimum password length

minimum password length allows you to customize the length of server-wide password values or per-login or per-role password values.

Summary Information	
Default value	6
Range of values	0 – 30
Status	Dynamic
Display level	10
Required role	System security officer
Configuration group	Security Related

See *Set or Change the Minimum Password Length* in the *Security Administration Guide*.

mnc_full_index_filter

mnc_full_index_filter prevents SAP ASE from considering noncovered indexes that do not have a limiting search argument at the server level, if there is a column in the index or a predicate that does not have a histogram.

Summary Information	
Default value	2
Range of values	0 – 2 • 0 – disable. • 1 – enable. • 2 – set according to the optimization goal setting.
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

You can use **mnc_full_index_filter** on data-only-locked (DOL) tables in which you have the intelligent index scan, even though the intelligent index scan manufactures search arguments.

Changing the value of **mnc_full_index_filter** does not increase the amount of memory SAP ASE uses.

mnc_full_index_filter is not enabled for any specific optional goal; the only way to obtain the behavior is to explicitly enable it.

msg confidentiality reqd

msg confidentiality reqd requires that all messages into and out of SAP ASE be encrypted. The **use security services** parameter must be 1 for messages to be encrypted.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

msg integrity reqd

msg integrity reqd requires that all messages be checked for data integrity. **use security services** must be 1 for message integrity checks to occur.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

If **msg integrity reqd** is set to 1, SAP ASE allows the client connection to succeed unless the client is using one of the following security services: **message integrity**, **replay detection**, **origin checks**, or **out-of-seq checks**.

net password encryption reqd

net password encryption reqd restricts login authentication to use only RSA encryption algorithm or the SAP proprietary algorithm.

Summary Information	
Default value	0
Range of values	0 – 3
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

Table 4. Values and descriptions for net password encryption reqd

Value	Description
0	Allows the client to choose the encryption algorithm used for login passwords on the network, including no password encryption.

Value	Description
1	Restricts clients to use either RSA or SAP proprietary encryption algorithms to encrypt login passwords on the network. This provides an incrementally restrictive setting that allows clients who have previously connect to reconnect with the SAP proprietary algorithm and new clients to connect with the stronger RSA algorithm. A client that attempts to connect without using password encryption fails.
2	Restricts clients to use only the RSA encryption algorithms to encrypt login passwords on the network. This provides strong RSA encryption of passwords. Clients that attempt to connect without using the RSA encryption fail.
3	SAP ASE allows only incoming clients that use the EPEP login protocol. The values 0, 1, and 2 also allow EPEP login protocol to be used when a client that supports the login protocol attempts to use it with an SAP ASE that implements the EPEP login protocol.

Note: Setting the value to 2 or 3 increases network memory to support the maximum configured connections using this protocol. **additional network memory** dynamically adds more memory to the network memory pool used by EPEP. When the value is set to 3, the KPP Handler goes into sleep status. This is because there is no need to provide new RSA key pair for every connection. Use **sp_who** to check the KPP Handler status.

When a connection is refused because network password encryption is required, the client receives:

```
Msg 1640, Level 16, State 2:
Adaptive Server requires encryption of the login
password on the network.
```

network polling mode

network polling mode configures the SAP ASE network polling mode.

Summary Information	
Default value	threaded
Range of values	threaded, inline, or compact
Status	Static
Display level	Basic
Required role	System administrator
Configuration group	Network Communication

When **network polling mode** is set to:

- **threaded** – SAP ASE spawns a separate thread for each network task configured that performs polling.

- **inline** – one of the engines performs the polling.
- **compact** – each engine creates its own network controller to perform its polling. SAP recommends that you set **network polling mode** to **compact** when there are multiple engine groups, and the load is distributed across the engines.

In-Line Network Polling

SAP ASE spawns separate threads to perform network polling.

However, separate threads may not show significant performance gains when running with a low engine configuration. SAP ASE may suffer from contention between engines and the network threads when it runs with a low engine configuration on machines with a lower number of CPUs.

Set the **network polling mode** configuration parameter to **inline** to avoid this contention. When you set **network polling mode** to **inline**, one engine performs polling, eliminating the need for a separate thread. On the Linux platform, setting the **network polling mode** to **inline** reduces CPU usage by using high-resolution sleep APIs.

Note: When you set **network polling mode** to **inline**, the **idle timeout** configuration parameter controls the number of times an SAP ASE engine loops while looking for a runnable task before releasing the CPU to the operating system.

number of alarms

number of alarms specifies the number of alarm structures allocated by SAP ASE.

Summary Information	
Default value	40
Range of values	40 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, SQL Server Administration

The Transact-SQL command **waitfor** defines a specific time, time interval, or event for the execution of a statement block, stored procedure, or transaction. SAP ASE uses alarms to correctly execute **waitfor** commands. Other internal processes require alarms.

When SAP ASE needs more alarms than are currently allocated, this message is written to the error log:

```
uasetalarm: no more alarms available
```

The number of bytes of memory required for each alarm structure is small. If you raise the **number of alarms value** significantly, adjust **max memory** accordingly.

number of aux scan descriptors

number of aux scan descriptors sets the number of auxiliary scan descriptors available in a pool shared by all users on a server.

Summary Information	
Default value	200
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, SQL Server Administration

Each user connection and each worker process has 48 scan descriptors exclusively allocated to it. Of these, 16 are reserved for user tables, 12 are reserved for worktables, and 20 are reserved for system tables (with 4 of these set aside for rollback conditions). A descriptor is needed for each table referenced, directly or indirectly, by a query. For user tables, a table reference includes:

- All tables referenced in the **from** clause of the query
- All tables referenced in a view named in the query (the view itself is not counted)
- All tables referenced in a subquery
- All tables that need to be checked for referential integrity (these are used only for inserts, updates, and deletes)
- A table created with **select...into**
- All worktables created for the query

If a table is referenced more than once (for example, in a self-join, in more than one view, or in more than one subquery) the table is counted each time. If the query includes a **union**, each **select** statement in the **union** query is a separate scan. If a query runs in parallel, the coordinating process and each worker process needs a scan descriptor for each table reference.

When the number of user tables referenced by a query scan exceeds 16, or the number of worktables exceeds 12, scan descriptors from the shared pool are allocated. Data-only-locked tables also require a system table descriptor for each data-only-locked table accessed with a table scan (but not those accessed with an index scan). If more than 16 data-only-locked tables are scanned using table scans in a query, auxiliary scan descriptors are allocated for them.

If a scan needs auxiliary scan descriptors after it has used its allotted number, and there are no descriptors available in the shared pool, SAP ASE displays an error message and rolls back the user transaction.

If none of your queries need additional scan descriptors, you may still want to leave **number of aux scan descriptors** set to the default value in case your system requirements grow. Set it to 0 only if you are sure that users on your system will never run queries on more than 16 tables and that your tables will always have few or no referential integrity constraints.

If your queries need more scan descriptors, use one of these methods to remedy the problem:

- Rewrite the query, or break it into steps using temporary tables. For data-only-locked tables, consider adding indexes if there are many table scans.
- Redesign the table's schema so that it uses fewer scan descriptors, if it uses a large number of referential integrity constraints. You can find how many scan descriptors a query would use by enabling **set showplan, noexec on** before running the query.
- Increase the **number of aux scan descriptors** setting.

The following sections describe how to use **sp_monitorconfig** to monitor the current and high-water-mark usage to avoid running out of descriptors, and how to estimate the number of scan descriptors you need.

Monitoring and Estimating Scan Descriptor Usage

sp_monitorconfig reports the number of unused (free) scan descriptors, the number of auxiliary scan descriptors currently being used, the percentage that is active, and the maximum number of scan descriptors used since the server was last started. Run it periodically, at peak periods, to monitor scan descriptor use.

sp_monitorconfig reports the number of unused (free) scan descriptors, the number of auxiliary scan descriptors currently being used, the percentage that is active, and the maximum number of scan descriptors used since the server was last started. Run it periodically, at peak periods, to monitor scan descriptor use.

This example shows scan descriptor use with 500 descriptors configured:

```
sp_monitorconfig "aux scan descriptors"
```

```
Usage information at date and time: Apr 22 2002  2:49PM.
Name                      Num_free    Num_active  Pct_act  Max_Used
Reuse_cnt   Instance_Name
------------------------  ----------  ----------  -------  --------
---
----------  --------------------
number of aux                  260        240    48.00         427
         0                 NULL
```

Only 240 auxiliary scan descriptors are being used, leaving 260 free. However, the maximum number of scan descriptors used at any one time since the last time SAP ASE was started is 427, leaving about 20 percent for growth in use and exceptionally heavy use periods. "Re-used" does not apply to scan descriptors.

To estimate the scan descriptor usage:

1. Determine the number of table references for any query that references more than 16 user tables, or for those that have a large number of referential constraints, by running the query with **set showplan** and **set noexec** enabled. If auxiliary scan descriptors are required, **showplan** reports the number needed:

```
Auxiliary scan descriptors required: 17
```

The reported number includes all auxiliary scan descriptors that are required for the query, including those for all worker processes. If your queries involve only referential constraints, you can also use **sp_helpconstraint**, which displays a count of the number of referential constraints per table.

2. For each query that uses auxiliary scan descriptors, estimate the number of users who would run the query simultaneously and multiply. If 10 users are expected to run a query that requires 8 auxiliary descriptors, a total of 80 will be needed at any one time.
3. Add the per-query results to calculate the number of needed auxiliary scan descriptors.

number of backup connections

number of backup connections sets the maximum number of user connections Backup Server establishes to dump or load in-memory databases.

Summary Information	
Default value	0
Range of values	1 – 32768
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	User Environment

The value of **number of backup connections** restricts the maximum number of stripes for an archived database because Backup Server requires one user connection per stripe when you run **dump** or **load database**, and requires an extra connection to run the **dump database** command.

number of backup connections is a limit, and does not consume any resources. Setting **number of backup connections** to 0 means that Backup Server can use the maximum number of user connections.

number of ccbs

Reserved for future use.

Summary Information	
Default value	0
Range of values	0 – 100
Status	Static
Display level	

Summary Information	
Required role	
Configuration group	Diagnostics

number of checkpoint tasks

number of checkpoint tasks configures parallel checkpoints. The value of number of checkpoint tasks must be less than or equal to the value of **number of engines at start-up**.

Summary Information	
Default value	1
Valid values	1– 8
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Backup/Recovery

The maximum value is limited by the value of the configuration parameters **number of engines online at startup** and **number of open databases**, with a maximum of 8. The default value sets serial checkpoints as the default behavior.

number of devices

number of devices controls the number of database devices SAP ASE can use. It does not include devices used for database or transaction log dumps.

Summary Information	
Default value	10
Range of values	1–2,147,483,647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Disk I/O, Memory Use

When you execute **disk init**, you can also assign the virtual device number (the `vdevno`), although this value is optional. If you do not assign the **vdevno**, SAP ASE assigns the next available virtual device number.

If you do assign the virtual device number, each device number must be unique among the device numbers used by SAP ASE. The number 0 is reserved for the master device. You can enter any unused device number that falls in the valid range of values.

To determine which numbers are currently in use, enter:

```
select vdevno from master..sysdevices
   where status & 2 = 2
```

Here, “status 2” specifies physical disk.

Setting the number of devices on UNIX

If you are using a large number of devices on UNIX platforms, SAP recommends that you set the appropriate number of devices and user connections in the configuration file. Attempting to configure a large number of devices dynamically using **sp_configure** may fail.

The **number of network connections** and **number of devices** configuration parameters use operating system file descriptors. You must increase the number of operating system file descriptors to use a large value for either configuration parameter.

What constitutes a large number of devices is platform-specific. These are the upper limits for the number of file descriptors for SAP ASE:

- HPIA 64, HPPA 64, AIX 64, Windows – Set by SAP ASE.
- IBM RISC 64-bit – 2000
- IBM PLinux – 2048
- AMD64 Linux – 10000
- IA32-bit Linux – 10000
- IA 64-bit Linux – 1024

Note: Changing the number of file descriptors requires you restart SAP ASE for the change to take affect.

If there are insufficient number of file descriptors, SAP ASE issues an error message similar to:

```
Msg 5893, Level 16, State 1:Procedure 'sp_configure', Line 1234:
The sum, ( ((number of user connections) + (number of remote sites) +
(max cisremote connections) + (number of java sockets)) / (max online
engines)) + (number of devices ) + (max number network listeners),
must be no greater than '958'.
Msg 5849, Level 16, State 1:Procedure 'sp_configure', Line 1234:
Verification failed for parameter 'number of devices'.
(return status = 1)
```

If you then set the number of file descriptors to 2048 with the command `ulimit -n 2048` and restart the server, SAP ASE reports this message in the error log when you increase the value for the **number of devices**:

```
kernel Using 2048 file descriptors.
```

number of disk tasks

number of disk tasks controls the number of tasks dedicated to polling and completing disk I/Os.

Summary Information	
Default value	1
Range of values	UNIX: 1– 64 Windows: 1
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Disk I/O, Processors

number of dtx participants

number of dtx participants sets the total number of remote transactions that the SAP ASE transaction coordination service can propagate and coordinate simultaneously.

Summary Information	
Default value	500
Valid values	100 – 2147483647
Status	Dynamic
Display level	10
Required role	System administrator
Configuration groups	DTM Administration, Memory Use

A DTX participant is an internal memory structure that the coordination service uses to manage a remote transaction branch. As transactions are propagated to remote servers, the coordination service must obtain new DTX participants to manage those branches.

Setting **number of dtx participants** to a number smaller than the default reduces the number of remote transactions that the server can manage. If no DTX participants are available, new distributed transactions cannot start. In-progress distributed transactions may abort if no DTX participants are available to propagate a new remote transaction.

Setting **number of dtx participants** to a number larger than the default increases the number of remote transaction branches that SAP ASE can handle, but also consumes more memory.

Optimizing the Number of DTX Participants for Your System

During a peak period, use **sp_monitorconfig** to examine the use of DTX participants.

For example:

```
sp_monitorconfig "number of dtx participants"
```

```
Usage information at date and time: Apr 22 2002  2:49PM.
Name                          Num_free    Num_active  Pct_act  Max_Used
Reuse_cnt   Instance_Name
------------------------   ----------  ----------  -------  --------
---
----------  --------------------
number of dtx                        80          20     4.00         210
         0                    NULL
```

If the `num_free` value is zero or very low, new distributed transactions may be unable to start due to a lack of DTX participants. Consider increasing the **number of dtx participants** value.

A low `Max_used` value may indicate that unused DTX participants are consuming memory that could be used by other server functions. Consider reducing the value of **number of dtx participants**.

number of dump threads

number of dump threads controls the number of threads that SAP ASE spawns to perform a memory dump.

Summary Information	
Default value	Disabled
Range of values	1 (disabled, no parallelism) – 8 (fully parallel)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Group Diagnostics

Using the appropriate value for number of dump threads can reduce the amount of time the engines are halted during the memory dump.

When you are determining the number of threads for memory:

- Use a value of 8 if the machine has enough free memory for the file system cache to hold the entire memory dump.
- If you do not know whether the machine has enough free memory, the value for number of dump threads depends on many factors, including the speed of the I/O system, the speed of

the disks, the controller’s cache, whether the dump file lives in a logical volume manager created on several disks, and so on.

- Disable parallel processing (by assigning a value of 1) if you do not halt the engines when performing memory dumps, described below.

  When SAP ASE performs a memory dump, the number of files it creates is the sum of the number of memory segments that it has allocated multiplied by the number of threads configured. SAP ASE uses separate threads to write on separate files. When this job completes, the engines are restarted, and the files are merged into the target dump file. Because of this, the time to dump the shared memory in parallel is greater than doing it serially.
- If you halt the engines during the memory dump, using a value other than 1 may reduce the amount of time the engines spend stopped while dumping the memory.

number of early send rows

number of early send rows configures the number of rows that are sent to the client in the first packet of a new result set.

Summary Information	
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Network Communication

SAP ASE assembles data and accumulates rows that wait in the send buffer, but does not send them until the buffer has reached its configured size. If the packet size is large and if the query is one that takes a long time to execute, these rows may sit for sometime before they are sent. When the rows are sent, they arrive at the client simultaneously, and the client then spends time processing them individually. For large data transfers, the larger the value for the network packet size the more efficient the overall transmission, but there is a time lag during which the client is idle while the server waits for the first buffer to fill. Reducing the size of the network packet size can improve response time, but at the cost of efficiency.

The **number of early send rows** and **early row send increment** configuration parameters allow you to configure the number of rows that sit in the send buffer, and can reduce the amount of wait time:

- **number of early send rows** – initial number of rows sent to a client.
- **early row send increment** – number by which the server increments the value for the number of early send rows when the packet fills and there are still rows waiting in the send buffer.

number of early send rows allows the server to send results to the client before the packet has reached its configured size. Each subsequent packet increments the number of rows until the packets has reached its configured size (as set by the network packet size parameter). This gives the benefit of the reduced latency for the first few rows without losing overall transmission efficiency. The cycle continues for each new result set, starting with smaller packets, then growing to larger ones.

For example, if you set the value for **number of early send rows** to 2, the server sends the rows in the buffer when it contains 2 rows. The client receives the initial rows more quickly and processes them while the server is working on the next set of rows in the query. If you subsequently set the value for **early row send increment** to 20,000, the server sends the second set of rows in the buffer when it contains 2 + 20,000 rows and adds the value 20,000 to the value for number of early send rows for each subsequent buffer: that is, 2 + 20,000 + 20,000 for the third buffer and so on.

See also

- *early row send increment* on page 102

number of engines at startup

number of engines at startup is used exclusively during start-up to set the number of engines brought online.

It allows great flexibility, subject to the restriction that you cannot set the value of **number of engines at startup** to a value greater than the number of CPUs on your machine, or to a value greater than the configuration of **max online engines**.

Summary Information	
Default value	1
Range of values	1 – number of CPUs on machine
Status	Static
Display level	Basic
Required role	System administrator
Configuration groups	Java Services, Memory Use, Processors

Note: When configured for threaded mode, SAP ASE ignores the **number of engines at startup** configuration parameter. In threaded mode, SAP ASE uses the size of the defined thread pools to determine the number of online engines at startup. If the configuration file contains no thread pool configuration information (for example, while you are upgrading from an earlier version), SAP ASE uses the existing value for **number of engines at startup** to determine the size the default engine pool.

Users who do not intend to bring engines online after start-up should set **max online engines** and **number of engines at startup** to the same value. A difference between **number of**

engines at startup and **max online engines** wastes approximately 1.8 MB of memory per engine.

SAP ASE allows users to take all engines offline, except engine zero.

number of histogram steps

number of histogram steps specifies the number of steps in a histogram.

Summary Information	
Default value	20
Range of values	3 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

number of index trips

number of index trips specifies the number of times an aged index page traverses the most recently used/least recently used (MRU/LRU) chain before it is considered for swapping out.

Summary Information	
Default value	0
Range of values	0–65535
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Cache Manager

As you increase the value of **number of index trips**, index pages stay in cache for longer periods of time.

A data cache is implemented as an MRU/LRU chain. As the user threads access data and index pages, these pages are placed on the MRU end of the cache's MRU/LRU chain. In some high transaction environments (and in some benchmarks), you may want to keep index pages in cache, since they will probably be needed again soon. Setting **number of index trips** higher keeps index pages in cache longer; setting it lower allows index pages to be swapped out of cache sooner.

You need not set the **number of index trips** for relaxed LRU pages. See *System Administration Guide: Volume 2 > Configuring Data Caches*.

Note: If the cache used by an index is relatively small (especially if it shares space with other objects) and you have a high transaction volume, do not set **number of index trips** too high. The cache can flood with pages that do not age out, and this may lead to the timing out of processes that are waiting for cache space. Before changing the value of **number of index trips** to a number other than 0, make sure the application has sufficient cache to store all index, OAM, and data pages. Consult SAP Technical Support before changing the value of **number of index trips**.

number of java sockets

number of java sockets enables the Java VM and the **java.net** classes SAP supports.

Summary Information	
Default value	0
Valid values	0 – 32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Java Services, Memory Use

number of large i/o buffers

number of large i/o buffers sets the number of allocation unit-sized buffers reserved for performing large I/O for certain SAP ASE utilities.

Summary Information	
Default value	6
Valid values	1–256
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Disk I/O, Memory Use, SQL Server Administration

These large I/O buffers are used primarily by the **load database** command, which uses one buffer to load the database, regardless of the number of stripes it specifies. **load database** then uses as many as 32 buffers to clear the pages for the database it is loading. These buffers are not used by **load transaction**. To perform more than six **load database** commands concurrently, configure one large I/O buffer for each **load database** command.

create database and **alter database** use these buffers for large I/O while clearing database pages. Each instance of **create database** or **load database** can use as many as 32 large I/O buffers.

These buffers are also used by disk mirroring and by some **dbcc** commands.

Note: In SAP ASE version 12.5.0.3 and later, the size of the large I/O buffers is one allocation (256 pages), not one extent (8 pages). The server thus requires more memory allocation for large buffers. For example, a disk buffer that required memory for 8 pages in earlier versions now requires memory for 256 pages.

number of locks

number of locks sets the total number of available locks for all users on SAP ASE.

Summary Information	
Default value	10000
Range of values	1000–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Lock Manager, Memory Use

The total number of locks needed by SAP ASE depends on the number of concurrent and parallel processes, and the types of actions performed by the transactions. To see how many locks are in use at a particular time, use **sp_lock**.

For serial operation, SAP suggests that you start by assigning 20 locks for each active, concurrent connection.

Parallel execution requires more locks than serial execution. For example, if you find that queries use an average of five worker processes, try increasing by one-third the **number of locks** configured for serial operation.

SAP ASE reports error 1204 when it runs out of locks. If users report lock errors, you may need to increase **number of locks**; but remember that locks use memory. See *System Administration Guide Volume 2 > Configuring Memory*. Use **sp_monitorconfig** to view the statistics for currently used locks. See *Troubleshooting: Error Messages Advanced Resolutions* for information about resolving error message 1204.

Note: Datarows locking may require that you change the value for **number of locks**. See the *Performance and Tuning Series: Locking and Concurrency Control*.

number of mailboxes

number of mailboxes specifies the number of mailbox structures allocated by SAP ASE.

Summary Information	
Default value	30

Summary Information	
Range of values	30–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, SQL Server Administration

Mailboxes, which are used with messages, are used internally by SAP ASE for communication and synchronization between kernel service processes. Mailboxes are not used by user processes. Do not modify this parameter unless instructed to do so by SAP Technical Support.

number of messages

number of messages specifies the number of message structures allocated by SAP ASE.

Summary Information	
Default value	64
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, SQL Server Administration

Messages, which are used with mailboxes, are used internally by SAP ASE for communication and synchronization between kernel service processes. Messages are also used to coordinate between a family of processes in parallel processing.

number of network tasks

number of network tasks controls the number of tasks dedicated to polling and completing network I/Os.

Summary Information	
Default value	1
Range of values	UNIX: 1– 64 Windows: 1
Status	Dynamic

Summary Information	
Display level	Basic
Required role	System administrator
Configuration groups	Disk I/O, Processors

number of network tasks functions only when **network polling mode** is set to threaded.

number of oam trips

number of oam trips specifies the number of times an *object allocation map* (OAM) page traverses the MRU/LRU chain before it is considered for swapping out. The higher the value of **number of oam trips**, the longer aged OAM pages stay in cache.

Summary Information	
Default value	0
Range of values	0–65535
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Each table, and each index on a table, has an OAM page, which holds information on pages allocated to the table or index and is checked when a new page is needed for the index or table. A single OAM page can hold allocation mapping for between 2,000 and 63,750 data or index pages.

The OAM pages point to the allocation page for each allocation unit where the object uses space. The allocation pages, in turn, track the information about extent and page usage within the allocation unit.

In some environments and benchmarks that involve significant allocations of space (that is, massive bulk copy operations), keeping OAM pages in cache longer improves performance. Setting **number of oam trips** to a higher value keeps OAM pages in cache.

Note: If the cache is relatively small and used by a large number of objects, do not set **number of oam trips** too high. This may result in the cache being flooded with OAM pages that do not age out, and user threads may begin to time out.

Before changing the value of **number of oam trips** to a number other than 0, make sure the application has sufficient cache to store all index, OAM, and data pages. Consult SAP Technical Support before changing the value of **number of oam trips**.

number of open databases

number of open databases sets the maximum number of databases that can be open simultaneously on SAP ASE.

Summary Information	
Default value	12
Range of values	6 –2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Meta-Data Caches, SQL Server Administration

When you calculate a value, include the system databases `master`, `model`, `sybsystemprocs`, and `tempdb`. If you have installed auditing, include the `sybsecurity` database. Also, count the sample databases `pubs2` and `pubs3`, the syntax database `sybsyntax`, and the **dbcc** database `dbccdb` if they are installed.

If you are planning to make a substantial change, such as loading a large database from another server, use **sp_helpconfig** to calculate an estimated metadata cache size by using **sp_helpconfig**. **sp_helpconfig** displays the amount of memory required for a given number of metadata descriptors, as well as the number of descriptors that can be accommodated by a given amount of memory. A database metadata descriptor represents the state of the database while it is in use or cached between uses.

Optimizing the number of open databases

If SAP ASE displays a message saying that you have exceeded the allowable number of open databases, adjust the value.

1. Use **sp_countmetadata** to find the total number of database metadata descriptors:

   ```
   sp_countmetadata "open databases"
   ```

   The best time to run **sp_countmetadata** is when there is little activity on the server. Running **sp_countmetadata** during a peak time can cause contention with other processes.

   Suppose SAP ASE reports the following information:

   ```
   There are 50 databases, requiring 1719 Kbytes of memory. The 'open
   databases' configuration parameter is currently set to 500.
   ```

2. Configure number of open databases with the value of 50:

   ```
   sp_configure "number of open databases", 50
   ```

This new configuration number is only a starting point; base the ideal size on the number of active metadata database cache descriptors, not the total number of databases.

3. During a peak period, find the number of active metadata descriptors:

```
sp_monitorconfig "open databases"
```

```
Usage information at date and time: Apr 22 2002  2:49PM.
Name                     Num_free    Num_active  Pct_act  Max_Used
Reuse_cnt   Instance_Name
------------------------   ----------  ----------  -------  -----
------
----------  --------------------
number of open                  50         20     .00         26
         0                  NULL
```

In this example, 20 metadata database descriptors are active; the maximum number of descriptors that have been active since the server was last started is 26.

See **sp_monitorconfig** in the *Reference Manual: Procedures* for more information.

4. Configure **number of open databases** to 26, plus additional space for 10 percent more (about 3), for a total of 29:

```
sp_configure "number of open databases", 29
```

If there is a lot of activity on the server, for example, if databases are being added or dropped, periodically run **sp_monitorconfig**. Reset the cache size as the number of active descriptors changes.

number of open indexes

number of open indexes sets the maximum number of indexes that can be used simultaneously on SAP ASE.

Summary Information	
Default value	500
Range of values	100–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Meta-Data Caches

If you are planning to make a substantial change, such as loading databases with a large number of indexes from another server, use **sp_helpconfig** to calculate an estimated metadata cache size. **sp_helpconfig** displays the amount of memory required for a given number of metadata descriptors, as well as the number of descriptors that can be accommodated by a given amount of memory. An index metadata descriptor represents the state of an index while it is in use or cached between uses.

Optimizing number of open indexes

If the default value of **number of open indexes** is insufficient, SAP ASE displays a message after trying to reuse active index descriptors, and you must adjust this value.

1. Use **sp_countmetadata** to find the total number of index metadata descriptors:

   ```
   sp_countmetadata "open indexes"
   ```

   The best time to run **sp_countmetadata** is when there is little activity in the server. Running **sp_countmetadata** during a peak time can cause contention with other processes.

   Suppose SAP ASE reports the following information:

   ```
   There are 698 user indexes in all database(s),
   requiring 286.289 Kbytes of memory. The 'open
   indexes' configuration parameter is currently set to
   500.
   ```
2. Configure the **number of open indexes** parameter to 698:

   ```
   sp_configure "number of open indexes", 698
   ```

   This new configuration is only a starting point; base the ideal size on the number of active index metadata cache descriptors, not the total number of indexes.
3. During a peak period, find the number of active index metadata descriptors:

   ```
   sp_monitorconfig "open indexes"
   ```

   ```
   Usage information at date and time: Apr 22 2002  2:49PM.
   Name                     Num_free    Num_active  Pct_act  Max_Used
   Reuse_cnt   Instance_Name
   ------------------------  ----------  ----------  -------  -----
   ------
   ----------  --------------------
   number of open                 182        516    73.92         590
            0                 NULL
   ```

   In this example, 590 is the maximum number of index descriptors that have been used since the server was last started.

   See **sp_monitorconfig** in the *Reference Manual: Procedures*.
4. Configure the **number of open indexes** configuration parameter to 590, plus additional space for 10 percent more (59), for a total of 649:

   ```
   sp_configure "number of open indexes", 649
   ```

   If there is a lot of activity on the server, for example, if tables are being added or dropped, periodically run **sp_monitorconfig**. Reset the cache size as the number of active descriptors changes.

number of open objects

number of open objects sets the maximum number of objects that can be open simultaneously on SAP ASE.

Summary Information	
Default value	500
Range of values	100–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Meta-Data Caches, SQL Server Administration

If you are planning to make a substantial change, such as loading databases with a large number of objects from another server, use **sp_helpconfig** to recalculate an estimated metadata cache size. **sp_helpconfig** displays the amount of memory required for a given number of metadata descriptors, as well as the number of descriptors that can be accommodated by a given amount of memory. An object metadata descriptor represents the state of an object while it is in use, or cached between uses.

Optimizing number of open objects

If the default **number of open objects** is insufficient, SAP ASE displays a message after trying to reuse active object descriptors.

1. Use **sp_countmetadata** to find the total number of object metadata cache descriptors:

   ```
   sp_countmetadata "open objects"
   ```

   The best time to run **sp_countmetadata** is when there is little activity in the server. Running **sp_countmetadata** during a peak time can cause contention with other processes.

   Suppose SAP ASE reports this information:

   ```
   There are 1340 user objects in all database(s),
   requiring 2894 Kbytes of memory. The 'open objects'
   configuration parameter is currently set to 500.
   ```
2. Configure **number of open objects** to account for the number of open objects:

   ```
   sp_configure "number of open objects", 1407
   ```

   1407 covers the 1340 user objects, plus 5 percent to accommodate temporary tables.

   This new configuration is only a starting point; base the ideal number on the active object metadata cache descriptors, not the total number of objects.
3. During a peak period, find the number of active metadata cache descriptors:

```
sp_monitorconfig "open objects"
```

```
Usage information at date and time: Aug 20 2007  1:32PM..
Name                    Num_free  Num_active  Pct_act  Max_Used  R
euse_cnt
   Instance_Name
---------------------  --------  ----------  -------  --------  -
--------
   --------------------
number of open
objects      560         847    71.40       1497         0
   NULL
```

In this example, 1497 is the maximum number of object descriptors that have been used since the server was last started.

4. Configure the **number of open objects** to 1497, plus 10 percent (150), for a total of 1647:

```
sp_configure "number of open objects", 1647
```

If there is a lot of activity on the server, for example, if tables are being added or dropped, periodically run **sp_monitorconfig**. Reset the cache size as the number of active descriptors changes. See **sp_monitorconfig** in the *Reference Manual: Procedures*.

number of open partitions

number of open partitions specifies the number of partitions that SAP ASE can access at one time.

Summary Information	
Default value	500
Range of values	100 – 2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Meta-Data Caches

Optimizing the number of open partitions parameter for your system

If the default value of **number of open partitions** is insufficient, SAP ASE displays a message after trying to reuse active partition descriptors. You must adjust this value.

This example assumes a system administrator has set the **number of open partitions** to 110:

1. Use **sp_countmetadata** to confirm the total number of open partitions:

```
sp_countmetadata "open partitions"
There are 42 user partitions in all database(s),
requiring 109 Kbytes of memory. The 'open
partitions' configuration parameter is currently set
to 110.
```

The best time to run **sp_countmetadata** is when there is little activity in the server. Running **sp_countmetadata** during a peak time can cause contention with other processes.

2. During a peak period, find the number of active metadata cache descriptors, for example:

```
sp_monitorconfig "open partitions"
```

```
Usage information at date and time: Jun 30 2008  3:15PM.

Name                     Num_free    Num_active  Pct_act  Max_Used
Reuse_cnt   Instance_Name
------------------------   ----------  ----------  -------  -----
------
----------  --------------------
number of open partitions    27            57    51.8         83
          0                  NULL
```

In this example, 83 is the maximum number of partition descriptors that have been used since the server was last started.

3. Configure the **number of open partitions** to 83, plus 10 percent (8), for a total of 91:

```
sp_configure "number of open partitions", 91
```

If there is a lot of activity on the server, for example, if tables are being added or dropped, periodically run **sp_monitorconfig**. Reset the cache size as the number of active descriptors changes. See **sp_monitorconfig** in the *Reference Manual: Procedures*.

number of pre-allocated extents

number of pre-allocated extents specifies the number of extents (eight pages) allocated in a single trip to the page manager.

Summary Information	
Default value	2
Range of values	1–32
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Currently, this parameter is used only by **bcp** to improve performance when copying in large amounts of data. By default, **bcp** allocates two extents at a time and writes an allocation record to the log each time.

Setting **number of pre-allocated extents** means that **bcp** allocates the specified number of extents each time it requires more space, and writes a single log record for the event.

An object may be allocated more pages than actually needed, so the value of **number of pre-allocated extents** should be low if you are using **bcp** for small batches. If you are using **bcp** for large batches, increase the value of **number of pre-allocated extents** to reduce the amount of overhead required to allocate pages and to reduce the number of log records.

Using a Value of 32 for the number of pre-allocated extents

Using a value of 32 for **number of pre-allocated extents** has a special significance for configuration and impacts the space allocations SAP ASE performs internally.

If you set **number of pre-allocated extents** to 32, SAP ASE reserves an entire allocation unit worth of extents for utility operations like **bcp-in** and **select into**, both of which use the large-scale allocation scheme of space reservation. This greatly improves the performance of these utilities, particularly when you run them concurrently on multiple nodes. Consequently, using a value of 32 guarantees that each node of a cluster is able to work independently on its own allocation unit without interference from the other nodes.

In earlier versions of SAP ASE, the **number of pre-allocated extents** parameter specified the number of extents reserved in a single allocation call for tables of all sizes.

With this version of SAP ASE, the value of **number of pre-allocated extents** is ignored for large tables with 240 or more pages for these commands only:

- **alter table** ***table_name*** **add** ***column_name*** . . .
- **alter table** ***table_name*** **modify** ***column_name*** . . .
- **alter table** ***table_name*** **drop** ***column_name*** . . .
- **alter table lock** . . .
- **reorg rebuild**

When you run these command on tables larger than 240 pages, SAP ASE reserves an entire allocation unit (32 extents), which greatly improves performance, particularly when you run them concurrently on multiple nodes.

The value of **number of pre-allocated extents** continues to be observed for the above commands for tables with fewer than 240 pages, and for all commands (such as **select into**, **bcp**, **alter table partition**) for tables of all sizes.

number of Q engines at startup

number of Q engines at startup, which specifies the number of Q engines that are online when the server starts, is required for MQ.

Summary Information	
Default value	0
Range of values	0 – 127
Status	Static
Display level	Comprehensive
Required role	System administrator

Summary Information	
Configuration group	Processors

You may need to increase **max online engines** to accommodate the number of **max online Q engines**.

number of remote connections

number of remote connections specifies the number of logical connections that can simultaneously be open to and from an SAP ASE.

Summary Information	
Default value	20
Range of values	5–32767
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, Network Communication

Each simultaneous connection to XP Server for ESP execution uses up to one remote connection each.

number of remote logins

number of remote logins controls the number of active user connections from SAP ASE to remote servers.

Summary Information	
Default value	20
Range of values	0–32767
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, Network Communication

Each simultaneous connection to XP Server for ESP execution uses up to one remote login each. Set this parameter to the same (or a lower) value as **number of remote connections**.

number of remote sites

number of remote sites determines the maximum number of remote sites that can simultaneously access SAP ASE.

Summary Information	
Default value	10
Range of values	0–32767
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, Network Communication

Each SAP ASE-to-XP Server connection uses one remote site connection.

Internally, **number of remote sites** determines the number of site handlers that can be active at any one time; all server accesses from a single site are managed with a single site handler. For example, if you set **number of remote sites** to 5, and each site initiates three remote procedure calls, **sp_who** shows 5 site handler processes for the 15 processes.

number of sort buffers

number of sort buffers specifies the amount of memory allocated for buffers used to hold pages read from input tables and perform index merges during sorts. **number of sort buffers** is used only for parallel sorting.

Summary Information	
Default value	500
Range of values	0–32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

SAP ASE allocates sort buffers from the cache to which the table is bound. If the cache is not bound to a table, SAP ASE allocates the sort buffers from the cache to which the database is bound. If the cache is not bound to a table or a database, SAP ASE allocates sort buffers from the default data cache.

Parallel sorts are used when you:

- Run **updates statistics**
- Create indexes

See *Performance and Tuning Series: Query Processing and Abstract Plans > Using Statistics to Improve Performance*.

The value you use for **number of sort buffers** depends on the page size of the server.

SAP recommends that you leave this parameter set to the default except when you are creating indexes in parallel.

Setting the value too high can rob nonsorting processes of access to the buffer pool in caches being used to perform sorts.

If you configure a high number of sort buffers, a sort on a large table may require more procedure cache. The effect is more pronounced with tables that have smaller row sizes, because the number of rows per page is higher.

This equation estimates the amount of procedure cache required (in bytes):

(*Number of sort buffers*) X (*rows per page*) X 100

If you do not configure enough procedure cache for the number of sort buffers, the sort may fail with error message 701. If this occurs, reconfigure SAP ASE with a lower number of sort buffers and retry the sort.

number of user connections

number of user connections sets the maximum number of user connections that can simultaneously be connected to SAP ASE. It does not refer to the maximum number of processes; that number depends not only on the value of this parameter but also on other system activity.

Summary Information	
Default value	25
Range of values	• UNIX and Linux: 5–32767 • Windows: 5–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, User Environment

Upper limit to the maximum number of user connections

The maximum allowable number of file descriptors per process is operating-system-dependent; see the configuration documentation for your platform.

The number of file descriptors available for SAP ASE connections is stored in the global variable *@@max_connections*. You can report the maximum number of file descriptors your system can use with:

```
select @@max_connections
```

The return value represents the maximum number of file descriptors allowed by the system for your processes, minus overhead. Overhead increases with the number of engines. For more information on how multiprocessing affects the number file descriptors available for SAP ASE connections, see *System Administration Guide: Volume 2 > Managing Mulitprocessor Servers*.

In addition, you must reserve a number of connections for the following items, which you also set with configuration parameters:

- The database devices, including mirror devices
- Site handlers
- Network listeners

The **number of user connections** + (**number of devices** * **max online engines** * 2) + **number of remote sites** + **max number network listeners** cannot be greater than the value of *@@max_connections*.

Reserved connections

One connection from the configured number of connections is reserved for temporary administrative tasks to make sure that database administrators can connect to SAP ASE. A reserved connection has a total login time of 15 minutes, and can be is allocated only to a user who has the **sa_role**. SAP ASE terminates the connection after 15 minutes to ensure the availability of the reserved connection at an installation with multiple database administrators.

SAP ASE also automatically uses this reserved connection when a client uses the last resource for connecting to SAP ASE.

If SAP ASE is using a reserved connection, the following informational message appears when the user logs in to SAP ASE:

```
There are not enough user connections available; you are being
connected
using a temporary administrative connection which will time out after
'15'
minutes. Increase the value of th 'number of user connections'
parameter
```

SAP ASE also prints a message similar to the following to the error log when the final connection to SAP ASE terminates due to a timeout:

```
00:00000:00008:2003/03/14 11:25:31.36 server  Process '16' has been
terminated as it exceeded the maximum login time allowed for such
processes.
This process used a connection reserved for system administrators and
has a
maximum login period of '15' minutes
```

Optimizing max number of user connections

There is no formula to determine how many connections to allow for each user. You must estimate this number, based on the system and user requirements. You must also take into account that on a system with many users, connections needed only occasionally or transiently can generally be shared among users. The following processes require user connections:

- One connection is needed for each user running **isql**.
- Application developers use one connection for each editing session.
- The number of connections required by users running an application depends on how the application has been programmed. Users executing Open Client programs need one connection for each open DB-Library **dbprocess** or Client-Library™ **cs_connection**.

Note: SAP suggests that you estimate the maximum number of connections used by SAP ASE and update **number of user connections** as you add physical devices or users to the system. Periodically use **sp_who** to determine the number of active user connections on your SAP ASE.

Certain other configuration parameters, including **stack size** and **default network packet size**, affect the amount of memory for each user connection.

User connections for shared memory—EJB Server

SAP ASE uses the value of **number of user connections** to establish the number of shared-memory connections for EJB Server. Thus, if **number of user connections** is 30, SAP ASE establishes 10 shared-memory connections for EJB Server. Shared-memory connections are not a subset of user connections, and are not subtracted from the number of user connections.

To increase the number of user connections for shared memory, you must:

1. Increase **number of user connections** to a number one-third of which is the number of desired shared-memory connections.
2. Restart SAP ASE.

Although **number of user connections** is a dynamic configuration parameter, you must restart the server to change the number of user connections for shared memory. See the *EJB Server Users Guide*.

With SAP ASE version 12.5.3 ESD #2, no sockets are automatically reserved for EJB. However, you can enable trace flag 1642 to revert to the functionality of earlier version, reserving one-third of the sockets for EJB. Enable traceflag 1642 to set up the EJB server. For this version of SAP ASE, you can ignore this message, `"hbc_ninit: No sockets available for HBC",` in the error log if the EJB server is not configured.

In SAP ASE version 12.5.3 and later, if the EJB server is enabled and HBC sockets are not available, `"hbc_ninit: No sockets available for HBC"` is reported. If traceflag 1642 is not enabled, set the flag, and restart SAP ASE. If the EJB server is not enabled, then no message is reported and SAP ASE automatically disables the sockets reserved for EJB server.

number of worker processes

number of worker processes specifies the maximum number of worker processes that SAP ASE can use at any one time for all simultaneously running parallel queries.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Query Tuning

SAP ASE issues a warning message at start-up if there is insufficient memory to create the specified number of worker processes. **memory per worker process** controls the memory allocated to each worker process.

If you have not configured **number of worker processes** for a sufficient number of threads from the worker thread pool, SAP ASE adjusts query plans at runtime to use fewer worker threads. If SAP ASE cannot adjust the queries at runtime, the queries recompile serially. However, **alter table** and **execute immediate** commands are aborted if they do not have sufficient worker threads.

o/s file descriptors

o/s file descriptors indicates the maximum per-process number of file descriptors configured for your operating system. This parameter is read-only and cannot be configured through SAP ASE.

Summary Information	
Default value	0
Range of values	Site-specific
Status	Read-only
Display level	Comprehensive
Required role	System administrator
Configuration group	O/S Resources

Many operating systems allow you to configure the number of file descriptors available per process. See your operating system documentation.

The number of file descriptors available for SAP ASE connections, which is less than the value of **o/s file descriptors**, is stored in the variable *@@max_connections*.

object lockwait timing

object lockwait timing controls whether SAP ASE collects timing statistics for requests of locks on objects.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

open index hash spinlock ratio

(Multiprocessing systems only) **open index hash spinlock ratio** sets the number of index metadata descriptor hash tables that are protected by one *spinlock*.

Summary Information	
Default value	100
Range of values	1–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, Meta-Data Cache

All the index descriptors belonging to a table are accessible through a hash table. When you run a query on the table, SAP ASE uses hash tables to look up the necessary index information in its `sysindexes` rows. A hash table is an internal mechanism used by SAP ASE to retrieve information quickly.

Usually, you do not need to change this parameter. In rare instances, however, you may need to reset it if SAP ASE demonstrates contention from hash spinlocks. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

For more information about configuring spinlock ratios, see *System Administration Guide: Volume 2 > Managing Multiprocessor Servers*.

open index spinlock ratio

open index spinlock ratio specifies the number of index metadata descriptors that are protected by one *spinlock*.

Summary Information	
Default value	100
Range of values	1–214748364
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Meta-Data Cache

SAP ASE uses a spinlock to protect an index descriptor, since more than one process can access the contents of the index descriptor. **open index spinlock ratio** is used only in multiprocessing systems.

The value specified for this parameter defines the ratio of index descriptors per spinlock.

If one spinlock is shared by too many index descriptors, it can cause spinlock contention. Use **sp_sysmon** to get a report on spinlock contention. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

If **sp_sysmon** output indicates an index descriptor spinlock contention of more than 3 percent, try decreasing the value of **open index spinlock ratio**.

See *System Administration Guide: Volume 2 > Managing Mulitprocessor Servers*.

open object spinlock ratio

(Multiprocessing systems only) **open object spinlock ratio** specifies the number of object descriptors that are protected by one *spinlock*.

Summary Information	
Default value	100
Range of values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Meta-Data Cache

SAP ASE uses a spinlock to protect an object descriptor, since more than one process can access the contents of the object descriptor.

The default value for this parameter is 100; 1 spinlock for each 100 object descriptors configured for your server. If your server is configured with only one engine, SAP ASE sets only 1 object descriptor spinlock, regardless of the number of object descriptors.

If one spinlock is shared by too many object descriptors, it causes spinlock contention. Use **sp_sysmon** to get a report on spinlock contention. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

If **sp_sysmon** output indicates an object descriptor spinlock contention of more than 3 percent, try decreasing the value of the **open object spinlock ratio** parameter.

See *System Administration Guide: Volume 2 > Managing Mulitprocessor Servers*.

optimization goal

optimization goal determines which optimization goal SAP ASE uses.

Summary Information	
Default value	**allrows_mix**
Range of values	**allrows_oltp**, **allrows_dss**
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

Optimization goals are a convenient way to match user query demands with the best optimization techniques, ensuring optimal use of the optimizer's time and resources. SAP ASE allows users to configure for two optimization goals, which you can specify at three tiers: server level, session level, and query level.

The server-level optimization goal is overridden at the session level, which is overridden at the query level.

These optimization goals allow you to choose an optimization strategy that best fits your query environment:

- **allrows_oltp** – the most useful goal for purely OLTP queries.
- **allrows_dss** – the most useful goal for operational DSS queries of medium-to-high complexity.

optimize temp table resolution

optimize temp table resolution allows stored procedures that reference temporary tables created outside the procedure to not require recompiling for each execution.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

Trace flag 299 previously supported this functionality.

optimization timeout limit

optimization timeout limit specifies the amount of time, as a fraction of the estimated execution time of the query, that SAP ASE can spend optimizing a query.

Summary Information	
Default value	10
Range of values	0 – 1000
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

A value of 0 indicates there is no optimization timeout.

optimizer level

optimizer level determines the level of optimization the query processor uses.

Summary Information	
Default value	**ase_default**

Summary Information	
Range of values	• **ase_current** – enables all optimizer changes through the current release. • **ase_default** – disables all optimizer changes since version 1503 ESD #1. • **ase1503esd2**– enables all optimizer changes through version 15.0.3 ESD #2. • **ase1503esd3** – enables all optimizer changes through version 15.0.3 ESD #3.
Status	Dynamic
Display level	
Required role	System administrator
Configuration group	Query Tuning

The values for **optimizer level** are character data, so you must use 0 as a placeholder for the second parameter, which must be numeric.

See *Performance and Tuning Series: Query Processing and Abstract Plans > Controlling Optimization* for information about optimization levels.

page lock promotion HWM

page lock promotion HWM (high-water mark), with **page lock promotion LWM** (low-water mark) and **page lock promotion PCT** (percentage), specifies the number of page locks permitted during a single scan session of a page-locked table or index before SAP ASE attempts to escalate from page locks to a table lock.

Summary Information	
Default value	200
Range of values	2–2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Lock Manager, SQL Server Administration

When the number of page locks acquired during a scan session exceeds **page lock promotion HWM**, SAP ASE attempts to acquire a table lock. **page lock promotion HWM** value cannot be higher than **number of locks**.

For more detailed information on scan sessions and setting up page lock promotion limits, see *Performance and Tuning Series: Locking and Concurrency Control > Locking Configuration and Tuning*.

The default value for **page lock promotion HWM** is appropriate for most applications. To avoid table locking, you may want to increase the value. For example, if you know that there are regular updates to 500 pages of an allpages-locked or datapages-locked table containing thousands of pages, increase concurrency for the tables by setting **page lock promotion HWM** to 500.

You can also configure lock promotion of page-locked tables and views at the object level. See **sp_setrowlockpromote** in the *Reference Manual: Procedures*.

Use **sp_sysmon** to see how changing **page lock promotion HWM** affects the number of lock promotions. **sp_sysmon** reports the ratio of exclusive page to exclusive table lock promotions and the ratio of shared page to shared table lock promotions. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

page lock promotion LWM

page lock promotion LWM (low-water mark), with **page lock promotion HWM** (high-water mark) and the **page lock promotion PCT**, specify the number of page locks permitted during a single scan session of a page locked table or an index before SAP ASE attempts to promote from page locks to a table lock.

Summary Information	
Default value	200
Range of values	2–value of **page lock promotion HWM**
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Lock Manager, SQL Server Administration

The **page lock promotion LWM** sets the number of page locks below which SAP ASE does not attempt to issue a table lock on an object. **page lock promotion LWM** must be less than or equal to **page lock promotion HWM**.

The default value for **page lock promotion LWM** is sufficient for most applications. If SAP ASE runs out of locks (except for an isolated incident), increase **number of locks**.

See the *Performance and Tuning Series: Locking and Concurrency Control*.

You can also configure page lock promotion at the object level. See **sp_setpglockpromote** in the *Reference Manual: Procedures*.

page lock promotion PCT

page lock promotion PCT sets the percentage of page locks (based on the table size) above which SAP ASE attempts to acquire a table lock. If the number of locks held on an object is between **page lock promotion LWM** (low-water mark) and **page lock promotion HWM** (high-water mark).

Summary Information	
Default value	100
Range of values	1–100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Lock Manager, SQL Server Administration

See *Performance and Tuning Series: Locking and Concurrency Control > Locking Configuration and Tuning*.

The default value for **page lock promotion PCT** is appropriate for most applications.

You can also configure lock promotion at the object level for page locked objects. See **sp_setpglockpromote** in the *Reference Manual: Procedures*.

page utilization percent

page utilization percent is used during page allocations to control whether SAP ASE scans a table's object allocation map (OAM) to find unused pages or simply allocates a new extent to the table.

Summary Information	
Default value	95
Range of values	1–100
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Disk I/O

The **page utilization percent** parameter is a performance optimization for servers with very large tables; it reduces the time needed to add new space.

If you set **page utilization percent** to 100, SAP ASE scans through all OAM pages to find unused pages allocated to the object before allocating a new extent. When this parameter is set lower than 100, SAP ASE compares the **page utilization percent** setting to the ratio of used and unused pages allocated to the table, as follows:

```
100 * used pages/(used pages + unused pages)
```

If **page utilization percent** is lower than the ratio, SAP ASE allocates a new extent instead of searching for the unused pages.

For example, when inserting data into a 10GB table that has 120 OAM pages and only 1 unused data page:

- A **page utilization percent** of 100 tells SAP ASE to scan through all 120 OAM pages to locate an unused data page.
- A **page utilization percent** of 95 allows SAP ASE to allocate a new extent to the object, because 95 is lower than the ratio of used pages to used and unused pages.

A low **page utilization percent** value results in more unused pages. A high **page utilization percent** value slows page allocations in very large tables, as SAP ASE performs an OAM scan to locate each unused page before allocating a new extent. This increases logical and physical I/O.

If page allocations (especially in the case of large inserts) seem to be slow, lower the value of **page utilization percent**, but reset it after inserting the data. A lower setting affects all tables on the server and results in unused pages in all tables.

Fast bulk copy ignores the **page utilization percent** setting and always allocates new extents until there are no more extents available in the database.

partition groups

partition groups specifies the maximum number of partition groups that can be allocated by SAP ASE.

Summary Information	
Default value	1024
Range of values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Meta-Data Cache

Partition groups are internal structures used by SAP ASE to control access to individual partitions of a table. Partition groups are used while upgrading or during a **load database** upgrade to unpartition SAP ASE 12.5.x and earlier partitions.

The default value allows a maximum 1024 open partition groups and a maximum of 2147483647 open partitions. The actual number of partitions may be slightly less, due to the grouping of partitions.

partition spinlock ratio

For SAP ASEs running with multiple engines, **partition spinlock ratio** sets the number of rows in the partition descriptors that are protected by one *spinlock*.

Summary Information	
Default value	10
Range of values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Meta-Data Cache

SAP ASE manages access to table partitions using partition descriptors. Each partition descriptor stores information about a partition (for example, the last page of the partition) that processes must use when accessing that partition. Configure partition descriptors using **number of open partitions**.

The default value of **partition spinlock ratio** sets 1 spinlock for every 10 partition caches. Decreasing the value of **partition spinlock ratio** may have little impact on the performance of SAP ASE. The default setting is correct for most servers.

See *System Administration Guide: Volume 2 > Managing Mulitprocessor Servers*.

pci memory size

pci memory size sets the size of the pluggable component interface (PCI) memory pool.

Summary Information	
Default value	64MB
Valid values	0 – 2147483647
Status	Dynamic
Display level	Intermediate
Required role	System Administrator
Configuration group	User Environment

All pluggable component adapter (PCA) and JVM plug-ins running under the PCI Bridge share a single dedicated PCI memory pool. If you set **pci memory size** to less than the default, SAP ASE uses the default size.

This memory pool is fully dedicated to the PCI bridge and any running pluggable component. Like all other memory pools, SAP ASE controls this memory pool. However, unlike other memory pools, the PCI memory pool is allocated when you initialize the PCI Bridge and does not grow after that time.

per object statistics active

per object statistic active controls whether SAP ASE collects statistics for each object.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

percent database for history

percent database for history specifies the percentage of the total space available in `sybmgmtdb` that is reserved for the `js_history` table.

Summary Information	
Default value	20
Valid values	0 – 100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

Increase **percent database for history** if there are more jobs running, or to store historical records about executed jobs for future queries.

percent database for output

percent database for output specifies the percentage of the total space available in `sybmgmtdb` that is reserved for job output.

Summary Information	
Default value	30
Valid values	0 – 100

Summary Information	
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

Increase the default value if there are more jobs running or jobs that produce lot of output that must be stored for querying.

percent history free

percent history free specifies the percentage of reserved space in `sybmgmtdb` to be kept free.

Summary Information	
Default value	30
Valid values	0 – 100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

For example, if you use the default value, SAP ASE starts purging the oldest history records to make room for new records when 70 percent of `sybmgmtdb` is filled.

percent output free

Specifies the percentage of reserved space kept free in `sybmgmtdb` that is reserved for Job Scheduler output.

Summary Information	
Default value	50
Valid values	0 – 100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

For example, if you use the default value, SAP ASE starts purging the oldest history records to make room for new records when 50 percent of `sybmgmtdb` is filled.

performance monitoring option

performance monitoring option enables the license for the BMC DBXray graphical performance monitoring and diagnostic tool .

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

permission cache entries

permission cache entries determines the number of cache protectors per task, increasing the amount of memory for each user connection and worker process.

Summary Information	
Default value	64
Range of values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, User Environment

Information about user permissions is held in the permission cache. When SAP ASE checks permissions, it looks first in the permission cache; if it does not find what it needs, it looks in the `sysprotects` table. This process is significantly faster if SAP ASE finds the information it needs in the permission cache and does not have to read `sysprotects`.

However, SAP ASE looks in the permission cache only when it is checking user permissions, not when permissions are being granted or revoked. When a permission is granted or revoked, the entire permission cache is flushed. This is because existing permissions have timestamps that become outdated when new permissions are granted or revoked.

If users on your SAP ASE frequently perform operations that require their permissions to be checked, you may see a small performance gain by increasing the value of **permission cache entries**. This effect is not likely to be significant enough to warrant extensive tuning.

If users on your SAP ASE frequently grant or revoke permissions, avoid setting **permission cache entries** to a large value. The space used for the permission cache would be wasted, since the cache is flushed with each **grant** and **revoke** command.

plan text pipe active

plan text pipe active determines whether SAP ASE collects query plan text.

Summary Information	
Default value	0
Range of values	0–1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

If both **plan text pipe active** and **plan text pipe max messages** are enabled, SAP ASE collects the plan text for each query. You can use `monSysPlanText` to retrieve the query plan text for all user tasks.

plan text pipe max messages

plan text pipe max messages determines the number of query plan text messages SAP ASE stores per engine.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Monitoring

The total number of messages in the `monSQLText` table is the value of **sql text pipe max messages** multiplied by the number of engines running.

print deadlock information

print deadlock information prints deadlock information to the error log.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on), 2 (on, print summary)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Lock Manager, SQL Server Administration

If you are experiencing recurring deadlocks, setting **print deadlock information** to 1 provides you with detailed information in the error log that can be useful in tracing the cause of the deadlocks. However, setting **print deadlock information** to 1 can degrade SAP ASE performance. For this reason, set **print deadlock information** on only when you are trying to determine the cause of deadlocks.

Use **sp_sysmon** output to determine whether deadlocks are occurring in your application. If they are, set **print deadlock information** to 1 to learn more about why they are occurring. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

A value of 2 allows you to print a summary of deadlock information to the error log (as opposed to the detailed information a value of 1 provides). For example:

```
Deadlock Id 34: Process (Familyid 0, Spid 70) was waiting for a
'exclusive page'
lock on page 10858346 of the 'equineline_job' table in database 18
but process
(Familyid 0, Spid 88) already held a 'exclusive page' lock on it.
Deadlock Id 34: Process (Familyid 0, Spid 88) was waiting for a
'exclusive page'
lock on page 11540986 of the 'equineline_job' table in database 18
but process
(Familyid 0, Spid 70) already held a 'update page' lock on it.
```

print recovery information

print recovery information determines what information SAP ASE displays on the console during recovery. (Recovery is performed on each database at SAP ASE start-up and when a database dump is loaded.)

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)

Summary Information	
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Backup/Recovery

The default value means that SAP ASE displays only the database name and a message saying that recovery is in progress. A value of 1 indicates that SAP ASE displays information about each individual transaction processed during recovery, including whether it was aborted or committed.

procedure cache size

procedure cache size specifies the size of the procedure cache, in 2K pages.

Summary Information	
Default value	7000
Range of values	7000 – 2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Memory Use, SQL Server Administration

SAP ASE uses the procedure cache while running stored procedures. If the server finds a copy of a procedure already in the cache, it does not need to read it from the disk. SAP ASE also uses space in the procedure cache to compile queries while creating stored procedures.

Since the optimum value for **procedure cache size** differs from application to application, resetting it may improve SAP ASE performance. For example, if you run many different procedures or ad hoc queries, your application uses the procedure cache more heavily, so you may want to increase this value.

Warning! If **procedure cache size** is too small, SAP ASE performance degrades.

If you are upgrading, **procedure cache size** is set to the size of the original procedure cache at the time of upgrade.

procedure deferred compilation

procedure deferred compilation enables or disables compiling statements that reference local variables or temporary tables inside a stored procedure until execution time, so that the optimization of those statements can use runtime values, instead of estimations.

Summary Information	
Default value	1(enabled)
Range of values	0 – 1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Query tuning

process wait events

process wait events controls whether SAP ASE collect statistics for each wait event for every task.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Monitoring

You can get wait information for a specific task using **monProcessWaits**.

See *Transact-SQL Users Guide > Using Stored Procedures*.

prod-consumer overlap factor

prod-consumer overlap factor affects optimization. SAP ASE changes the group by algorithm, and you cannot use set statistics I/O with parallel plans.

Summary Information	
Default value	20
Range of values	
Status	dynamic

Summary Information	
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

quorum heartbeat interval

quorum heartbeat interval specifies the number of seconds between quorum heartbeats.

Summary Information	
Default value	5
Valid values	1 – 60
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

Setting **quorum heartbeat interval** to a lower number increases the heartbeat overhead but speeds the detection of a lost disk link, resulting in a quicker termination of an instance for which you have set I/O fencing or that has lost its SAN link. Setting **quorum heartbeat interval** to a high number reduces heartbeat overhead, but delays the detection of a lost disk link.

quorum heartbeat retries

quorum heartbeat retries specifies the number of times an instance attempts to detect a quorum heartbeat before determining that the quorum device is no longer running, and exiting.

Summary Information	
Default value	2
Valid values	0 – 32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

A value of 0 indicates that the instance should terminate upon the first quorum heartbeat failure.

Tuning this to a lower number causes an instance to fail over more quickly when access to the quorum device is lost, potentially improving application recovery times. Tuning this to a higher number degrades application recovery, reducing the chances that a transient disk access problem causes an instance failure.

quoted identifier enhancements

quoted identifier enhancements enables and disables quoted identifiers use in SAP ASE.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

When enabled, SAP ASE can use quoted identifiers for:

- Tables
- Views
- Column names
- Index names
- System procedure parameters

Note: The default value for **quoted identifier enhancements** depends on the value to which **enable functionality group** is set. If you set **enable functionality group** to:

- 0 – the default value for **quoted identifier enhancements** is 0.
- 1 – the default value for **quoted identifier enhancements** is 1.

However, if you set **quoted identifier enhancements** to 1, it uses a value of 1 regardless of what you set **enable functionality group** to.

See **enable functionality group**.

See the *Reference Manual: Blocks*.

rapidlog buffer size

Specifies the buffer size for the output of diagnostic for measuring 'Proc Cache Header' performance.

Summary Information	
Default value	1024

Summary Information	
Valid values	1024–2000000
Status	Dynamic
Display level	
Required role	System administrator
Configuration groups	Error Log

See *Performance and Turning Series: Basics > Memory Use Performance > Procedure Cache > Diagnostic for 'Proc Cache Header' Memory Pool*.

rapidlog max files

Specifies the maximum number of fies for the output of diagnostic for measuring 'Proc Cache Header' performance.

Summary Information	
Default value	99999999
Valid values	1–99999999
Status	Dynamic
Display level	
Required role	System administrator
Configuration groups	Error Log

See *Performance and Turning Series: Basics > Memory Use Performance > Procedure Cache > Diagnostic for 'Proc Cache Header' Memory Pool*.

read committed with lock

read committed with lock determines whether an SAP ASE using transaction isolation level 1 (read committed) holds shared locks on rows or pages of data-only-locked tables during **select** queries.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary Information	
Configuration group	Lock Manager

For cursors, **read committed with lock** applies only to read-only cursors declared.

For transaction isolation level 1, **select** queries on allpages-locked tables continue to hold locks on the page at the current position. Any updatable cursor on a data-only-locked table also holds locks on the current page or row. See the *Performance and Tuning Series: Basics*.

recovery interval in minutes

recovery interval in minutes sets the maximum number of minutes per database that SAP ASE uses to complete its recovery procedures in case of a system failure.

Summary Information	
Default value	5
Range of values	1–32767
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Backup/Recovery

The recovery procedure rolls transactions backward or forward, starting from the transaction that the checkpoint process indicates as the oldest active transaction. The recovery process has more or less work to do, depending on the value of **recovery interval in minutes**.

SAP ASE estimates that 6000 rows in the transaction log require 1 minute of recovery time. However, different types of log records can take more or less time to recover. If you set **recovery interval in minutes** to 3, the checkpoint process writes changed pages to disk only when `syslogs` contains more than 18,000 rows since the last checkpoint.

Note: The recovery interval has no effect on long-running, minimally logged transactions (such as **create index**) that are active when SAP ASE fails. It may take as much time to reverse these transactions as it took to run them. To avoid lengthy delays, dump each database after index maintenance operations.

SAP ASE uses the **recovery interval in minutes** setting and the amount of activity on each database to decide when to checkpoint each database. When SAP ASE checkpoints a database, it writes all *dirty pages* (data pages in cache that have been modified) to disk. This may create a brief period of high I/O, called a checkpoint spike. The checkpoint also performs other maintenance tasks, including truncating the transaction log for each database for which the **truncate log on chkpt** option has been set. About once per minute, the sleeping checkpoint process “wakes up,” checks the **truncate log on chkpt** setting, and checks the recovery interval to determine if a checkpoint is needed. The figure below shows the logic used by SAP ASE during this process.

Figure 5: The checkpoint process

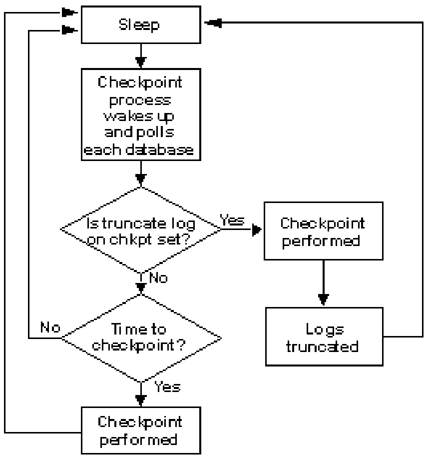

You may want to change the recovery interval if your application and its use change. For example, you may want to shorten the recovery interval when there is an increase in update activity on SAP ASE. Shortening the recovery interval causes more frequent checkpoints, with smaller, more frequent checkpoint spikes, and slows the system slightly. However, setting the recovery interval too high may cause the recovery time to be unacceptably long. You can reduce the spikes caused by checkpointing by reconfiguring the **housekeeper freewrite percent** parameter. For more information on the performance implications of **recovery interval in minutes**, see *Performance and Tuning Series: Basics > Memory Use and Performance*.

Use **sp_sysmon** to determine how a particular recovery interval affects the system. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

recovery prefetch size

recovery prefetch size sets the look-ahead size (in numbers of log records) to be used by the recovery prefetch scan.

Summary Information	
Default value	0 (use dynamic prefetch)
Range of values	0–20,000
Status	Dynamic
Display level	Comprehensive

Summary Information	
Required role	System administrator
Configuration group	SQL Server Administration

Set to 0 if the scan is to determine the look-ahead size dynamically, or to a value > 0 if the look-ahead size is to be set to a specific number of log records to look-ahead.

remote server pre-read packets

remote server pre-read packets determines the number of packets that are pre-read by a site handler during connections with remote servers.

Summary Information	
Default value	3
Range of values	3–255
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, Network Communication

To reduce the required number of connections, communication between two servers is managed through a single site handler. The site handler can pre-read and keep track of data packets for each user process before the receiving process is ready to accept them.

The default value for **remote server pre-read packets** is appropriate for most servers. Increasing the value uses more memory; decreasing the value can slow network traffic between servers.

restricted decrypt permission

restricted decrypt permission enables or disables restricted decrypt permission in all databases. You must have the **sso_role** to set this parameter.

Summary Information	
Default value	0
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System security officer
Configuration group	Security Related

When **restricted decrypt permission** is set to 0 (off), decrypt permission on encrypted columns acts the same as in versions earlier than 15.0.2:

- The table owner or the SSO explicitly grants decrypt permission. However, **with grant** option on decrypt permission is supported.
- Decrypt permission is granted implicitly to table owners and the SSO, as well as to any user through a chain of ownership. For example, if user Fred owns the `proc1` stored procedure, which selects data from the encrypted column `fred.table1.col1`, and if Fred grants **exec** permission on `proc1` to Harry, then Harry has implicit decrypt permission on `fred.table1.col1`
- Decrypt permission is not needed for **alter table decrypt**. because the table owner has implicit decrypt permission on encrypted columns.

When **restricted decrypt permission** is set to 1 (on):

- Decrypt permission is granted implicitly only to the SSO.
- The SSO can grant decrypt permission using the **with grant option** parameter. This allows the SSO to decide who can grant decrypt permission in the system. For example, if the SSO wants user1 to be able to grant **decrypt permission** on `user3.user3_tab`, the SSO issues:

  ```
  grant decrypt on user3.user3_tab to user1
  with grant option
  ```

  If you use a system encryption password, SAP recommends that, to protect data privacy, you do not grant decrypt permission to the DBO to. Access to keys through user passwords prevents the DBO and other parties from accessing the data unless they have a key's password; however, you may find it convenient for the DBO to decide which users should see the decrypted data. If you are not protecting keys and data with user-specified passwords, the SSO should retain the sole responsibility to grant decrypt permission.
- Table ownership does not give a user implicit decrypt permission. That is, if you create a table with encrypted columns, you do not have decrypt permission on them unless it is explicitly granted to you.
- No user is implicitly granted decrypt permission through an ownership chain. For example, if Fred owns the **proc1** stored procedure, which selects data from the encrypted column `fred.table1.col1`, and if Fred grants **exec** permission on **proc1** to Harry, then Harry must also have explicit decrypt permission on `fred.table1.col1` to see the data.
- Aliased users assume the permissions of the user to whom they are aliased. Similarly, a user with **sa_role**, who is implicitly aliased to the DBO in any database, inherits any decrypt permissions that have been explicitly granted to the DBO.
- Decrypt permission is required for **alter table decrypt** statement because the table owner does not have implicit decrypt permission on the table.

If you change **restricted decrypt permission** from 0 to 1, currently executing statements that use implicit decrypt permission finish; however any subsequent statements that use implicit decrypt permission fail with this error until the SSO grants the user decrypt permission on the necessary columns:

```
Msg 10330 "DECRYPT permission denied on object object_name, database
database_name, owner owner_name."
```

If you change **restricted decrypt permission** from 1 to 0, the rows that reflect explicit grants remain in the `sysprotects` system table. However, these rows have no effect on implicitly granted decrypt permissions because SAP ASE does not check `sysprotects` to make sure decrypt permission can be implicitly granted. **sp_helprotect** displays misleading information for only those users who were granted or revoked explicit decrypt permission before you reconfigure the system, and who now have implicit decrypt permission.

SAP recommends that, to keep the system consistent, you revoke any explicit decrypt permissions granted to users before you switch between enabling or disabling **restricted decrypt permission** to keep the system consistent.

See the *Encryption Users Guide* for more information about decrypt permissions.

row lock promotion HWM

row lock promotion HWM (high-water mark), with **row lock promotion LWM** (low-water mark) and **row lock promotion PCT** specifies the maximum number of row locks permitted during a single scan session of a table or an index before SAP ASE attempts to escalate from row locks to a table lock.

Summary Information	
Default value	200
Range of values	2–2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Lock Manager, SQL Server Administration

When the number of locks acquired during a scan session exceeds **row lock promotion HWM**, SAP ASE attempts to acquire a table lock. The **lock promotion HWM** value cannot be higher than the **number of locks** value.

See *Performance and Tuning Series: Locking and Concurrency Control > Locking Configuration and Tuning*.

The default value for **row lock promotion HWM** is appropriate for most applications. To avoid table locking, you may want to increase the value of row lock promotion HWM.. For example, if you know that there are regular updates to 500 rows on a table that has thousands of rows, you can increase concurrency for the tables by setting **row lock promotion HWM** to around 500.

You can also configure row lock promotion at the object level. See **sp_setpglockpromote** in the *Reference Manual: Procedures*.

row lock promotion LWM

row lock promotion LWM (low-water mark), with the **row lock promotion HWM** (high-water mark) and **row lock promotion PCT** specifies the number of row locks permitted during a single scan session of a table or an index before SAP ASE attempts to promote from row locks to a table lock.

Summary Information	
Default value	200
Range of values	2–value of **row lock promotion HWM**
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Lock Manager, SQL Server Administration

row lock promotion LWM sets the number of locks below which SAP ASE does not attempt to acquire a table lock on the object. The **row lock promotion LWM** must be less than or equal to **row lock promotion HWM**.

The default value for **row lock promotion LWM** is sufficient for most applications. If SAP ASE runs out of locks (except for an isolated incident), increase **number of locks**.

See the *Performance and Tuning Series: Locking and Concurrency Control*.

You can also configure lock promotion at the object level. See **sp_setpglockpromote** in the *Reference Manual: Procedures*.

row lock promotion PCT

If the number of locks held on an object is between **row lock promotion LWM** (low-water mark) and **row lock promotion HWM** (high-water mark), **row lock promotion PCT** sets the percentage of row locks (based on the number of rows in the table) above which SAP ASE attempts to acquire a table lock.

Summary Information	
Default value	100
Range of values	1–100
Status	Dynamic
Display level	Intermediate
Required role	System administrator

Summary Information	
Configuration groups	Lock Manager, SQL Server Administration

The default value for **row lock promotion PCT** is appropriate for most applications.

For more information on setting up lock promotion limits, see *Performance and Tuning Series: Locking and Concurrency Control > Locking Configuration and Tuning*.

You can also configure row lock promotion at the per-object level. See **sp_sterowlockpromote** in the *Reference Manual: Procedures*.

rtm thread idle wait period

rtm thread idle wait period defines the time, in seconds, a native thread used by SAP ASE waits when it has no work to do. When the time set for a native thread is reached, the thread automatically fades out.

Summary Information	
Default value	600
Range of values	600 – 4026531839
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	SQL Server Administration

runnable process search count

runnable process search count specifies the number of times an engine loops while looking for a runnable task before relinquishing the CPU to the operating system.

Summary Information	
Default value	2000 (default value of 3 for the Cluster Edition)
Range of values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Note: runnable process search count functions only when you configure SAP ASE for process kernel mode; it is nonfunctional for threaded kernel mode. Use **alter thread pool *pool_name* with idle timeout = *n*** instead.

SAP ASE engines check the run queue for runnable tasks whenever a task completes or exceeds its allotted time on the engine. At times, there are no tasks in the run queues. An engine can either relinquish the CPU to the operating system or continue to check for a task to run. Setting **runnable process search count** higher causes the engine to loop more times, thus holding the CPU for a longer time. Setting the **runnable process search count** lower causes the engine to release the CPU sooner.

If your machine is a uniprocessor that depends on helper threads to perform I/O, you may see some performance benefit from setting **runnable process search** to perform network I/O, disk I/O, or other operating system tasks. If a client, such as a bulk-copy operation, is running on the same machine as a single CPU server that uses helper threads, you may need to allow both the server and the client access to the CPU.

Note: If you are having performance problems, try setting **runnable process search count** to 3.

For SAP ASEs running on uniprocessor machines that do not use helper threads, and for multiprocessor machines, the default value should provide good performance.

With a **runnable process search count** value of 3, the Cluster Edition can better share the system CPU with other processes running on the same machine. However, if you set **runnable process search count** to 3 and SAP ASE is running as a standalone process, users may experience delays in server response times. In this case, reset **runnable process search count** to 2000.

Use **sp_sysmon** to determine how the **runnable process search count** parameter affects the SAP ASE use of CPU cycles, engine yields to the operating system, and blocking network checks. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

runnable process search count versus idle timeout

runnable process search count and **alter thread pool ...idle timeout** both indicate how SAP ASE looks for work.

- **runnable process search count** specifies the number of loops SAP ASE spends looking for work, and is a server-wide parameter.
- **alter thread pool ...idle timeout** specifies the period of time SAP ASE spends looking for work, and is tuned according to individual thread pools. **idle timeout** is more consistent across processors with varying speeds.

sampling percent

sampling percent is the numeric value of the sampling percentage, such as 5 for 5%, 10 for 10%, and so on.

Summary Information	
Default value	0

Summary Information	
Range of values	0 – 100
Status	Dynamic
Display level	Comprehensive
Required role	System or database administrator
Configuration group	Query Tuning

To reduce I/O contention and resources, run **update statistics** using a sampling method, which can reduce the I/O and time when your maintenance window is small and the data set is large. If you are updating a large data set or table that is in constant use, being truncated and repopulated, you may want to perform a statistical sampling to reduce the time and the size of the I/O.

Remember that with sampling, results are not fully accurate. Balance changes to histogram values against the savings in I/O.

Although a sampling of the data set may not be completely accurate, usually the histograms and density values are reasonable within an acceptable range.

When you are deciding whether or not to use sampling, consider the size of the data set, the time constraints you are working with, and if the histogram produced is as accurate as needed.

The percentage to use when sampling depends on your needs. Test various percentages until you receive a result that reflects the most accurate information on a particular data set.

Statistics are stored in the system tables `systabstats` and `sysstatistics`.

secure default login

secure default login specifies a default login for all users who are preauthenticated but who do not have a login in `master..syslogins`.

Summary Information	
Default value	0
Range of values	0 (followed by another parameter naming the default login)
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

Establish the secure default login with:

```
sp_configure "secure default login", 0, default_login_name
```

where:

- **secure default login** – is the name of the parameter.
- **0** – is a required parameter because the second parameter of **sp_configure** must be a numeric value.
- *default_login_name* – is the name of the default login for a user who is unknown to SAP ASE, but who has already been authenticated by a security mechanism. The login name must be a valid login in `master..syslogins`.

select for update

select for update enables SAP ASE to exclusively lock rows for subsequent updates within the same transaction, and for updatable cursors, preventing other concurrent tasks from updating these rows and from blocking the subsequent update.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

See *Transact-SQL Users Guide > Queries: Selecting Data from a Table*.

Note: The default value for **select for update** depends on the value to which **enable functionality group** is set. If you set **enable functionality group** to:

- 0 – the default value for **select for update** is 0.
- 1 – the default value for **select for update** is 1.

However, if you set **select for update** to 1, it uses a value of 1 regardless of what you set **enable functionality group** to.

See **enable functionality group**.

select on syscomments.text

select on syscomments.text enables protection of the text of database objects through restriction of the **select** permission on the `text` column of the `syscomments` table.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)

Summary Information	
Status	Dynamic
Display level	Comprehensive
Required role	System security officer
Configuration group	Security Related

The default value sets **select** permission to “public.” Set the value to 0 to restrict **select** permission to the object owner and the system administrator.

send doneinproc tokens

send doneinproc tokens enables or disables SAP ASE for sending `doneinproc` packets (these are TDS messages that are sent after various statements, in particular, non-**select** statements like **insert**, **update**, and so on).

Summary Information	
Default value	1 (on)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Network Communication

When disabled (set to 0), SAP ASE does not send `doneinproc` tokens for non-**select** statements.

send doneinproc tokens replaces **dbcc tune ‘doneinproc’** and trace flag 292.

Setting **send doneinproc tokens** to 1 is safe in most cases. However, some stored procedures are executed using asynchronous commands from CT-Lib, and using a value of 0 may cause state-machine errors in some CT-Lib applications.

session migration timeout

session migration timeout specifies the amount of time available for a client to complete a migration by connecting to the target instance. If the client does not migrate to the target instance in the time alloted, SAP ASE fails the connection.

Summary Information	
Default value	600
Valid values	0 – 32767

Summary Information	
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

session tempdb log cache size

session tempdb log cache size specifies the size for each session `tempdb` log cache.

Summary Information	
Default value	The logical page size, in bytes
Range of values	The logical page size up to 2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

SAP ASE uses the `tempdb` log cache to buffer the user transaction log records for the session involving `tempdb` objects. The size for **session tempdb log cache** is determined by the SAP ASE page size. Configuring a reasonable size for **session tempdb log cache size** can help prevent context switches due to a session having to flush the **tempdb** user log cache.

SAP ASE includes two log caches for each user connection: the session `tempdb` log cache, and the user log cache, which is determined by the **user log cache size** parameter.

shared memory starting address

shared memory starting address determines the virtual address where SAP ASE starts its shared memory region. It is unlikely that you will ever reconfigure **shared memory starting address**. (Do so only after consulting with SAP Technical Support.)

Summary Information	
Default value	0
Range of values	Platform-specific
Status	Static
Display level	Comprehensive
Required role	System administrator

Summary Information	
Configuration group	Physical Memory

size of auto identity column

size of auto identity column sets the precision of IDENTITY columns that are automatically created with the **sp_dboption auto identity** and **unique auto_identity index** options.

Summary Information	
Default value	10
Range of values	1–38
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

The maximum value that can be inserted into an IDENTITY column is $10^{precision}$ -1. After an IDENTITY column reaches its maximum value, all further **insert** statements return an error that aborts the current transaction.

If you reach the maximum value of an IDENTITY column, you can increase it with a modify operation in the **alter table** command. See the *Transact-SQL Users Guide* for examples.

You can also use the **create table** command to create a table that is identical to the old one, but with a larger precision for the IDENTITY column. After you have created the new table, use the **insert** command or **bcp** to copy data from the old table to the new one.

size of global fixed heap

size of global fixed heap specifies the memory space for internal data structures and other needs.

Summary Information	
Default values	150 pages (32-bit version) 300 pages (64-bit version)
Minimum values	10 pages (32-bit version) 20 pages (64-bit version)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary Information	
Configuration groups	Java Services, Memory Use

If you change **size of the global fixed heap**, change the total logical memory by the same amount.

size of process object heap

size of process object fixed heap specifies the total memory space for all processes using the Java VM.

Summary Information	
Default values	1500 pages (32-bit version) 3000 pages (64-bit version)
Minimum values	45 pages (32-bit version) 90 pages (64-bit version)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration groups	Java Services, Memory Use

If you change **size of process object fixed heap**, change the **total logical memory** by the same amount.

size of shared class heap

size of shared class heap specifies the shared memory space for all Java classes that are called into the Java VM. SAP ASE maintains the shared class heap server-wide for both user-defined and system-provided Java classes.

Summary Information	
Default values	1536 pages (32-bit version) 3072 pages (64-bit version)
Minimum values	650 pages (32-bit version) 1300 pages (64-bit version)
Status	Dynamic
Display level	Basic
Required role	System administrator

Summary Information	
Configuration groups	Java Services, Memory Use

If you change the **size of shared class heap**, change the **total logical memory** by the same amount.

size of unilib cache

size of unilib cache specifies the memory used in bytes rounded up to the nearest 1K in addition to the minimum overhead size, which provides enough memory to load a single copy of the largest Unilib conversion table plus the largest Unilib sort table.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Unicode

Asian clients may need to increase **size of unilib cache** by an extra 100K for every additional character set they want to support via Unicode-based conversion.

solaris async i/o mode

solaris async i/o mode allows you to select various asynchronous IO modes on the Solaris platform.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	
Required role	System administrator
Configuration group	Disk I/O

Allows you to select various asynchronous IO modes on the Solaris platform. This parameter is effective if SAP ASE is running in threaded kernel mode. This parameter is static, therefore is effective after restarting the server.

0 – (Default) Use this mode if the Solaris patch containing the fix for Oracle BugID 16054425 is not installed. You may see sub-optimal IO performance.

1 – (Recommended) You must have the Solaris patch containing the fix for Oracle BugID 16054425 installed.

Install the following Oracle patch for your platform:

- For Solaris 10 SPARC: 148888-03
- For Solaris 10 x86/x64: 148889-03
- For Solaris 11, latest SRU containing fix for Oracle Bug 16054425

Note: If **solaris async i/o mode** is set to 1 without the patch for Oracle BugID 16054425, Adaptive Server may report 694 or 823 errors and require restarting the server. Oracle Bug 15868517 refers to backport of Oracle Bug 16054425 for S10U11.

sproc optimize timeout limit

sproc optimize timeout limit specifies the amount of time SAP ASE can spend optimizing a stored procedure as a fraction of the estimated execution time.

Summary Information	
Default value	40
Range of values	0 – 4000
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

SQL batch capture

SQL batch capture controls whether SAP ASE collects SQL text. If both **SQL batch capture** and **max SQL text monitored** are enabled, SAP ASE collects the SQL text for each batch for each user task.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

SQL Perfmon Integration

(Windows only)**SQL Perfmon Integration** enables and disables the ability to monitor SAP ASE statistics from the Windows Performance Monitor.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

SAP ASE must be registered as a Windows Service to support monitor integration. This occurs automatically when you start SAP ASE using the Services Manager in the SAP for the Windows program group, and is the default configuration when you use the SAP installer or the **syconfig** utility to create an SAP ASE.

The 15.7 and later version of **sybperf** exposes a set of SAP ASE counters that are more useful for monitoring SAP ASE performance.

sql server clock tick length

sql server clock tick length specifies the duration of the server's clock tick, in microseconds.

Summary Information	
Default value	Platform-specific
Range of values	Platform-specific minimum–1000000, in multiples of default value
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Both the default value and the minimum value are platform-specific. SAP ASE rounds values up to an even multiple of *n*, where *n* is the platform-specific clock-tick default value. Use **sp_helpconfig** or **sp_configure** to find the current values for **sql server clock tick length**.

In mixed-use applications with some CPU-bound tasks, decrease the value of **sql server clock tick length** to:

- Help I/O-bound tasks – a value of 20,000 is reasonable for this. Shortening the clock-tick length means that CPU-bound tasks exceed the allotted time on the engine more frequently per unit of time, which allows other tasks greater access to the CPU
- Marginally increase response times – SAP ASE runs its service tasks once per clock tick. Decreasing the clock-tick length means that the service tasks are run more frequently per unit of time

Increasing **sql server clock tick length** favors CPU-bound tasks, because they execute longer between context switches. The maximum value of 1,000,000 may be appropriate for primarily CPU-bound applications. However, any I/O-bound tasks may suffer as a result. You can mitigate this somewhat by tuning **cpu grace time**.

Note: Changing the value of **sql server clock tick length** can have serious effects on SAP ASE performance. Consult with SAP Technical Support before resetting this value.

sql text pipe active

sql text pipe active controls whether SAP ASE collects SQL text.

Summary Information	
Default value	0
Range of values	0–1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

If this option is enabled and **sql text pipe max messages** is set, SAP ASE collects the SQL text for each query. Use `monSysSQLText` to retrieve the SQL text for all user tasks.

sql text pipe max messages

sql text pipe max messages specifies the number of SQL text messages SAP ASE stores per engine.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary Information	
Configuration groups	Memory Use, Monitoring

The total number of messages in the `monSQLText` table is the value of **sql text pipe max messages** multiplied by the number of engines running.

stack guard size

stack guard size sets the size, in bytes, of the stack guard area, which is an overflow stack of configurable size at the end of each stack.

Summary Information	
Default value	4096
Range of values	0–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, User Environment

SAP ASE allocates one stack for each user connection and worker process when it starts. These stacks are located contiguously in the same area of memory, with a guard area at the end of each stack. At the end of each stack guard area is a guardword, which is a 4-byte structure with a known pattern. The figure below illustrates how a process can corrupt a stack guardword.

Note: (On UNIX platforms only) SAP recommends that you include an additional 4096 bytes when you configure stack guard size to a nondefault value to increase the usable portion of the stack guard area.

Figure 6: Process about to corrupt stack guardword

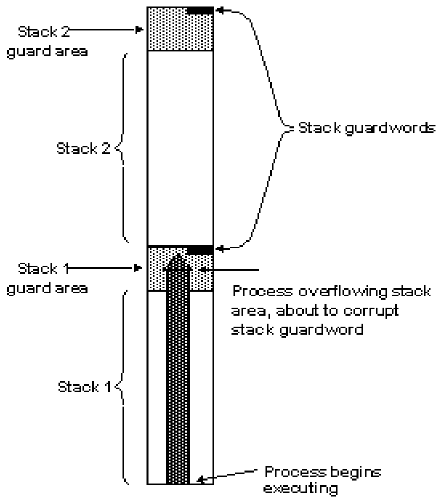

SAP ASE periodically checks to see whether the stack pointer for a user connection has entered the stack guard area associated with that user connection's stack. If it has, SAP ASE aborts the transaction, returns control to the application that generated the transaction, and generates error 3626:

```
The transaction was aborted because it used too much
stack space. Either use sp_configure to increase the
stack size, or break the query into smaller pieces.
spid: %d, suid: %d, hostname: %.*s, application name:
%.*s
```

SAP ASE also periodically checks the guardword pattern to see if it has changed, thus indicating that a process has overflowed the stack boundary. When this occurs, SAP ASE prints these messages to the error log and shuts down:

```
kernel: *** Stack overflow detected: limit: 0x%lx sp: 0x%lx
kernel: *** Stack Guardword corrupted
kernel: *** Stack corrupted, server aborting
```

In the first message, “limit” is the address of the end of the stack guard area, and “sp” is the current value of the stack pointer.

In addition, SAP ASE periodically checks the stack pointer to see whether it is completely outside both the stack and the stack guard area for the pointer's process. If it is, SAP ASE shuts down, even if the guardword is not corrupted. When this happens, SAP ASE prints the following messages to the error log:

```
kernel: *** Stack overflow detected: limit: 0x%lx sp: 0x%lx
kernel: *** Stack corrupted, server aborting
```

The default value for **stack guard size** is appropriate for most applications. However, if you experience server shutdown from either stack guardword corruption or stack overflow, increase **stack guard size** by a 2K increment. Each configured user connection and worker process has a stack guard area; thus, when you increase **stack guard size**, you use up that amount of memory, multiplied by the number of user connections and worker processes you have configured.

Rather than increasing **stack guard size** to avoid stack overflow problems, consider increasing **stack size** . The stack guard area is intended as an overflow area, not as an extension to the regular stack.

SAP ASE allocates stack space for each task by adding the values of the **stack size** and **stack guard size** parameters. **stack guard size** must be configured in multiples of 2K. If the value you specify is not a multiple of 2K, **sp_configure** verification routines round the value up to the next highest multiple.

stack size

stack size specifies the size, in bytes, of the execution stacks used by each user process on SAP ASE.

Summary Information	
Default value	Platform-specific
Range of values	Platform-specific minimum–2147483647
Status	Static
Display level	Basic
Required role	System administrator
Configuration group	User Environment

To find the **stack size** values for your platform, use **sp_helpconfig** or **sp_configure**. **stack size** must be configured in multiples of 2K. If the value you specify is not a multiple of 2K, **sp_configure** verification routines round the value up to the next highest multiple.

An execution stack is an area of SAP ASE memory where user processes keep track of their process context and store local data.

Certain queries can contribute to the probability of a stack overflow. Examples include queries with extremely long **where** clauses, long select lists, deeply nested stored procedures, and multiple selects and updates that **holdlock**. When a stack overflow occurs, SAP ASE prints an error message and rolls back the transaction. See the *Troubleshooting and Error Messages Guide* for more information on specific error messages.

The two options for remedying stack overflows are to break the large queries into smaller queries and to increase **stack size**. Changing **stack size** affects the amount of memory required for each configured user connection and worker process.

If you have queries that exceed the size of the execution stack, you may want to rewrite them as a series of smaller queries, especially if there are only a small number of such queries, or if you run them infrequently.

There is no way to determine how much stack space a query requires without actually running the query. Stack space for each user connection and worker process is preallocated at start-up.

Therefore, determining the appropriate value for **stack size** is an empirical process. Test your largest and most complex queries using the default value for **stack size**. If they run without generating error messages, the default is probably sufficient. If they generate error messages, begin by increasing **stack size** by a small amount (2K). Re-run your queries and see if the amount you have added is sufficient. If it is not, continue to increase **stack size** until queries run without generating error messages.

If you are using CIS, or if Java is enabled in the database and you want to use methods that call JDBC, SAP recommends that you increase the default by 50 percent. If you are not using JDBC or CIS, the standard default value is usually sufficient.

start mail session

(Windows only) **start mail session** enables and disables the automatic initiation of an SAP ASE mail session when you start SAP ASE.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Extended Stored Procedure

A value of 1 configures SAP ASE to start a mail session the next time SAP ASE is started. A value of 0 configures SAP ASE not to start a mail session at the next restart.

If **start mail session** is 0, you can start an SAP ASE mail session explicitly, using the **xp_startmail** system ESP.

Before setting **start mail session**, you must prepare your Windows system by creating a mailbox and mail profile for SAP ASE. Then create an SAP ASE account for Sybmail. See the *Configuration Guide for Windows*.

start xp server during boot

start xp server during boot determines whether XP Server starts when SAP ASE starts.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Static
Display level	
Required role	
Configuration group	Extended Stored Procedures

When set to 1, XP Server starts when SAP ASE starts. If you set **start xp server during boot** to 0, XP Server does not start until you run **xp_cmdshell**.

startup delay

startup delay controls when RepAgent is started during the server start.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

By default, RepAgent starts at the same time as SAP ASE. SAP ASE writes a message to the error log stating the wait time.

statement cache size

statement cache size increases the server allocation of procedure cache memory and limits the amount of memory from the procedure cache pool used for cached statements.

Summary Information	
Default value	0
Valid values	Size of cache in 2K pages
Status	Dynamic

Summary Information	
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, SQL Server Administration

Note: If you enable the statement cache, you must configure **set chained on/off** in its own batch.

Because cached statements are transformed into lightweight stored procedures, statement caching requires additional open object descriptors.

statement pipe active

statement pipe active controls whether SAP ASE collects statement-level statistics.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

If both **statement pipe active** and **statement pipe max messages** are enabled, SAP ASE collects the statement statistics for each query. Use `monSysStatement` to retrieve the statistics for all executed statements.

statement pipe max messages

statement pipe max messages determines the number of statement statistics messages SAP ASE stores per engine.

Summary Information	
Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Monitoring

The total number of messages in the `monSQLText` table is the value of **sql text pipe max messages** multiplied by the number of engines running.

statement statistics active

statement statistic active controls whether SAP ASE collects monitoring table statement-level statistics.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

Use `monProcessStatement` to get statement statistics for a specific task.

streamlined dynamic SQL

streamlined dynamic SQL enables the statement cache to store dynamic SQL statements.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

Note: The default value for **streamlined dynamic SQL** depends on the value to which **enable functionality group** is set. If you set **enable functionality group** to:

- 0 – the default value for **streamlined dynamic SQL** is 0.
- 1 – the default value for **streamlined dynamic SQL** is 1.

However, if you set **streamlined dynamic SQL** to 1, it uses a value of 1 regardless of what you set **enable functionality group** to.

See **enable functionality group**.

See *Performance and Tuning Series: Basics > Memory Use and Performance*.

strict dtm enforcement

strict dtm enforcement determines whether or not SAP ASE transaction coordination services strictly enforce the ACID properties (atomicity, consistency, integrity and durability) of distributed transactions.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Static
Display level	10
Required role	System administrator
Configuration group	DTM Administration

In environments where SAP ASE should propagate and coordinate transactions only to other SAP ASEs that support transaction coordination, set **strict dtm enforcement** on. If a transaction attempts to update data in a server that does not support transaction coordination services, SAP ASE aborts the transaction.

In heterogeneous environments, you may want to make use of servers that do not support transaction coordination. This includes earlier versions of SAP ASE and non-SAP database stores configured using CIS. Under these circumstances, set **strict dtm enforcement** off to allow SAP ASE to propagate transactions to legacy SAP ASEs and other data stores. This does not, however, ensure that the remote work of these servers is rolled back or committed with the original transaction.

suppress js max task message

suppress js max task message prevents SAP ASE from printing the Job Scheduler **js maxtask** error messages to the error log.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	Application Functionality

Note: The default value for **suppress js max task message** depends on the value to which **enable functionality group** is set. If you set **enable functionality group** to:

- 0 – the default value for **suppress js max task message** is 0.
- 1 – the default value for **suppress js max task message** is 1.

However, if you set **suppress js max task message** to 1, it uses a value of 1 regardless of what you set **enable functionality group** to.

See **enable functionality group**.

suspend audit when device full

suspend audit when device full determines what SAP ASE does when an audit device becomes completely full.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

Note: If you have two or more audit tables, each on a separate device other than the master device, and you have a threshold procedure for each audit table segment, the audit devices should never become full. Only if a threshold procedure is not functioning properly does the “full” condition occur.

Choose one of these values:

- 0 – truncates the next audit table and starts using it as the current audit table when the current audit table becomes full. If you set **suspend audit when device full** to 0, you ensure that the audit process is never suspended. However, you incur the risk that older audit records are lost if they have not been archived.
- 1 – suspends the audit process and all user processes that cause an auditable event. To resume normal operation, the system security officer must log in and set up an empty table as the current audit table. During this period, the system security officer is exempt from normal auditing. If the system security officer’s actions would generate audit records under normal operation, SAP ASE sends an error message and information about the event to the error log.

syb_sendmsg port number

syb_sendmsg port number specifies the port number that SAP ASE uses to send messages to a UDP (User Datagram Protocol) port with **sp_sendmsg** or **syb_sendmsg**.

Summary Information	
Default value	0
Valid values	0, or 1024–65535, or system limit
Status	Dynamic
Display level	Comprehensive
Required role	System security officer
Configuration group	Network Communication

If more than one engine is configured, a port is used for each engine, numbered consecutively from the port number specified. If the port number is set to the default value, 0 SAP ASE assigns port numbers.

Note: Sending messages to UDP ports is not supported on Windows.

A system security officer must set the **allow sendmsg** configuration parameter to 1 to enable sending messages to UDP ports. To enable UDP messaging, a system administrator must set **allow sendmsg** to 1. For more information on UDP messaging, see **sp_sendmsg** in the *Reference Manual: Procedures*.

sysstatistics flush interval

sysstatistics flush interval determines the length of the interval, in minutes, between flushes of `sysstatistics`.

Summary Information	
Default value	0
Valid values	0 – 32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

SAP ASE dynamically maintains the statistics for the number of rows and columns modified in a table as part of any DML statement and flushes them according to the value of **sysstatistics flush interval**.

SAP ASE uses these statistics for query optimization since they are more accurate. The **datachange** function determines the amount of data that is changed at the table, column, or partition level since the last **update statistics**, and initiates updating statistics on the object.

The in-memory statistics are always flushed to disk during a polite shutdown of the server. You can configure **sysstatistics flush interval** to flush these in-memory statistics to disk by the house keeper task at regular intervals. Set **sysstatistics flush interval** to 0 to disable this housekeeper task.

systemwide password expiration

systemwide password expiration sets the number of days that passwords remain in effect after they are changed.

Summary Information	
Default value	0
Range of values	0–32767
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

If **systemwide password expiration** is set to 0, passwords do not expire.

The password expires when the number of specified days passes. For example, if you create a new login on August 1, 2013 at 10:30 a.m., with a password expiration interval of 30 days, the password expires on August 31, 2013 at 10:30 a.m.

An account's password is considered expired if an interval greater than *number_of_days* has passed since the last time the password for that account was changed.

When the number of days remaining before expiration is less than 25 percent of the value of **systemwide password expiration** or 7 days, whichever is greater, each time the user logs in, a message displays, giving the number of days remaining before expiration. Users can change their passwords anytime before expiration.

systemwide password expiration is superseded by the password policy settings. For more information, see *Choose and Create a Password* in the *Security Administration Guide*.

When an account's password has expired, the user can still log in to SAP ASE but cannot execute any commands until he or she has used **alter login** to change his or her password.

This restriction applies only to login sessions established after the password has expired. Users who are logged in when their passwords expire are not affected until the next time they log in.

tape retention in days

tape retention in days specifies the number of days you intend to retain each tape after it has been used for either a database or a transaction log dump. This parameter can keep you from accidentally overwriting a dump tape.

Summary Information	
Default value	0
Range of values	0–365
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Backup/Recovery

For example, if you have set **tape retention in days** to 7 days, and you attempt to use the tape before 7 days have elapsed since the last time you dumped to that tape, Backup Server issues a warning message.

You can override the warning using the **with init** option when executing the dump command. Doing this causes the tape to be overwritten and all data on the tape to be lost.

Both the **dump database** and **dump transaction** commands provide a **retaindays** option, which overrides the **tape retention in days** value for a particular dump. See *System Administration Guide: Volume 2 > Backing Upa and Restoring User Databases*.

tcp no delay

tcp no delay controls TCP (Transmission Control Protocol) packet batching. The default value means that TCP packets are not batched.

Summary Information	
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration groups	Network Communication, O/S Resources

TCP normally batches small logical packets into single, larger, physical packets, and fills physical network frames with as much data as possible, which improves network throughput in terminal emulation environments where users mostly send keystrokes across the network.

However, applications that use small TDS (Tabular Data Stream) packets may benefit from disabling TCP packet batching.

Note: Disabling TCP packet batching means that packets are sent, regardless of size; this increases the volume of network traffic.

text prefetch size

text prefetch size limits the number of pages of `text`, `unitext`, and `image` data that can be prefetched into an existing buffer pool.

Summary Information	
Default value	16
Valid values	0–65535
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration Group	Network Communications

SAP ASE prefetches only `text`, `unitext`, and `image` data that was created with SAP ASE 12.x or was upgraded using **dbcc rebuild_text**.

threshold event max messages

Determines the number of events SAP ASE stores in the `monThresholdEvent` table. Once the number of events in the `monThresholdEvent` monitoring table exceed this value, SAP ASE overwrites the oldest unread events with new events.

Table 5. Summary Information

Default value	0
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System Administrator
Configuration group	Memory User, Monitoring

threshold event monitoring

Enable or disables SAP ASE from recording threshold events.

Summary Information	
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration Group	Monitoring

time slice

time slice sets the number of milliseconds that the SAP ASE scheduler allows a task to run.

Summary Information	
Default value	100
Range of values	50–1000
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

If **time slice** is set too low, SAP ASE may spend too much time switching between tasks, which increases response time. If it is set too high, CPU-intensive tasks may monopolize engines, which also increases response time. The default value allows each task to run for 1/10 of a second before relinquishing the CPU to another task.

See *Performance and Tuning Series: Basics > Using Engines and CPUs*.

Use **sp_sysmon** to determine how **time slice** affects voluntary yields by SAP ASE engines. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

total data cache size

total data cache size reports the amount of memory, in kilobytes, that is currently available for data, index, and log pages. This parameter is a calculated value that is not directly user-configurable.

Summary Information	
Default value	0
Range of values	0 – 2147483647
Status	Calculated
Display level	Basic
Required role	System administrator
Configuration groups	Cache Manager, Memory Use

The amount of memory available for the data cache can be affected by a number of factors, including:

- The amount of physical memory available on your machine
- The values to which the following parameters are set:
  - **total logical memory**
  - **number of user connections**
  - **total procedure cache percent**
  - **number of open databases**
  - **number of open objects**
  - **number of open indexes**
  - **number of devices**

A number of other parameters also affect the amount of available memory, but to a lesser extent.

total logical memory

total logical memory displays the total logical memory for the current configuration of SAP ASE.

Summary Information	
Default value	N/A
Range of values	N/A
Status	Read-only
Display level	Intermediate

Summary Information	
Required role	System administrator
Configuration groups	Memory Use, Physical Memory

The total logical memory is the amount of memory that the SAP ASE current configuration uses. **total logical memory** displays the memory that is required to be available, but which may or may not be in use at any given moment. For information about the amount of memory in use at a given moment, see **total physical memory**. You cannot use **total logical memory** to set any of the memory configuration parameters.

total physical memory

total physical memory is a read-only configuration parameter that displays the total physical memory for the current configuration of SAP ASE.

Summary Information	
Default value	N/A
Range of values	N/A
Status	Read-only
Display level	Intermediate
Required role	System administrator
Configuration group	Memory Use

The total physical memory is the amount of memory that SAP ASE is using at a given moment in time. Configure SAP ASE so that the value for **max memory** is larger than the value for **total logical memory**, and the value for **total logical memory** is larger than the value for **total physical memory**.

transfer utility memory size

SAP ASE maintains a memory pool for the **transfer table** command and for tables marked for incremental transfer. This pool provides memory for maintaining state information about current and past transfers, and for memory used to write to and read from transfer files. **transfer utility memory size** determines the size of this memory pool.

Summary Information	
Default value	4096
Range of values	0 – 2147483647
Status	Dynamic
Display level	Intermediate

Summary Information	
Required role	System administrator
Configuration group	SQL Server Administration

The units for this pool are in memory pages, which are blocks of 2048 bytes. The default size is large enough to accommodate over 100 tables marked for incremental transfer, all transferring simultaneously.

If your installation does not include tables marked for incremental transfer, and does not use the **transfer table** command, you may set the size of this memory pool to zero to reclaim this memory.

txn to pss ratio

txn to pss ratio determines the total number of transaction descriptors that are available to the server.

Summary Information	
Default value	16
Valid values	1 – 2147483647
Status	Static
Display level	1
Required role	System administrator
Configuration groups	DTM Administration, Memory Use

SAP ASE manages transactions as configurable server resources. Each time a new transaction begins, SAP ASE must obtain a free *transaction descriptor* from a global pool that is created when the server is started. Transaction descriptors are internal memory structures that SAP ASE uses to represent active transactions.

SAP ASE requires one free transaction descriptor for:

- The outer block of each server transaction. The outer block of a transaction may be created explicitly when a client executes a new **begin transaction** command. SAP ASE may also implicitly create an outer transaction block when clients use Transact-SQL to modify data without using **begin transaction** to define the transaction.

  Note: Subsequent, nested transaction blocks, created with additional **begin transaction** commands, do not require additional transaction descriptors.

- Each database accessed in a *multidatabase transaction*. SAP ASE must obtain a new transaction descriptor each time a transaction uses or modifies data in a new database.

At start-up, this ratio is multiplied by the number of PSS structures to create the transaction descriptor pool:

```
# of transaction descriptors = PSS structures * txn to pss ratio
```

The default value, 16, ensures compatibility with versions of SAP ASE earlier than 12.x which also allocated 16 transaction descriptors for each user connection. In version 12.x and later, the number of simultaneous transactions is limited only by the number of transaction descriptors available in the server.

Note: You can have as many databases in a user transaction as there are in your SAP ASE installation. For example, if your SAP ASE has 25 databases, you can include 25 databases in your user transactions.

Optimizing the txn to pss ratio for your system

During a peak period, use **sp_monitorconfig** to examine the use of transaction descriptors.

For example:

```
sp_monitorconfig "txn to pss ratio"
```

```
Usage information at date and time: Apr 22 2002  2:49PM.
Name                           Num_free    Num_active  Pct_act  Max_Used
Reuse_cnt   Instance_Name
------------------------   ----------  ----------  -------  --------
---
----------  --------------------
txn to pss ratio                    784          80   10.20         523
         0                    NULL
```

If the `num_used` value is zero or very low, transactions may be delayed as SAP ASE waits for transaction descriptors to become free in the server. In this case, consider increasing the value of **txn to pss ratio**.

If the `Max_Used` value is too low, unused transaction descriptors may be consuming memory that can be used by other server functions. Consider reducing the value of **txn to pss ratio**.

unified login required

unified login required requires that all users who log in to SAP ASE be authenticated by a security mechanism.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

The **use security services** parameter must be 1 to use the unified login security service.

update statistics hashing

update statistics hashing enables SAP ASE to gather hash-based statistics.

Summary Information	
Default value	**off**
Range of values	One of: • **off** – no hashing • **on** – hashing on all columns • **partial** – hashing only for low unique count columns • **default** – off
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	General Information

The values for **update statistics** are character data, so you must use 0 as a placeholder for the second parameter, which must be numeric, and specify **off**, **on**, **partial**, or **default** for the third parameter. For example:

```
sp_configure "update statistics hashing", 0, partial
```

upgrade version

upgrade version reports the version of the upgrade utility that upgraded your master device. The upgrade utility checks and modifies this parameter during an upgrade.

Summary Information	
Default value	15000
Range of values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Warning! Do not reset **upgrade version**. Doing so may cause serious problems with SAP ASE.

You can determine whether an upgrade has been done on your master device by using **upgrade version** without specifying a value:

```
sp_configure "upgrade version"
```

use security services

use security services specifies that SAP ASE uses network-based security services.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Static
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

If the parameter is set to 0, none of the network-based security services can be used.

user log cache queue size

Determines whether a queueing strategy is used for logging.

Summary Information	
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System Administrator
Configuration group	User Environment

Setting **user log cache queue size** to:

- 1 – enables queuing for user log caches. The user log cache is divided into multiple cachelets, the number of which is dependent on the value of **user log cache size**.
- 0 – disables queuing for user log caches. The user log cache is not divided into multiple cachelets regardless of the value to which you set **user log cache size**.

user log cache size

user log cache size specifies the size, in bytes, for each user's log cache. Its size is determined by the server's logical page size.

Summary Information	
Default value	Logical page size

Summary Information	
Range of values	2048[a] –2147483647[1]
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, User Environment

There is one user log cache for each configured user connection and worker process. SAP ASE uses these caches to buffer the user transaction log records, which reduces the contention at the end of the transaction log.

When a user log cache becomes full or another event occurs (such as when the transaction completes), SAP ASE “flushes” all log records from the user log cache to the database transaction log. By first consolidating the log records in each user’s log cache, rather than immediately adding each record to the database’s transaction log, SAP ASE reduces contention of processes writing to the log, especially for SMP systems that are configured with more than one engine.

Note: For transactions using a database with mixed data and log segments, the user log cache is flushed to the transaction log after each log record. No buffering takes place. If your databases do not have dedicated log segments, do not increase the user log cache size.

Do not configure **user log cache size** to be larger than the maximum amount of log information written by an application’s transaction. Since SAP ASE flushes the user log cache when the transaction completes, any additional memory allocated to the user log cache is wasted. If no transaction in your server generates more than 4000 bytes of transaction log records, set **user log cache size** no higher than that value. For example:

```
sp_configure "user log cache size", 4000
```

Setting **user log cache size** too high wastes memory. Setting it too low can cause the user log cache to fill up and flush more than once per transaction, increasing the contention for the transaction log. If the volume of transactions is low, the amount of contention for the transaction log may not be significant.

Use **sp_sysmon** to understand how this parameter affects cache behavior. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

1 a. Minimum determined by server’s logical page size

user log cache spinlock ratio

For SAP ASEs running with multiple engines, **user log cache spinlock ratio** specifies the ratio of user log caches per user log cache *spinlock*. There is one user log cache for each configured user connection.

Summary Information	
Default value	20
Range of values	1–2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration groups	Memory Use, User Environment

The default specifies 1 spinlock for each 20 user connections configured for your server.

Use **sp_sysmon** to understand how this parameter affects cache behavior. See the *Performance and Tuning Series: Monitoring SAP Adaptive Server with sp_sysmon*.

utility statistics hashing

utility statistics hashing enables the gathering of index attributes using hash-based statistics when creating an index.

Summary information	
Default value	off
Range of values	0 (followed by parameter option)
Status	Dynamic
Display level	Intermediate
Required role	When granular permissions is enabled, `manage server configuration` is required. When granular permissions is disabled, sa_role is required.
Configuration groups	Lock Manager

Enable hashing with:

```
sp_configure "utility statistics hashing", 0, literal
```

where:

- **on** – index attributes are gathered using hash-based statistics.
- **off** – the sort-based algorithm from versions earlier than 15.7 ESD #2 is used.

- **new** – hashing is gathered for minor attributed columns that have not had statistics previously gathered.
- **default** – same as off.

wait event timing

wait event timing controls whether SAP ASE collects statistics for individual wait events.

Summary Information	
Default value	0
Range of values	0–1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	Memory Use, Monitoring

A task may have to wait for a variety of reasons (for example, waiting for a buffer read to complete). The `monSysWaits` table contains the statistics for each wait event. The `monWaitEventInfo` table contains a complete list of wait events.

wait on uncommitted insert

Allows you to control the wait behavior of the **update**,**insert**, and **delete** commands for an uncommitted insert.

Summary Information	
Default value	0
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration groups	SQL Server Administration

wait on uncommitted insert can only be used for DOL tables on transaction isolation 0, 1, and 2, and does not affect the behavior of transaction isolation 3.

When **wait on uncommitted insert** is set to 1:

- at isolation level 0, delete and update queries block on uncommitted inserted rows with the key value of interest, and select queries read dirty and do not block on uncommitted inserts

- at isolation levels 1 and 2, select, delete, and update queries, all block on uncommitted inserted rows

When **wait on uncommitted insert** is set to 0, the behavior of select, update, delete, and insert is the same as in previous releases.

Note: When **wait on uncommitted insert** is set to 1 (to wait for uncommitted inserted rows) concurrency might descend and deadlocks can be encountered where they were not encountered before.

---

workload manager cache size

workload manager cache size specifies the maximum amount of memory, in 2K pages, that the workload manager can use.

Summary Information	
Default value	80
Valid values	80 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

See *Clusters Users Guide > Managing the Workload*.

xact coordination interval

xact coordination interval defines the length of time between attempts to resolve transaction branches that have been propagated to remote servers.

Summary Information	
Default value	60 (seconds)
Valid values	1 – 2147483647 (seconds)
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	DTM Administration

The coordinating SAP ASE makes regular attempts to resolve the work of remote servers participating in a distributed transaction. The coordinating server contacts each remote server participating in the distributed transaction in a serial manner, as shown in the figure below. The coordination service may be unable to resolve a transaction branch for a variety of reasons. For example, if the remote server is not reachable due to network problems, the coordinating server reattempts the connection after the time specified by **xact coordination level**.

Figure 7: Resolving remote transaction branches

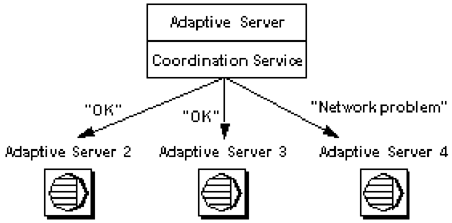

With the default value of **xact coordination interval**, 60, SAP ASE attempts to resolve remote transactions once every minute. Decreasing the value may speed the completion of distributed transactions, but only if the transactions are themselves resolved in less than a minute. Under normal circumstances, there is no performance penalty to decreasing the value of **xact coordination interval**.

Setting **xact coordination interval** to a higher number can slow the completion of distributed transactions, and cause transaction branches to hold resources longer than they normally would. Under normal circumstances, do not increase the value of **xact coordination interval** beyond its default.

xp_cmdshell context

xp_cmdshell context sets the security context for the operating system command to be executed using the **xp_cmdshell** system ESP.

Summary Information	
Default value	1
Valid values	0, 1, 2
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Extended Stored Procedure

The values for the context determines under which account the command runs:

- 0 – command runs under XP Server's account.
- 1 – command runs under user's account.
- 2 – command runs under XP Server's account only if the user has administrator privileges.

Setting **xp_cmdshell context** to 1 restricts the **xp_cmdshell** security context to users who have accounts at the operating system level. Its behavior is platform-specific. If **xp_cmdshell**

context is set to 1, to use an **xp_cmdshell** ESP, an operating system user account must exist for the SAP ASE user name. For example, an SAP ASE user named “sa” cannot use **xp_cmdshell** unless he or she has an operating-system-level user account named “sa”.

Starting XP Server as root automatically sets **xp_cmdshell** to 1 so secure access is automatically enabled.

On Windows, when **xp_cmdshell context** is set to 1, **xp_cmdshell** succeeds only if the user name of the user logging in to SAP ASE is a valid Windows user name with Windows system administration privileges on the system on which SAP ASE is running.

On other platforms, when **xp_cmdshell context** is set to 1, **xp_cmdshell** succeeds only if SAP ASE was started by a user with “superuser” privileges at the operating system level. When SAP ASE gets a request to execute **xp_cmdshell**, it checks the `uid` of the user name of the ESP requestor and runs the operating system command with the permissions of that `uid.`

If **xp_cmdshell context** is 0, the permissions of the operating system account under which SAP ASE is running are the permissions used to execute an operating system command from **xp_cmdshell**. This allows users to execute operating commands that they would not ordinarily be able to execute under the security context of their own operating system accounts.

CHAPTER 6 Disk Resource Issues

Many SAP ASE defaults are set to reasonable values for aspects of storage management, such as database, table, and index location, and how much space is allocated for each one. Responsibility for storage allocation and management is often centralized, and usually, the system administrator has ultimate control over the allocation of disk resources to SAP ASE and the physical placement of databases, tables, and indexes on those resources.

Device Allocation and Object Placement

When configuring a new system, the system administrator must consider several issues that have a direct impact on the number and size of disk resources required. These device allocation issues refer to commands and procedures that add disk resources to SAP ASE.

Table 6. Device allocation documentation

Task	See
Initialize and allocate a default pool of database devices	*System Administration Guide: Volume 1 > Initializing Database Devices*
Mirror database devices for recovery	*System Administration Guide: Volume 2 > Mirroring Database Devices*

After the initial disk resources have been allocated to SAP ASE, the system administrator, database owner, and object owners should consider how to place databases and database objects on specific database devices. These object placement issues determine where database objects reside on your system and whether or not the objects share devices.

Table 7. Object placement

Task	See
Place databases on specific database devices	*System Administration Guide: Volume 2 > Creating and Managing User Databases*
Place tables and indexes on specific database devices	*System Administration Guide: Volume 2 > Creating and Managing User Databases*

Do not consider allocating devices separately from object placement. For example, if you decide that a particular table must reside on a dedicated pair of devices, first allocate those devices to SAP ASE.

Commands for Managing Disk Resources

SAP ASE offers many commands and strategies for managing disk resources.

Command	Task	See
`disk init name = "dev_name" physname = "phys_name"...`	Makes a physical device available to a particular SAP ASE. Assigns a database device name (*dev_name*) that is used to identify the device in other SAP ASE commands.	*System Administration Guide: Volume 1 > Initializing Database Devices*
`sp_deviceattr logical-name, optname, optvalue`	Changes the *dsync* setting of an existing database device file.	*System Administration Guide: Volume 1 > Initializing Database Devices*
`sp_diskdefault "dev_name"...`	Adds *dev_name* to the general pool of default database space.	*System Administration Guide: Volume 1 > Initializing Database Devices*
`disk resize name = “device_name”, size = additional_space`	Dynamically increases the size of database devices.	*System Administration Guide: Volume 1 > Initializing Database Devices*
`disk mirror name = "dev_name" mirror = "phys_name"...`	Mirrors a database device on a specific physical device.	*System Administration Guide: Volume 2 > Mirroring Database Devices*

The table below lists the commands used in object placement. For information about how object placement affects performance, see *Performance and Tuning Series: Physical Database Tuning > Controlling Physical Data Placement*.

Table 8. Commands for placing objects on disk resources

Command	Task	See
`create database...on dev_name` or `alter database...on dev_name`	Makes database devices available to a particular SAP ASE database. The **log on** clause to **create database** places the database’s logs on a particular database device.	*System Administration Guide: Volume 2 > Creating and Managing User Databases*

Command	Task	See

`create database...`
or
`alter database...`
	When used without the **on** *dev_name* clause, these commands allocate space on the default database devices.	*System Administration Guide: Volume 2 > Creating and Managing User Databases*

`sp_addsegment` *`seg_name`*`,` *`dbname`*`,` *`devname`* and `sp_extendsegment` *`seg_name`*`,` *`dbname`*`,` *`devname`*	Creates a segment—a named collection of space—from the devices available to a particular database.	*System Administration Guide: Volume 2 > Creating and Using Segments*
`create table...on` *`seg_name`* or `create index...on` *`seg_name`*	Creates database objects, placing them on a specific segment of the database's assigned disk space.	*System Administration Guide: Volume 2 > Creating and Using Segments*

`create table...`
or
`create index...`
	When used without **on** *seg_name*, tables and indexes occupy the general pool of space allocated to the database (the default devices).	*System Administration Guide: Volume 2 > Creating and Using Segments in System*

Considerations in Storage Management Decisions

The system administrator must make many decisions regarding the physical allocation of space to SAP ASE databases.

The major considerations are:

- Recovery – disk mirroring and maintaining logs on a separate physical device provide two mechanisms for full recovery in the event of physical disk failures.
- Performance – for tables or databases where speed of disk reads and writes is crucial, properly placing database objects on physical devices yields performance improvements. Disk mirroring slows the speed of disk writes.

Recovery

Recovery is the key motivation for using several disk devices. You can ensure full recovery by storing a database's log on a separate physical device. You can also mirror database devices to achieve nonstop recovery.

- Storing logs on a separate device – unless a database device is mirrored, full recovery requires that a database's transaction log be stored on a different device from the actual data (including indexes) of a database.

  In the event of a hard disk failure, you can create an up-to-date database by loading a dump of the database and then applying the log records that were safely stored on another device. See *System Administration Guide: Volume 2 > Creating and Managing User Databases* for information about the **log on** clause of **create database**.
- Disk mirroring – Nonstop recovery in the event of a hard disk failure is guaranteed by mirroring all SAP ASE devices to a separate physical disk. See *System Administration Guide: Volume 2 > Mirroring Database Devices*.

Performance

You can improve system performance by placing logs and database objects on separate devices.

For example:

- Place a table on one hard disk and nonclustered indexes on another to ensure that physical reads and writes are faster, since the work is split between two disk drives.
- Split large tables across two disks to improve performance, particularly for multiuser applications.
- When log and data share devices, disable user log cache buffering of transaction log records.
- Use partitioning to provide multiple insertion points for a heap table, add a degree of parallelism to systems configured to perform parallel query processing, and make it possible to distribute a table's I/O across multiple database devices.

See *Performance and Tuning Series: Physical Database Tuning > Controlling Physical Data Placement* for a detailed discussion of how object placement affects performance.

Status and Defaults at Installation

The installation program and scripts initialize the master device and set up the `master`, `model`, `sybsystemprocs`, `sybsecurity`, and temporary databases for you.

When you install SAP ASE, the system databases, system-defined segments, and database devices are organized as follows:

- The `master`, `model`, and `tempdb` databases are installed on the master device.
- The `sybsystemprocs` database is installed on a device that you specified.
- Three segments are created in each database: `system`, `default`, and `logsegment`.
- The master device is the default storage device for all user-created databases.

Note: After initializing new devices for default storage, remove the master device from the default storage area with **sp_diskdefault**. Do not store user databases and objects on the master device.

- If you install the audit database, `sybsecurity`, it is located on its own device.

System Tables That Manage Storage

Two system tables in the `master` database, `sysusages` and `sysdevices`, and three more in each user database (`syssegments`, `sysindexes`, and `syspartitions`) track the placement of databases, tables (including the transaction log table, `syslogs`), and indexes.

The relationship between the tables is illustrated below.

Figure 8: System tables that manage storage

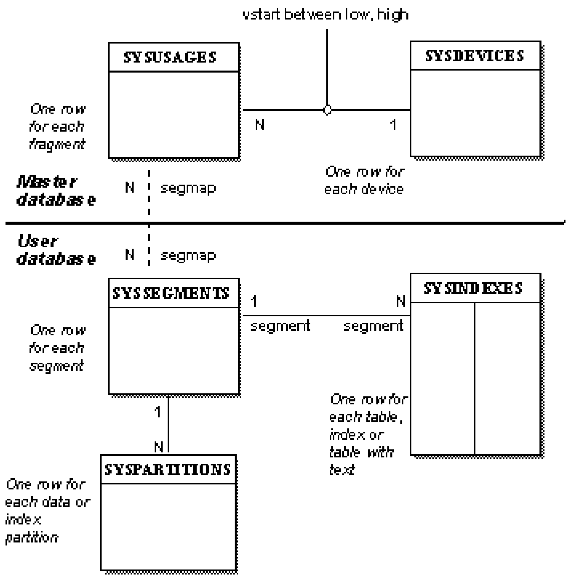

The sysdevices Table

The `sysdevices` table in the `master` database contains one row for each *database device* and may contain a row for each dump device (tape, disk, or operating system file) available to SAP ASE.

The **disk init** command adds entries for database devices to `master..sysdevices`. Dump devices, added using **sp_addumpdevice**, are discussed in *System Administration Guide: Volume 2 > Developing a Backup and Recovery Plan*.

`sysdevices` stores two names for each device:

- A *logical name* or *device name*, used in all subsequent storage-management commands, is stored in the `name` column of `sysdevices`. This is usually a user-friendly name, perhaps indicating the planned use for the device, for example, “logdev” or “userdbdev.”
- The *physical name* is the actual operating system name of the device. Use this name only in the **disk init** command; after that, all SAP ASE data storage commands use the logical name.

Place a database or transaction log on one or more devices by specifying the logical name of the device in the **create database** or **alter database** statement. The **log on** clause to **create database** places a database’s transaction log on a separate device to ensure full recoverability. The log device must also have an entry in `sysdevices` before you can use **log on**.

A database can reside on one or more devices, and a device can store one or more databases. See *System Administration Guide: Volume 2 > Creating and Managing User Databases* for information about creating databases on specific database devices.

The sysusages Table

The `sysusages` table in the `master` database keeps track of the space you assign to all SAP ASE databases.

create database and **alter database** allocate new space to the database by adding a row to `sysusages` for each database device or device fragment. When you allocate only a portion of the space on a device with **create** or **alter database**, that portion is called a *fragment*.

sp_addsegment, **sp_dropsegment**, and **sp_extendsegment** change the `segmap` column in `sysusages` for the device that is mapped or unmapped to a segment. See *System Administration Guide: Volume 2 > Creating and Using Segments*.

The syssegments Table

The `syssegments` table, one in each database, lists the segments in a database.

A *segment* is a collection of the database devices and fragments available to a particular database. Tables and indexes can be assigned to a particular segment—and therefore to a particular physical device—or can span a set of physical devices.

create database makes default entries in `syssegments`. **sp_addsegment** and **sp_dropsegment** to add and remove entries from `syssegments`.

The sysindexes Table

The `sysindexes` table lists each table and index and the segment where each table, clustered index, nonclustered index, and chain of text pages is stored. It also lists other information, such as the **max_rows_per_page** setting for the table or index.

The **create table**, **create index**, and **alter table** commands create new rows in `sysindexes`. Partitioning a table changes the function of `sysindexes` entries for the table.

The syspartitions Table

The `syspartitions` table lists each table and index partition and the segment where the partition is stored.

`syspartitions` maintains key storage management information such as the first page of a data or index page chain, the last page of a heap, the root page of an index partition, and so on.

Use **create table**, **create index** and **alter table** to create new rows in `syspartitions`.

CHAPTER 7 Managing Remote Servers

Users on a local SAP ASE can execute stored procedures on a remote SAP ASE. Executing an *remote procedure calls* (RPC) sends the results of the remote process to the calling process, which usually appears on the user's screen.

Users on a local SAP ASE can execute stored procedures on a remote SAP ASE. Executing an RPC sends the results of the remote process to the calling process, which usually appears on the user's screen.

To enable RPCs, the system administrator and system security officer of each SAP ASE must execute the following steps:

- On the local server:
  - System security officer – use **sp_addserver** to list the local server and remote server in the system table `master..sysservers`.
  - List the remote server in the `interfaces` file or directory service for the local server.
  - Restart the local server so the global variable *@@servername* is set to the name of the local server. If this variable is not set properly, users cannot execute RPCs from the local server on any remote server.
- On the remote server:
  - System security officer – use **sp_addserver** to list the server originating the RPC in the system table `master..sysservers`.
  - To allow the user who is originating the remote procedure access to the server, a system security officer uses **create login**, and a system administrator uses **sp_addremotelogin**.
  - Add the remote login name as a user of the appropriate database and grant that login permission to execute the procedure. (If **execute** permission is granted to "public," the user does not need to be granted specific permission.)

The figure below shows how to set up servers for remote access.

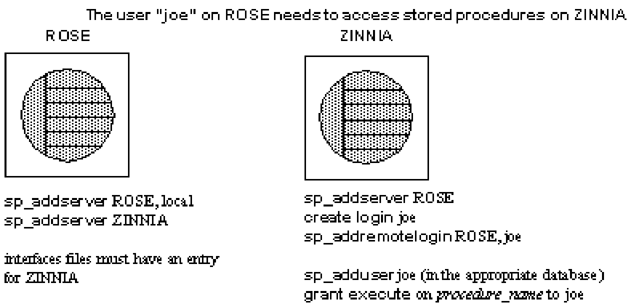

For operating-system-specific information about handling remote servers, see the installation documentation for your platform.

Adding Remote Logins

The system security officer and system administrator of any SAP ASE share control over which remote users can access the server, and the identity that remote users assume.

The system administrator uses **sp_addremotelogin** to add remote logins and **sp_dropremotelogin** to drop remote logins. The system security officer uses **sp_remoteoption** to control whether password checking is required.

Map Users' Server IDs

You can map logins from a remote server to a local server.

You can map:

- A particular remote login to a particular local login name. For example, user “joe” on the remote server might be mapped to “joesmith”.
- All logins from one remote server to one local name. For example, all users sending remote procedure calls from the MAIN server might be mapped to “remusers”.
- All logins from one remote server to use their remote names.

The first option can be combined with the other two options, and its specific mapping takes precedence over the other two more general mappings. The second and third options are mutually exclusive; you can use either of them, but not both.

Changing the mapping option

Use **sp_dropremotelogin** to remove the old mapping.

Use **sp_addremotelogin** to add remote logins:

```
sp_addremotelogin remoteserver [, loginame
     [, remotename]]
```

If the local names are not listed in `master..syslogins`, use **create login** to add them as SAP ASE logins before you add the remote logins.

Only a system administrator can execute **sp_addremotelogin**. See the *Reference Manual: Procedures*.

Map Remote Logins to Particular Local Names

You can map remote logins to specific local names.

This example maps the login “pogo” from a remote system to the local login “bob”.

The user logs in to the remote system as “pogo”. When “pogo” executes remote procedure calls from GATEWAY, the local system maps the remote login name to “bob”.

```
create logn bob with password itsA8secret
sp_addremotelogin GATEWAY, bob, pogo
```

Map All Remote Logins to One Local Name

You can create entries that map all remote logins to a single name.

This example creates an entry that maps all remote login names to the local name “albert”. All names are mapped to “albert”, except those with specific mappings. For example, if you mapped “pogo” to “bob”, and then the rest of the logins to “albert”, “pogo” still maps to “bob”.

```
create logn albert with password itsA8secret
sp_addremotelogin GATEWAY, albert
```

If you use **sp_addremotelogin** to map all users from a remote server to the same local name, use **sp_remoteoption** to specify the “trusted” option for those users. For example, if all users from server GATEWAY that are mapped to “albert” are to be trusted, specify:

```
sp_remoteoption GATEWAY, albert, NULL, trusted, true
```

If you do not specify logins as trusted, they cannot execute RPCs on the local server unless they specify passwords for the local server when they log in to the remote server. Users can run **ct_remote_pwd** to specify a password for server-to-server connections when they use Open Client Client-Library. **isql** and **bcp** do not permit users to specify a password for RPC connections.

Warning! Do not map more than one remote login to a single local login, as it reduces individual accountability on the server. Audited actions can be traced only to the local server login, not to the individual logins on the remote server.

If you are using network-based security

If users are logged in to the remote server using unified login, the logins must be designated as trusted on the local server, or they must specify passwords for the server when they log in to the remote server.

Warning! Using the **trusted** mode of **sp_remoteoption** reduces the security of your server, as passwords from such “trusted” users are not verified.

Keeping Remote Login Names for Local Servers

You can retain remote login names for your local server.

1. Use **create login** to create a login for each login from the remote server.
2. Use **sp_addremotelogin** for the server to create an entry in `master..sysremotelogins` with a null value for the remote login name and a value of -1 for the `suid`. For example:

   ```
   sp_addremotelogin GATEWAY
   ```

Example of Remote User Login Mapping

The remote user mapping procedures and the ability to set permissions for individual stored procedures give you control over which remote users can access local procedures.

For example, you can allow the "vp" login from the CORPORATE server to execute certain local procedures and all other logins from CORPORATE to execute the procedures for which the "admin" login has permission.

Note: Typically, the passwords for users on the remote server must match passwords on the local server.

This statement displays the local and remote server information recorded in `master..sysservers`:

```
select srvid, srvname from sysservers
```

```
srvid  srvname
-----  ----------
    0  SALES
    1  CORPORATE
    2  MARKETING
    3  PUBLICATIONS
    4  ENGINEERING
```

The SALES server is local. The other servers are remote.

This statement displays information about the remote servers and users stored in `master..sysremotelogins`:

```
select remoteserverid, remoteusername, suid
  from sysremotelogins
```

```
remoteserverid   remoteusername   suid
--------------   --------------   ------
1                joe              1
1                nancy            2
1                NULL             3
3                NULL             4
4                NULL             -1
```

By matching the value of `remoteserverid` in this result and the value of `srvid` in the previous result, you can find the name of the server for which the `remoteusername` is valid. For example, in the first result, `srvid` 1 indicates the CORPORATE server; in the second result, `remoteserverid` 1 indicates that same server. Therefore, the remote user login names "joe" and "nancy" are valid on the CORPORATE server.

The following statement shows the entries in `master..syslogins`:

```
select suid, name from syslogins
```

```
suid    name
------   ------------
    1   sa
```

```
    2   vp
    3   admin
    4   writer
```

The results of all three queries together show:

- The remote user name “joe” (`suid` 1) on the remote CORPORATE server (`srvid` and `remoteserverid` 1) is mapped to the “sa” login (`suid` 1).
- The remote user name “nancy” (`suid` 2) on the remote CORPORATE server (`srvid` and `remoteserverid` 1) is mapped to the “vp” login (`suid` 2).
- The other logins from the CORPORATE server (`remoteusername` “NULL”) are mapped to the “admin” login (`suid` 3).
- All logins from the PUBLICATIONS server (`srvid` and `remoteserverid` 3) are mapped to the “writer” login (`suid` 4).
- All logins from the ENGINEERING server (`srvid` and `remoteserverid` 4) are looked up in `master..syslogins` by their remote user names (`suid` -1).
- There is no `remoteserverid` entry for the MARKETING server in `sysremotelogins`. Therefore, users who log in to the MARKETING server cannot run remote procedure calls from that server.

Password Checking for Remote Users

A system security officer can use **sp_remoteoption** to determine whether passwords are checked when remote users log in to the local server.

By default, passwords are verified (this is the “untrusted” mode). In **trusted** mode, the local server accepts remote logins from other servers and front-end applications without user-access verification for the particular login.

When **sp_remoteoption** is used with arguments, it changes the mode for the named user:

```
sp_remoteoption [remoteserver, loginame, remotename,
     optname, {true | false}]
```

For example, to set **trusted** mode for the user “bob”, enter

```
sp_remoteoption GATEWAY, pogo, bob, trusted,
    true
```

Effects of Using the Untrusted Mode

The effects of the untrusted mode depend on the user’s client program.

isql and some user applications require that logins have the same password on the remote server and the local server. You can write Open Client applications to allow local logins to have different passwords on different servers.

To change your password in “untrusted” mode, you must first change it on all the remote systems you access before you can change it on your local server. If you change your password on the local server first, when you issue the remote procedure call to execute **alter login** on the remote server, your passwords no longer match.

The syntax for changing your password on the remote server is:

```
remote_server...alter login login_name
    with password, caller_passwd
    modify password [immediately] new_password ]
```

See *Security Administration Guide > Managing SAP ASE Logins and Database Users*.

Getting Information About Remote Logins

sp_helpremotelogin displays information about the remote logins on a server.

This example shows the remote login “pogo” mapped locally to login name “bob”, with all other remote logins keeping their remote names:

```
sp_helpremotelogin
```

```
server     remote_user_name     local_user_name      options
---------    -----------------    ------------------    ---------
GATEWAY      **mapped locally**   **use local name**    untrusted
GATEWAY      pogo                 bob                   untrusted
```

Configuration Parameters for Remote Logins

Some configuration parameters affect RPCs.

This table lists the configuration parameters that affect RPCs. Configuration parameters are set using **sp_configure**, and most on this list do not take effect until you restart SAP ASE.

Configuration parameter	Default
allow remote access	1
number of remote logins	20
number of remote sites	10
number of remote connections	20
remote server pre-read packets	3

CHAPTER 8 Initializing Database Devices

A database device stores the objects that make up databases.

The term *device* does not necessarily refer to a distinct physical device: it can refer to any piece of a disk (such as a disk partition) or a file in the file system that is used to store databases and their objects.

Each database device or file must be prepared and made known to SAP ASE before it can be used for database storage. This process is called *initialization*.

After a database device has been initialized, it can be:

- Allocated to the default pool of devices for the **create** and **alter database** commands
- Assigned to the pool of space available to a user database
- Assigned to a user database and used to store one or more database objects
- Assigned to store a database's transaction logs

Using the disk init command

System administrators use the disk init command to initialize new database devices with the **disk init** command.

The **disk init** command:

- Maps the specified physical disk device or operating system file to a database device name
- Lists the new device in `master..sysdevices`
- Prepares the device for database storage

Note: Before you run **disk init**, see the installation documentation for your platform for information about choosing a database device and preparing it for use with SAP ASE. You may want to repartition the disks on your computer to provide maximum performance for your SAP databases.

disk init divides the database devices into *allocation units*, groups of 256 logical pages. The size of the allocation unit depends on which logical page size your server is configured for (2, 4, 8, or 16K). In each allocation unit, the **disk init** command initializes the first page as the allocation page, which contains information about the database (if any) that resides on the allocation unit.

Warning! After you run the **disk init** command, dump the `master` database. This makes recovery easier and safer in case `master` is damaged. See *System Administration Guide: Volume 2 > Restoring the System Databases*.

The **disk init** command initializes a physical device or file and makes it usable by SAP ASE.

See the *Reference Manual: Commands* for the **disk init** syntax.

Getting Information about Devices

sp_helpdevice provides information about the devices in the `sysdevices` table.

When used without a device name, **sp_helpdevice** lists all the devices available on SAP ASE. When used with a device name, it lists information about that device. Here, **sp_helpdevice** is used to report information about the master device:

```
sp_helpdevice master
device_name  physical_name  description
-----------  -------------- -----------------------------------------
--
master       d_master       special, default disk, physical disk, 30
MB
status      cntrltype    vdevno    vpn_low   vpn_high
------      ----------   -------    -------   ---------
3           0            0          0         10239
```

Each row in `master..sysdevices` describes:

- A dump device (tape, disk, or file) to be used for backing up databases, or
- A database device to be used for database storage.

The initial contents of `sysdevices` are operating-system-dependent. `sysdevices` entries usually include:

- One for the master device
- One for the `sybsystemprocs` database, which you can use to store additional databases such as `pubs2` and `sybsyntax`, or for user databases and logs
- Two for tape dump devices

If you installed auditing, there is a separate device for `sybsecurity`.

The `vpn_low` and `vpn_high` columns represent the page numbers that have been assigned to the device. For dump devices, these columns represent the media capacity of the device.

The `status` field indicates the type of device, whether a disk device is used as a default storage device when users issue a **create** or **alter database** command without specifying a database device, disk mirroring information, and **dsync** settings.

Table 9. Status bits in sysdevices

Bit	Meaning
1	Default disk (may be used by any **create** or **alter database** command that does not specify a location)
2	Physical disk

Bit	Meaning
4	Logical disk (not used)
8	Skip header (used with tape dump devices)
16	Dump device
32	Serial writes
64	Device mirrored
128	Reads mirrored
256	Secondary mirror side only
512	Mirror enabled
2048	Used internally; set after **disk unmirror, side = retain**
4096	Primary device needs to be unmirrored (used internally)
8192	Secondary device needs to be unmirrored (used internally)
16384	UNIX file device uses **dsync** setting (writes occur directly to physical media)

For more information about dump devices and **sp_addumpdevice**, See *System Administration Guide: Volume 2 Developing a Backup and Recovery Plan*.

Dropping devices

Use **sp_dropdevice** to drop database and dump devices.

The syntax is:

```
sp_dropdevice logicalname
```

You cannot drop a device that is in use by a database. You must drop the database first.

sp_dropdevice removes the device name from `sysdevices`. **sp_dropdevice** does not remove an operating system file; it only makes the file inaccessible to SAP ASE. Use operating system commands to delete a file after using **sp_dropdevice**.

Designating Default Devices

To create a pool of default database devices to be used by all SAP ASE users for creating databases, use **sp_diskdefault** after the devices are initialized.

sp_diskdefault marks these devices in `sysdevices` as default devices. Whenever users create (or alter) databases without specifying a database device, new disk space is allocated from the pool of default disk space.

The syntax for **sp_diskdefault** is:

```
sp_diskdefault logicalname, {defaulton | defaultoff}
```

After adding user devices, use the **defaultoff** option to remove the master device from the pool of default space:

```
sp_diskdefault master, defaultoff
```

The following designates `sprocdev`, the device that holds the `sybsystemprocs` database, a default device:

```
sp_diskdefault sprocdev, defaulton
```

SAP ASE can have multiple default devices. They are used in the order in which they appear in the `sysdevices` table (that is, alphabetical order). When the first default device is filled, the second default device is used, and so on.

Note: After initializing a set of database devices, you may want to assign them to specific databases or database objects rather than adding them to the default pool of devices. For example, you may want to make sure a table never grows beyond the size of a particular device.

Choosing Default and Nondefault Devices

sp_diskdefault lets you plan space usage for performance and recovery, while allowing users to create or alter databases.

Do not use these devices as default devices:

- The master device
- The device used for `sybsecurity`
- Any device intended solely for logs
- Devices where high-performance databases reside

You can use the device that holds `sybsystemprocs` for other user databases.

Note: If you are using disk mirroring or segments, exercise caution in deciding which devices you add to the default list. In most cases, devices that are to be mirrored, or databases that contain objects placed on segments should specifically allocate devices, rather than being made part of default storage.

Increasing the Size of Devices with disk resize

The **disk resize** command allows you to increase the size of your database devices dynamically, rather than initializing a new device.

For example, if `/SAP_ASE/testdev.dat` requires an additional 10MB of space, you can run **disk resize** and allocate this amount of space to the device. The **create** and **alter database** commands can use this added space.

Use **disk resize** to increase the size for both devices on raw partitions and for file systems. The minimum amount of space by which you can increase a device is 1MB or an allocation unit, whichever is greater.

Page Size	Allocation Unit Size	Minimum Incremental Size
2K	0.5MB	1MB
4K	1MB	1MB
8K	2MB	2MB
16K	4MB	4MB

You cannot use **disk resize** on dump or load devices.

Any properties that are set on the device continue to be set after you increase its size. That is, if a device has **dsync** set before you increase its size, it has **dsync** set afterwards. Also, any access rights that were set before you increased the size of the device remain set.

A user with the **sa_role** can execute the **disk resize** command, which:

- Updates the high value in `master....sysdevices`, and
- Prepares the additional space for database storage.

Use audit trails on **disk resize** to track the number of times a device is resized. The device being resized is always online and available for users during the resize operation.

Resizing a disk requires that:

- You have already initialized the device with **disk init**.
- *device_name* must refer to a valid logical device name.
- You disable mirroring while the resize operation is in progress. You can reestablish mirroring when the resize operation is complete.

In this example, the configuration of the device `testdev` is:

```
sp_helpdevice testdev
device_name  physical_name          description
  status  cntrltype  vdevno          vpn_low          vpn_high
-----------  -------------------    -------------
  ------- ---------  -------------   -------------    --------------
testdev      /SAP_ASE/dev/testdev.dat   special, dsync on, directio
off, physical disk, 10.00MB
  16386    0          1                      0              5119
```

To increase the size of `testdev` by 4MB using **disk resize**, enter:

```
disk resize
name = "test_dev",
size = "4M"
```

`testdev.dat` is now 14MB:

```
sp_helpdevice testdev
device_name  physical_name          description
```

```
  status  cntrltype  vdevno          vpn_low          vpn_high
-----------  -------------------    -------------
  ------- ---------  -------------   -------------    --------------
testdev      /SAP_ASE/dev/testdev.dat   special, dsync on, directio
off, physical disk, 14.00MB
  16386           0         1          0                  7167
```

See the *Reference Manual: Commands* for **disk resize** syntax.

Insufficient Disk Space

During the physical initialization of the disk, if an error occurs due to insufficient disk space, **disk resize** extends the database device to the largest size possible.

For example, on a server that uses 4K logical pages, if you try to increase the size of the device by 40MB, but only 39.5MB is available, the device is extended only by 39.5MB.

You cannot use **disk resize** to decrease the size of a device.

CHAPTER 9 Setting Database Options

Use **sp_dboption** to change settings, which remain in effect until they are changed, for an entire database.

sp_dboption:

- Displays a complete list of the database options when it is used without a parameter
- Changes a database option when used with parameters

You can change options only for user databases. You cannot change options for the `master` database. To change a database option in a user database (or to display a list of the database options), execute **sp_dboption** while using the `master` database.

The syntax is:

```
sp_dboption [dbname, optname, {true | false}]
```

Note: Changes to `model`'s database options do not affect `tempdb` or current user-defined multiple temporary databases when you restart SAP ASE. These changes appear only in databases that you create after you change the **model** database. Restarting SAP ASE clears objects and data contained in the temporary databases, but does not reset database options.

Database Option Descriptions

All users with access to the `master` database can execute **sp_dboption** with no parameters to display a list of the database options.

The report from **sp_dboption** looks like this:

```
sp_dboption
```

```
Settable database options.
--------------------
abort tran on log full
allow nulls by default
allow wide dol rows
async log service
auto identity
dbo use only
ddl in tran
delayed commit
enforce dump tran sequence
full logging for all
full logging for alter table
full logging for reorg rebuild
full logging for select into
identity in nonunique index
```

```
no chkpt on recovery
no free space acctg
read only
select into/bulkcopy/pllsort
single user
trunc log on chkpt
trunc. log on chkpt.
unique auto_identity index
```

For a report on which options have been set in a particular database, execute **sp_helpdb** in that database.

See the *Commands Reference: Procedures* for information about each database option in detail.

Viewing the Options on a Database

Use **sp_helpdb** to determine the options that are set for a particular database. **sp_helpdb** lists each active option in the "status" column of its output.

The following example shows that the **read only** option is turned on in `mydb`:

```
sp_helpdb mydb
```

```
name           db_size  owner  dbid   created          status
-----          -------  -----  ----   ------------     -------------
------
mydb           20.0 MB    sa     5    Mar 05, 2005     read only

device_fragments      size    usage              created        fre
e kbytes
-----------------     ------  ------------       ---------      ----
-------
master                10.0 MB  data and log      Mar 05
2005           1792

device                         segment
------------------------------ ------------------------------
master                         default
master                         logsegment
master                         system
```

To display a summary of the options for all databases, use **sp_helpdb** without specifying a database:

```
sp_helpdb
```

```
name            db_size  owner dbid   created             status
-------------   -------- ----- ----   ------------       -----------
--------
master         48.0 MB   sa      1   Apr 12, 2005      mixed log and
data
model           8.0 MB   sa      3   Apr 12, 2005      mixed log and
data
```

```
pubs2           20.0 MB   sa      6   Apr 12, 2005       select into/
     bulkcopy/pllsort, trunc log on chkpt, mixed log and data
sybsystemdb      8.0 MB   sa      5   Apr 12, 2005      mixed log and
data
sybsystemprocs 112.0 MB   sa      4   Apr
12, 2005       trunc log on chkpt,
     mixed log and data
tempdb           8.0 MB   sa      2   Apr 12, 2005       select into/
     bulkcopy/pllsort, trunc log on chkpt, mixed log and data
```

Displaying Currently Set Switches with sysoptions

The `number` column of `sysoptions` contains the switch ID for currently set switches.

`sysoptions` shows these switches:

- Trace flag set in the runserver file with the `-T` flag
- Trace flag set with **dbcc traceon(flag_number)** or **set switch serverwide on**
- Trace flags and switches set for a specific system process ID (spid) with **set switch on**

`sysoptions` shows only the switches that are visible to the user. Users cannot see switches set privately by other spids. The value for `number` is `Null` for all option categories other than switches.

`sysoptions` displays this information for switches:

- `spid` – spid for the current session.
- `name` – name of the switch. If an unnamed switch is set, `name` contains the number of the switch converted to a string.
- `category` – specifies the string `Switch`.
- `currentsetting` – set to 1 if the switch is set, 0 if the switch is not set.
- `defaultsetting` – contains 0.
- `scope` – indicates whether the switch is set server-wide or for the session. Values are:
  - 0 – switch is not set.
  - 8 – switch is set server-wide.
  - 16 – switch is set privately to the current spid.
  - 24 – switch is set server-wide and privately.
- `number` – contains the switch ID as an integer.

This query displays all `sysoptions` entries for switches:

```
select * from sysoptions
where category = 'Switch'
```

This query displays the switches set for the current session:

```
select * from sysoptions
where category = 'Switch'
and currentsetting = '1'
```

This shows the `sysoptions` output after setting trace flag 3604:

```
spid       name                                 category         curre
ntsetting      defaultsetting          scope                  number
------
----------------------------      ---------------  ---------------
      ----------------       ------------         --------------
    11           print_output_to_client              Switch
        1
                  0                8                          3604
```

CHAPTER 10 Configuring Character Sets, Sort Orders, and Languages

SAP provides both internationalization and localization support.

Internationalization is enabling an application to support multiple languages and cultural conventions.

An internationalized application uses external files to provide language-specific information at execution time. Because it contains no language-specific code, an internationalized application can be deployed in any native language environment without code changes. A single version of a software product can be adapted to different languages or regions, conforming to local requirements and customs without engineering changes. This approach to software development saves significant time and money over the lifetime of an application.

Localization is 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. One version of a software product may have many localized versions.

SAP ASE includes the character set definition files and sort order definition files required for data processing support for the major business languages in Western Europe, Eastern Europe, the Middle East, Latin America, and Asia.

SAP Language Modules provide translated system messages and formats for Chinese (Simplified), French, German, Japanese, Korean, Brazilian Portuguese, and Spanish. By default, SAP ASE comes with U.S. English message files.

This chapter describes the available character sets and language modules and summarizes the steps necessary to change the default character set, sort order, or message language for SAP ASE.

Advantages of Internationalized Systems

Designing an application to work outside its country of origin can seem daunting. Often, programmers think that internationalizing means hard-coding dependencies based on cultural and linguistic conventions for just one country.

A better approach is to write an internationalized application: that is, one that examines the local computing environment to determine what language to use and loads files containing language-specific information at runtime.

When you use an internationalized application, a single application can be deployed in all countries. This has several advantages:

- You write and maintain one application.
- The application can be deployed, without change, in new countries as needed. You need only supply the correct localization files.
- All sites can expect standard features and behavior.

A Sample Internationalized System

An internationalized system may include internationalized client applications, gateways, and servers running on different platforms in different native language environments.

For example, an international system might include the following components:

- Order processing applications in New York City, Mexico City, and Paris (Client-Library applications)
- An inventory control server in Germany (SAP ASE)
- An order fulfillment server in France (SAP ASE)
- A central accounting application in Japan (an Open Server application working with an SAP ASE)

In this system, the order processing applications:

- Query the inventory control server to determine if requested items are in stock
- Place orders with the order fulfillment server
- Send financial information to the accounting application

The inventory control server and the order fulfillment server respond to queries, and the accounting application collects financial data and generates reports.

The system looks like this:

Figure 9: Example of an international system

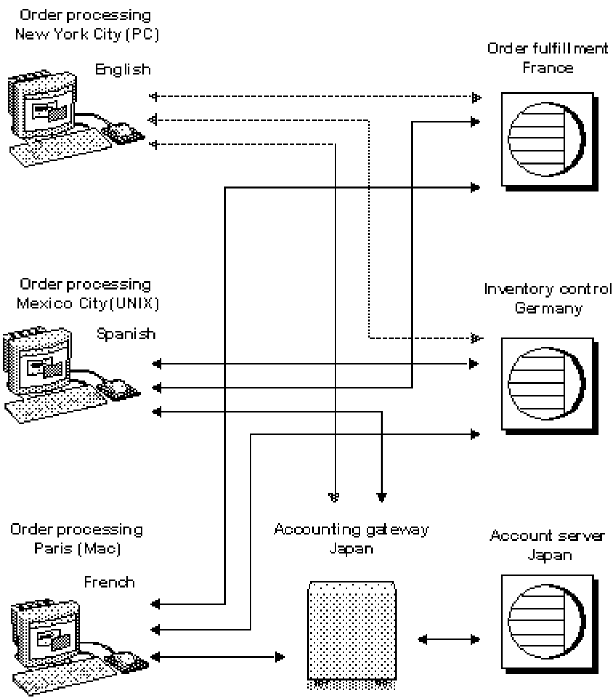

In this example, all applications and servers use local languages and character sets to accept input and output messages.

Elements of an Internationalized System

In an internationalized environment, you can manipulate the character set, the sort order, and system messages to configure your server language. SAP suggests that you review each of these elements and carefully plan the client/server network you want to create.

- Character set – the language in which the server sends and receives data to and from the client servers. Select the character set after carefully planning and analyzing the language needs of all client servers.
- Sort order – sort order options are dependent on the language and character set you select.
- System messages – messages display in one of several languages provided by SAP. If your server language is not one of the languages provided, your system messages display in English, the default.

Selecting the Character Set for Your Server

In your server, all data is encoded in a special code. For example, the letter “a” is encoded as “97” in decimal. A *character set* is a specific collection of characters (including alphabetic and numeric characters, symbols, and nonprinting control characters) and their assigned numerical values, or codes.

A character set generally contains the characters for an alphabet, for example, the Latin alphabet used in the English language, or a script such as Cyrillic used with languages such as Russian, Serbian, and Bulgarian. Character sets that are platform-specific and support a subset of languages, for example, the Western European languages, are called *native* or *national character sets*. All character sets that come with SAP ASE, except for Unicode UTF-8, are native character sets.

A *script* is a writing system, a collection of all the elements that characterize the written form of a human language—for example, Latin, Japanese, or Arabic. Depending on the languages supported by an alphabet or script, a character set can support one or more languages. For example, the Latin alphabet supports the languages of Western Europe (see table below). On the other hand, the Japanese script supports only one language, Japanese. Therefore, the Group 1 character sets support multiple languages, while many character sets, such as those in Group 101, support only one language.

The language or languages that are covered by a character set is called a *language group*. A language group can contain many languages or only one language; a native character set is the platform-specific encoding of the characters for the language or languages of a particular language group.

Within a client/server network, you can support data processing in multiple languages if all the languages belong to the same language group (see the table below). For example, if data in the server is encoded in a Group 1 character set, you could have French, German, and Italian data and any of the other Group 1 languages in the same database. However, you cannot store data from another language group in the same database. For example, you cannot store Japanese data with French or German data.

Unlike the native character sets just described, *Unicode* is 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.

Since all character sets support the Latin script, and therefore English, a character set always supports at least two languages—English and one other language.

Many languages are supported by more than one character set. The character set you install for a language depends on the client's platform and operating system.

Table 10. Supported languages and character sets

Language group	Languages	Character sets
Group 1	*Western European*: Albanian, Catalan, Danish, Dutch, English, Faeroese, Finnish, French, Galician, German, Icelandic, Irish, Italian, Norwegian, Portuguese, Spanish, Swedish	ASCII 8, CP 437, CP 850, CP 860, CP 863, CP 1252 , ISO 8859-1, ISO 8859-15, Macintosh Roman, ROMAN8, ROMAN9, ISO-15, CP 858 CP 1252 is identical to ISO 8859-1 except for the 0x80–0x9F code points which are mapped to characters in CP 1252.
Group 2	*Eastern European*: Croatian, Czech, Estonian, Hungarian, Latvian, Lithuanian, Polish, Romanian, Slovak, Slovene (and English)	CP 852, CP 1250, ISO 8859-2, Macintosh Central European
Group 4	Baltic (and English)	CP 1257
Group 5	*Cyrillic*: Bulgarian, Byelorussian, Macedonian, Russian, Serbian, Ukrainian (and English)	CP 855, CP 866, CP 1251, ISO 8859-5, Koi8, Macintosh Cyrillic
Group 6	Arabic (and English)	CP 864, CP 1256, ISO 8859-6
Group 7	Greek (and English)	CP 869, CP 1253, GREEK8, ISO 8859-7, Macintosh Greek
Group 8	Hebrew (and English)	CP 1255, ISO 8859-8
Group 9	Turkish (and English)	CP 857, CP 1254, ISO 8859-9, Macintosh Turkish, TURKISH8
Group 101	Japanese (and English)	CP 932 DEC Kanji, EUC-JIS, Shift-JIS
Group 102	Simplified Chinese (PRC) (and English)	CP 936, EUC-GB, GB18030
Group 103	Traditional Chinese (ROC) (and English)	Big 5, CP 950, EUC-CNS, Big 5 HKSCS CP 950 is identical to Big 5.
Group 104	Korean (and English)	EUC-KSC, cp949
Group 105	Thai (and English)	CP 874, TIS 620
Group 106	Vietnamese (and English)	CP 1258
Unicode	Over 650 languages	UTF-8

Note: The English language is supported by all character sets because the first 128 (decimal) characters of any character set include the Latin alphabet (defined as “ASCll-7”). The characters beyond the first 128 differ between character sets and are used to support the characters in different native languages. For example, code points 0-127 of CP 932 and CP 874 both support English and the Latin alphabet. However, code points 128-255 support Japanese characters in CP 932 and code points 128-255 support Thai characters in CP 874.

The following character sets support the European currency symbol, the "euro": CP 1252 (Western Europe); CP 1250 (Eastern Europe); CP 1251 (Cyrillic); CP 1256 (Arabic); CP 1253 (Greek); CP 1255 (Hebrew); CP 1254 (Turkish); CP 874 (Thai); iso15, roman9 and CP858. Unicode UTF-8 also supports:

- Traditional Chinese on the Windows and Solaris platforms
- Arabic, Hebrew, Thai, and Russian on the Linux platform

Note: iso_1 and ISO 8859-1 are different names for the same character set.

To mix languages from different language groups you *must* use Unicode. If your server character set is Unicode, you can support more than 650 languages in a single server and mix languages from any language group.

Unicode

Unicode enables all the world's languages to be encoded in the same data set.

Prior to the introduction of Unicode, if you wanted to store data in, for example, Chinese, you had to choose a character set appropriate for that language—to the exclusion of most other languages. It was either impossible or impractical to mix character sets, and thus diverse languages, in the same data set.

SAP supported Unicode in the form of three datatypes: `unichar`, `univarchar`, and `unitext`. These datatypes store data in the UTF-16 encoding of Unicode.

UTF-16 is an encoding wherein Unicode scalar values are represented by a single 16-bit value (or, in rare cases, as a pair of 16-bit values). The three encodings are equivalent insofar as either encoding can be used to represent any Unicode character. The choice of UTF-16 datatypes, rather than a UTF-16 server default character set, promotes easy, step-wise migration for existing database applications.

SAP ASE supports Unicode literals in SQL queries and a wide range of sort orders for UTF-8.

The character set model used by SAP ASE is based on a single, configurable, server-wide character set. All data stored in SAP ASE, using any of the "character" datatypes (`char`, `varchar`, `nchar`, `nvarchar`, and `text`), is interpreted as being in this character set. Sort orders are defined using this character set, as are language modules—collections of server messages translated into local languages.

During the connection dialog, a client application declares its native character set and language. If properly configured, the server thereafter attempts to convert any character data between its own character set and that of the client (character data includes any data stored in the database, as well as server messages in the client's native language).This works well as long as the server's and client's character sets are compatible. It does not work well when characters are not defined in the other character set, as is the case for the character sets SJIS, used for Japanese, and KOI8, used for Russian and other Cyrillic languages. Such incompatibilities are the reason for Unicode, which can be thought of as a character superset, including definitions for characters in all other character sets.

The Unicode datatypes `unichar`, `univarchar`, and `unitext` are completely independent of the traditional character set model. Clients send and receive Unicode data independently of whatever other character data they send and receive.

Configuration Parameters

The UTF-16 encoding of Unicode includes “surrogate pairs,” which are pairs of 16-bit values that represent infrequently used characters.

Additional checking is built in to SAP ASE to ensure the integrity of surrogate pairs. You can switch this checking off by setting the **enable surrogate processing** configuration parameter to 0. This yields slightly higher performance, although the integrity of surrogate pairs is no longer guaranteed.

Unicode also defines “normalization,” which is the process by which all possible representations of a single character are transformed into a single representation. Many base characters followed by combining diacritical marks are equivalent to precomposed characters, although their bit patterns are different. For example, the following two sequences are equivalent:

```
0x00E9  -- é (LATIN SMALL LETTER E WITH ACUTE)
```

```
0x00650301  -- e (LATIN SMALL LETTER E), ´ (COMBINING ACUTE ACCENT)
```

The **enable unicode normalization** configuration parameter controls whether or not SAP ASE normalizes incoming Unicode data.

Significant performance increases are possible when the **default Unicode sortorder** is set to “binary” and the **enable Unicode normalization** configuration parameter is set to 1. This combination allows SAP ASE to make several assumptions about the nature of the Unicode data, and code has been implemented to take advantage of these assumptions.

Functions

All functions that take `char` parameters accept `unichar` as well. Functions with more than one parameter, when called with at least one `unichar` parameter, results in implicit conversion of any non-`unichar` parameters to `unichar`.

To guarantee the integrity of surrogate pairs when **enable surrogate processing** is set to 1 (the default), the string functions do not allow surrogate pairs to be split. Positions fall at the beginning of a surrogate pair.

Several functions round out the `unichar` support. Included are the functions **to_unichar()** and **uscalar()**, which are analogous to **char()** and **ascii()**. The functions **uhighsurr()** and **ulowsurr()** allow the explicit handling of surrogate pairs in user code.

There are restrictions when using unitext with functions. For information, see the restriction description under the “Usage” section for each function in the *Reference Manual: Blocks*.

Using unichar Columns

When using the **isql** or **bcp** utilities, Unicode values display in hexadecimal form unless the `-Jutf8` flag is used, indicating the client's character set is UTF-8. In this case, the utility converts any Unicode data it receives from the server into UTF-8. For example:

```
% isql -Usa -P -Jiso_1
1> select unicode_name from people where unicode_name = 'Jones'
2> go
```

```
unicode_name
-------------------------------------------------------------------|
0x004a006f006e00650073
(1 row affected)
```

whereas:

```
% isql -Usa -P -Jutf8
1> select unicode_name from people where unicode_name = 'Jones'
2> go
```

```
unicode_name
-------------------------------------------------------------------
Jones
(1 row affected)
```

This facilitates ad hoc queries. Not all terminal windows are capable of displaying the full repertoire of Unicode characters, but simple tests involving ASCII characters are greatly simplified.

Using unitext

The variable-length `unitext` datatype can hold up to 1,073,741,823 Unicode characters (2,147,483,646 bytes). You can use `unitext` anywhere you use the `text` datatype, with the same semantics. `unitext` columns are stored in UTF-16 encoding, regardless of the SAP ASE default character set.

Open Client Interoperability

The Open Client libraries support the datatype `cs_unichar`, which can be bound to user variables declared as an array of short integers. This Open Client datatype interfaces directly with the server's `unichar`, `unitext`, and `univarchar`.

Java Interoperability

The internal JDBC driver efficiently transfers `unichar` data between SQL and Java contexts.

Going from SQL to Java, the class **java.sql.ResultSet** provides a number of "get" methods to retrieve data from the columns of a result set. Any of these get methods work with columns defined as `unichar`, `unitext`, or `univarchar`. The method **getString()** is particularly efficient since no conversion needs to be performed.

Use the **setString()** method of the class **java.sql.PreparedStatement** to go from Java to SQL. The internal JDBC driver copies Java string data directly into the SQL parameter defined as `unichar`, `unitext`, or `univarchar`.

The external JDBC driver (jConnect) has been modified to support the same seamless interface as the internal driver.

Selecting the Server Default Character Set

You must specify a default character set, which is the one in which the server stores and manipulates data, when you configure the server. Each server can have only one default character set.

By default, the installation tool assumes that the native character set of the platform operating system is the server's default character set. However, you can select any character set supported by SAP ASE as the default on your server (see the table below).

For example, if you are installing the server on IBM RS/6000 running AIX, and you select one of the Western European languages to install, the installation tool assumes the default character set to be ISO 8859-1.

If you are installing a Unicode server, select UTF–8 as your default character set.

For non-Unicode servers, determine what platform most of your client systems use and use the character set for this platform as the default character set on the server.

This has two advantages:

- The number of unmappable characters between character sets is minimized.
  Since there is usually not a complete one-to-one mapping between the characters in two character sets, there is a potential for some data loss. This is usually minor because most unconverted characters are special symbols that are not commonly used or are specific to a platform.
- This minimizes the character set conversion that is required.
  When the character set on the client system differs from the default character set on the server, data must be converted in order to ensure data integrity. Although the measured performance decrease that results from character set conversion is insignificant, it is good practice to select the default character set that results in the fewest conversions.

For example, if most of your clients use CP 850, specify CP 850 on your server. You can do this even if your server is on an HP-UX system (where its native character set for the Group 1 languages is ROMAN8).

Note: SAP strongly recommends that you decide which character set to use as your default before you create any databases or make any changes to the SAP-supplied databases.

In the example below, 175 clients all access the same SAP ASE. The clients are on different platforms and use different character sets. The critical factor that allows these clients to function together is that *all* of the character sets in the client/server system belong to the same language group. The default language for the SAP ASE is CP 850, which is the character set used by the largest number of clients. This allows the server to operate most efficiently, with the least amount of character set conversion.

Figure 10: Clients using different character sets in the same language group

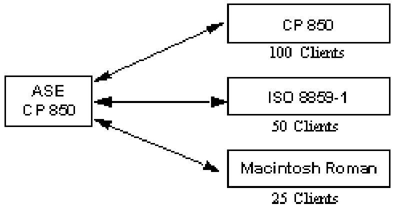

To help you choose the default character set for your server, the following tables list the most commonly used character sets by platform and language.

Table 11. Most Widely Used Western European Client Platforms

Platform	Language	Character set
Win 95, 98	U.S. English, Western Europe	CP 1252
Windows	U.S. English, Western Europe	CP 1252
Win 2000	U.S. English, Western Europe	CP 1252
Sun Solaris	U.S. English, Western Europe	ISO 8859-1
HP-UX 10,11	U.S. English, Western Europe	ROMAN8
IBM AIX 4.x	U.S. English, Western Europe	ISO 8859-1

Table 12. Most Widely Used Japanese Client Platforms

Platform	Language	Character set
Win 95, 98	Japanese	CP 932 for Windows
Win NT 4.0	Japanese	CP 932 for Windows
Win 2000	Japanese	CP 932 for Windows
Sun Solaris	Japanese	EUC-JIS
HP-UX 10,11	Japanese	EUC-JIS
IBM AIX 4.x	Japanese	EUC-JIS

Table 13. Most Widely Used Chinese Client Platforms

Platform	Language	Character set
Win 95, 98	Chinese (simplified)	CP 936 for Windows
Win NT 4.0	Chinese (simplified)	CP 936 for Windows
Win 2000	Chinese (simplified)	CP 936 for Windows
Sun Solaris	Chinese (simplified)	EUC-GB
HP-UX 10,11	Chinese (simplified)	EUC-GBS
IBM AIX 4.x	Chinese (simplified)	EUC-GB

Selecting the Sort Order

Different languages sort the same characters differently. SAP ASE uses sort orders to create indexes, store date into indexed tables, and specify an **order by** clause.

For example, in English, *Cho* is sorted before *Co*, whereas in Spanish, the opposite is true. In German, β is a single character, however in dictionaries it is treated as the double character *ss* and sorted accordingly. Accented characters are sorted in a particular order so that *aménité* comes before *amène*, whereas if you ignored the accents, the reverse would be true. Therefore, language-specific sort orders are required so that characters are sorted correctly.

Each character set comes with one or more sort orders that SAP ASE uses to collate data. A sort order is tied to a particular language or set of languages and to a specific character set. The same sort orders can be used for English, French, and German because they sort the same characters identically, for example, *A*, *a*, *B*, *b*, and so on. Or the characters are specific to one of the languages—for example, the accented characters, *é*, *à*, and *á*, are used in French but not in English or German—and therefore, there is no conflict in how those characters are sorted. The same is not true for Spanish however, where the double letters *ch* and *ll* are sorted differently. Therefore, although the same character sets support all four languages, there is one set of sort orders for English, French and German, and a different set of sort orders for Spanish.

In addition, a sort order is tied to a particular character set. Therefore, there is one set of sort orders for English, French, and German in the ISO 8859-1 character set, another set in the CP 850 character set, and so on. The sort orders available for a particular character set are located in sort order definition files (`*.srt` files) in the character set directory.

Different Types of Sort Orders

All character sets are offered with a binary sort order at a minimum, which blindly sorts all data based only on the arithmetic value of the code assigned to represent each letter (the "binary" code) in the character set.

Binary sort order works well for the first 128 characters of each character set (ASCII English) and for Asian languages.When a character set supports more than one language (for example, Group 1 or Unicode) the binary sort order most likely give incorrect results, and you should select another sort order.

Character sets may also have one or more of the following dictionary sort orders:

- *Dictionary order, case-sensitive, accent-sensitive* – sorts uppercase and lowercase letters separately. Dictionary order recognizes the various accented forms of a letter and sorts them after the associated unaccented letter.
- *Dictionary order, case-insensitive, accent-sensitive* – sorts data in dictionary order but does not recognize case differences. Uppercase letters are equivalent to their lowercase counterparts and are intermingled in sorting results. Useful for avoiding duplicate entries in tables of names.
- *Dictionary order, case-insensitive, accent-sensitive, order with preference* – does not recognize case difference in determining equivalency of items. A word in uppercase is equivalent to the same word in lowercase. Preference is given to uppercase letters (they appear first) if all other conditions are equal.

  Using case-insensitive with preference may cause poor performance in large tables when the columns specified in an **order by** clause match the key of the table's clustered index. Do not select case-insensitive order with preference unless your installation requires that uppercase letters be sorted before lowercase letters in otherwise equivalent strings for **order by** clauses.
- *Dictionary order, case-insensitive, accent-insensitive* – treats accented forms of a letter as equivalent to the associated unaccented letter. It intermingles accented letters in sorting results.

Selecting the Default Sort Order

SAP servers can support only one default sort order at a time. If your users all use the same language, or their languages all use the same sort order, selecting the default sort order is straightforward. .

For example, if your users are using French data and expect French sorting, then you can pick one of the French dictionary sort orders. Or if your users are using data in multiple languages and the languages use the same sort order, for example English, French, and German, you can pick one sort order and it works for all your users in all languages.

However, if you have users using different languages that require different sort orders, for example French and Spanish, then you must select one of the sort orders as the default. If you pick, for example, a French sort order, your Spanish users will not see the *ch* and *ll* double characters sorted as they would expect. The installation procedure, by default, configures the server with the binary sort order.

You can use the **sortkey** function to set up customized alternative sort orders for your data—one for each language.These sort orders can be selected dynamically to meet the needs of different users. The **sortkey** function is separate from the default sort order, but can coexist in the same server. The range and depth of sort orders provided by the **sortkey** function is better than those provided by the default sort order mechanism. For more information, see **sortkey** and **compare** in the *Reference Manual: Building Blocks*.

Table 14. Available sort orders

Language or script	Character sets	Sort orders
All languages	UTF-8	Multiple sort orders
Cyrillic: Bulgarian, Byelorussian, Macedonian, Russian, Serbian, Ukrainian	CP 855, CP 866, CP 1251, ISO 8859-5, Koi8, Macintosh Cyrillic	Dictionary order, case sensitive, accent sensitive
Eastern European: Czech, Slovak	CP 852, ISO 8859-2, CP 1250	Dictionary order, case sensitive, accent sensitive Dictionary order, case insensitive, accent sensitive Dictionary order, case sensitive, accent sensitive, with preference Dictionary order, case insensitive, accent insensitive
English, French, German	ASCII 8, CP 437, CP850, CP 860, CP 863, CP 1252a, ISO 8859-1, ISO 8859-15, Macintosh Roman, ROMAN8, ROMAN9, ISO 15	Dictionary order, case sensitive, accent sensitive Dictionary order, case insensitive, accent sensitive Dictionary order, case sensitive, accent sensitive, with preference Dictionary order, case insensitive, accent insensitive
English, French, German	CP 850, CP 858	Alternate dictionary order, case sensitive Alternate dictionary order, case sensitive, accent insensitive Alternate dictionary order, case sensitive, with preference
Greek	ISO 8859-7	Dictionary order, case sensitive, accent sensitive

Language or script	Character sets	Sort orders
Hungarian	ISO 8859-2	Dictionary order, case sensitive, accent sensitive Dictionary order, case insensitive, accent sensitive Dictionary order, case insensitive, accent insensitive
Japanese	EUCJIS, SJIS, DECKANJI	General purpose case-insensitive dictionary ordering
Kazakh	87	50
Russian	CP 866, CP 1251, ISO 8859-5, Koi8, Macintosh Cyrillic	Dictionary order, case sensitive, accent sensitive Dictionary order, case insensitive, accent sensitive
Scandinavian	CP 850	Dictionary order, case sensitive, accent sensitive Dictionary order, case insensitive, with preference
Simplified Chinese	EUC-GB, GB-18030, CP936	General purpose case-insensitive dictionary ordering
Spanish	ASCII 8, CP 437, CP850, CP 860, CP 863, CP 1252, ISO 8859-1, ISO 8859-15, Macintosh Roman, ROMAN8	Dictionary order, case sensitive, accent sensitive Dictionary order, case insensitive, accent sensitive Dictionary order, case insensitive, accent insensitive
Thai	CP 874, TIS 620	Dictionary order
Turkish	ISO 8859-9	Dictionary order, case sensitive, accent sensitive Dictionary order, case insensitive, accent insensitive Dictionary order, case insensitive, accent sensitive
Western European	CP 1252	Dictionary order, case insensitive, case sensitive, with preference, accent insensitive, Spanish dictionary, Spanish case insensitive, Spanish accent insensitive

If your language does not appear here, there is no language-specific sort order for your language. Select a binary sort order and then investigate whether the **sortkey** function meets your needs. As this table illustrates, many languages have more than one sort order.

Chinese Pinyin Sort Order

Pinyin, more formally known as "Hanyu Pinyin," uses the Roman alphabet to represent the standard Chinese pronunciation system.

Pinyin consists of a system of transliteration to Roman alphabets for reading and writing Mandarin without Chinese characters. Pinyin uses accents to represent the four tones of Mandarin.

Earlier versions of SAP ASE used the Simplified Chinese (GB) sort orders, `gbpinyin` and `gbpinyinocs`, using the Unilib character set, significantly impacting the performance of databases using the GB character sets.

SAP ASE version 15.0.3 automatically uses the `gbpinyin` and `gbpinyinocs` sort orders, eliminating a processing step and significantly improving performance.

In earlier versions, the default **size of unilib cache** configuration parameter was 268 KB. In version 15.0.3, the default has been increased to 302 KB.

Improved performance occurs in queries that access ASCII and `gbpinyin` data. However, if the data set has a mixture of other characters, you may not see any performance improvement.

Selecting Case-Insensitive Sort Orders for Chinese and Japanese Character Sets

Use the **sp_helpsort** and **sp_configure** system procedures to select case-insensitive sort orders.

sp_helpsort lists the available case-insensitive sort orders.

```
sp_helpsort
-----------
Name                   ID
-------------------------
nocase_eucgb           52
nocase_cp936           52
nocase_gb18030         52
nocase_eucjis          52
nocase_sjis            52
nocase_deckanji        52
```

Use **sp_configure** to switch to a case-insensitive sort order:

```
sp_configure 'default sortorder id', 52
```

Selecting the Default Unicode Sort Order

The value for **default unicode sort order** database option is different than the sort order for the server's default character set.

This separate configuration parameter is a static parameter that requires that you restart your server and reindex the `unichar` data if it is changed. This sort order is identified using a string parameter, rather than a numeric parameter, to guarantee that the sort order is unique.

This table lists the available default Unicode sort orders:

Name	ID	Description
defaultml	20	Default Unicode multilingual ordering
thaidict	21	Thai dictionary ordering
iso14651	22	Ordering as per ISO14651 standard
utf8bin	24	Ordering for UTF-16 that matches the UTF-8 binary
binary	25	Binary sort
altnoacc	39	Alternate accent-insensitive
altdict	45	Alternate dictionary ordering
altnocsp	46	Alternate case-insensitive with preference
scandict	47	Scandinavian dictionary ordering
scannocp	48	Scandinavian case-insensitive with preference
bin_utf8	50	UTF-8 binary sort order
dict	51	General-purpose dictionary ordering
nocase	52	General-purpose case-insensitive dictionary ordering
nocasep	53	General-purpose case-insensitive with preference
noaccent	54	Dictionary order, case-insensitive, accent-insensitive
espdict	55	Spanish dictionary ordering
espnocs	56	Spanish case-insensitive dictionary ordering
espnoac	57	Spanish accent-insensitive dictionary ordering
rusnocs	59	Russian case-insensitive dictionary ordering
cyrnocs	64	Cyrillic case-insensitive dictionary ordering
elldict	65	Greek dictionary ordering
hundict	69	Hungarian dictionary ordering
hunnoac	70	Hungarian accent-insensitive dictionary ordering
hunnocs	71	Hungarian case-insensitive dictionary ordering
turknoac	73	Turkish accent-insensitive dictionary ordering

This table lists the loadable sort orders:

Name	ID	Description
cp932bin	129	Ordering that matches the binary ordering of CP932
dynix	130	Chinese phonetic ordering
gb3213bn	137	Ordering that matches the binary ordering of GB2312
cyrdict	140	Common cyrillic dictionary ordering
turdict	155	Turkish Dictionary ordering
euckscbn	161	Ordering that matches the binary ordering of EUCKSC
gbpinyin	163	Chinese phonetic ordering
rusdict	165	Russian dictionary ordering
sjisbin	179	Ordering that matches the binary ordering of SJIS
eucjisbn	192	Ordering that matches the binary ordering of EUCJIS
big5bin	194	Ordering that matches the binary ordering of BIG5

To view this sort order list in SAP ASE, use **sp_helpsort**. See the *Reference Manual: Procedures*.

You can add sort orders using external files in the `$/collate/Unicode` directory. The names and collation IDs are stored in `syscharsets`. The names of external Unicode sort orders do not have to be in `syscharsets` before you can set the default Unicode sort order.

Note: External Unicode sort orders are provided by SAP. Do not attempt to create external Unicode sort orders.

Sort order associated with Unicode data is completely independent of the sort order associated with traditional character data. All relational expressions involving the Unicode datatypes are performed using the Unicode sort order. This includes mixed-mode expressions involving Unicode and non-Unicode data. For example, in the following query the `varchar` character constant ‘Mü’ is implicitly cast to `unichar` and the comparison is performed according to the Unicode sort order:

```
select * from authors where unicode_name > 'Mü'
```

The same holds true for all other comparison operators, as well as the concatenation operator “+”, the operator “in”, and the operator “between.” Once again, the goal is to retain compatibility with existing database applications.

Tables joins based on equality (equijoins) deserve special mention. These are generally optimized by the server to take advantage of indexes that defined on the participating columns. When a `unichar` column is joined with a `char` column, the latter requires a conversion, and since the character sort order and the Unicode sort order are distinct, the optimizer will ignore the index on the `char` column.

In SAP ASE version 12.5.1 and later, when the server's default character set is configured to UTF-8, you can configure the server's default sort order (for `char` data) to be any of the above sort orders. Prior to this version, the binary sort order "bin_utf8" (ID=50) was the only well-behaved sort order for UTF-8. Although not required, the sort order for `char` data in UTF-8 can be selected so that it corresponds with the sort order for `unichar`.

There is a potential confusion regarding choice of binary sort orders for Unicode. The sort order named "binary" is the most efficient one for `unichar` data (UTF-16), and is thus the default. This order is based on the Unicode scalar value, meaning that all 32-bit surrogate pairs are placed after all 16-bit Unicode values. The sort order named "utf8bin" is designed to match the order of the default (most efficient) binary order for UTF-8 `char` data, namely "bin_utf8". The recommended matching combinations are thus "binary" for `unichar` and "binary" for UTF-8 `char`, or "utf8bin" for `unichar` and "bin_utf8" for UTF-8 `char`. The former favors `unichar` efficiency, while the latter favors `char` efficiency. Avoid using "utf8bin" for UTF-8 `char`, since it is equivalent to "bin_utf8" but less efficient.

Select a Language for System Messages

Any installation of SAP ASE can use Language Modules, which contains files of messages in different languages.

SAP ASE provides Language Modules for messages in the following languages: English, Chinese (Simplified), French, German, Japanese, Korean, Brazilian Portuguese, and Spanish. If your client language is *not* one of these languages, you see system messages in English, the default language.

Each client can choose to view messages in their own language at the same time, from the same server; for example, one client views system messages in French, another in Spanish, and another in German. To do this, however, all selected languages must be part of the same language group. For example, French, Spanish and German are all part of language group 1. Japanese, on the other hand, is part of language group 101, which contains no other languages. Therefore, if Japanese is your server language, you can display system messages only in Japanese or English. Remember that all language groups can display messages in English. There is also a server-wide default language, used if the user has not selected a specific language. If you use Unicode, you can view system messages in any of the supported languages.

You can select the language for your system messages in one of two ways:

- Select a language as part of your user profile
- Enter a language in the `locales.dat` file

This table displays the supported system message languages and their language groups. Each user can select only one language per session for system messages:

Language group	System message languages	Character sets
Group 1	French, German, Spanish, Brazilian Portuguese	ASCII 8, CP 437, CP 850, CP 860, CP 863, CP 1252, ISO 8859-1, ISO 8859-15, Macintosh Roman, ROMAN8
Group 2	Polish	Cp 1250, CP 852, ISO 8859-2
Group 101	Japanese	CP 932, DEC Kanji, EUC-JIS, Shift-JIS
Group 102	Simplified Chinese (PRC)	CP 936, EUC-GB, GB18030
Group 104	Korean	EUC-KSC, CP 949
Group 105	Thai	CP 874, TIS 620
Unicode	French, German, Spanish, Brazilian Portuguese, Japanese, Simplified Chinese, Korean	UTF-8
All Other Language Groups	English	

Install Language Modules for all languages in which clients will receive messages. These Language Modules, located in the `locales` subdirectory of the SAP ASE installation directory, are part of a group of files called localization files.

A Spanish-Version Server

This server requires only a Spanish language group.

1. Select the server language, in this case, Spanish, which is part of language group 1. Based on your platform, select a character set from language group 1. SAP recommends that you select the character set used by the greatest number of clients. Or, if you think your company might someday expand into other countries and languages, you might consider installing Unicode.
2. Install the Spanish Language Module in the server. This allows clients to view system messages in Spanish.
3. Select the default sort order. Spanish has three possible sort orders, in addition to binary sort order. Select a sort order.
4. Restart the server.

A US-based company in Japan

This US-based company needs to support both English and Japanese.

1. Select the default character set for your server. If you install a character set from language group 101 (Japanese), you can support both Japanese and English data in the same server.
2. Install the Japanese Language Module so that system messages are available in Japanese.
3. Select the sort order. Because a binary sort order is the only sort order available for Japanese, both the English and Japanese clients have a default binary sort order. Consider using the **sortkey** function to provide solutions for both audiences.
4. Make sure that each Japanese user requests Japanese messages by default. Since you are using a character set from language group 101, and you have already installed the Japanese Language Module, your client in Japan sees messages in Japanese, while clients in the U.S. can choose to see messages in either English or Japanese.

A Japan-Based Company with Multinational Clients

This company is located in Japan, and has clients in France, Germany, and Spain, so you need to mix European and Asian languages in the same server.

1. Select the default server language and character set. Since your company is based in Japan and most of your clients are located in Japan, the default server language should be Japanese. But you also want your clients in France, Germany, and Spain to be able to send and receive data in their native languages. Japanese is part of language group 101, while French, German, and Spanish are part of language group 1. Since the languages you need are not part of the same language group, the only way you can have all of these languages on the same server is to select Unicode as your default character set.
2. Install the Language Modules for Japanese, French, German, and Spanish.
3. Select the binary sort order, since this is the only sort order available for the Unicode character set. (You can, however, consider using the **sortkey** function inside your application code to supply data sorted according to each user's preference.)
4. Select Japanese as the default language for system messages. Clients in other countries can select their own native language for messages.

Changing the character set, sort order, or message language

Even after you have configured your server, a system administrator can change the default character set, sort order, or message language used by SAP ASE.

Because a sort order is built on a specific character set, changing character sets always involves a change in sort order. However, you can change the sort order without changing character sets, because more than one sort order may be available for a character set.

Note: You cannot change the SAP ASE default character set and sort order if it includes an archived database.

To display SAP ASE's default sort order, character set, and a table of its primary sort orders, enter:

```
sp_helpsort
```

Changing the default character set

SAP ASE can have only one default character set, in which data is stored in its databases. When you install SAP ASE, you specify a default character set.

Warning! Read the following carefully, and exercise caution when changing the default character set in SAP ASE. SAP strongly recommends that you perform backups before you change a default character set.

When you change the default character set in SAP ASE, you must convert any existing data to the new default character set. Conversion is unnecessary only if:

- There is no user data in the server.
- It is acceptable to destroy user data in the server.
- You are absolutely certain that data in the server uses only ASCll-7. In this case, you can change the default without first copying your data out of the server.

To change the default character set:

1. Copy the data out using **bcp**.
2. Change the default character set.
3. Use **bcp** with the appropriate flags for data conversion to copy the data back into the server.

Next

See the *Utility Guide* for more information about using **bcp** to copy data.

Warning! After converting data to a different character set (particularly to UTF-8), the data may be too large for the allocated column size. Re-create the columns affected with a larger size.

Code conversion between the character set of the existing data and the new default character set must be supported. If it is not, conversion errors will occur and the data is not converted correctly.

Even if conversions are supported between the character sets, some errors may occur due to minor differences between the character sets, or because some characters do not have equivalents in other character sets. Rows containing problematic data may not get copied back into the database, or data may contain partial or invalid characters.

Changing the sort order with a resources file

SAP ASE character sets can be changed using the resource file.
Use the resource file located in `$SYBASE/ASE-15_0/init/sample_resource_files/`. The resource file looks similar to:

```
sybinit.release_directory: USE_DEFAULT
sybinit.product: sqlsrv
sqlsrv.server_name: PUT_YOUR_SERVER_NAME_HERE
sqlsrv.sa_password: PUT_SA_PASSWORD_HERE
sqlsrv.new_config: yes
sqlsrv.do_add_server: yes
sqlsrv.network_protocol_list: tcp
sqlsrv.network_hostname_list: PUT_YOUR_HOSTNAME_HERE
sqlsrv.network_port_list: PUT_YOUR_PORT_NUMBER_HERE
sqlsrv.application_type: USE_DEFAULT
sqlsrv.server_page_size: USE_DEFAULT
sqlsrv.force_buildmaster: no
sqlsrv.master_device_physical_name:
PUT_THE_PATH_OF_YOUR_MASTER_DEVICE_HERE
sqlsrv.master_device_size: USE_DEFAULT
sqlsrv.master_database_size: USE_DEFAULT
sqlsrv.errorlog: USE_DEFAULT
sqlsrv.do_upgrade: no
sqlsrv.sybsystemprocs_device_physical_name:
PUT_THE_PATH_OF_YOUR_SYBSYSTEMPROCS_DEVICE_HERE
sqlsrv.sybsystemprocs_device_size: USE_DEFAULT
sqlsrv.sybsystemprocs_database_size: USE_DEFAULT
sqlsrv.sybsystemdb_device_physical_name:
PUT_THE_PATH_OF_YOUR_SYBSYSTEMDB_DEVICE_HERE_OR_REMOVE_THIS_LINE
sqlsrv.sybsystemdb_device_size: USE_DEFAULT
sqlsrv.sybsystemdb_database_size: USE_DEFAULT
sqlsrv.tempdb_device_physical_name:
PUT_THE_PATH_OF_YOUR_TEMPDB_DEVICE_HERE_OR_REMOVE_THIS_LINE
sqlsrv.tempdb_device_size: USE_DEFAULT
sqlsrv.tempdb_database_size: USE_DEFAULT
sqlsrv.default_backup_server: PUT_YOUR_BACKUP_SERVER_NAME_HERE
#sqlsrv.addl_cmdline_parameters:
PUT_ANY_ADDITIONAL_COMMAND_LINE_PARAMETERS_HERE
sqlsrv.do_configure_pci: no
```

```
sqlsrv.sybpcidb_device_physical_name:
PUT_THE_PATH_OF_YOUR_SYBPCIDB_DATA_DEVICE_HERE
sqlsrv.sybpcidb_device_size: USE_DEFAULT
sqlsrv.sybpcidb_database_size: USE_DEFAULT
# If sqlsrv.do_optimize_config is set to yes, both
sqlsrv.avail_physical_memory
and sqlsrv.avail_cpu_num need to be set.
sqlsrv.do_optimize_config: no
sqlsrv.avail_physical_memory:
PUT_THE_AVAILABLE_PHYSICAL_MEMORY_FOR_ASE_IN_OPTIMIZATION
sqlsrv.avail_cpu_num:
PUT_THE_AVAILABLE_NUMBER_CPU_FOR_ASE_IN_OPTIMIZATION
~
~
```

Change the Default Sort Order

SAP ASE can have only one default sort order, which is the collating sequence it uses to order data.

When you consider changing the sort order for character data on a particular SAP ASE, keep in mind that all of your organization's SAP ASEs should have the same sort order. A single sort order enforces consistency and makes distributed processing easier to administer.

You may have to rebuild your indexes after changing the default sort order.

Reconfiguring the Character Set, Sort Order, or Message Language

Changing SAP ASE's default character set, sort order, or message language involves extracting the data, installing the new character set, sort order, or message language, shutting down and restarting the server, and reloading the data. For procedures on how to configure the character set, sort order, or message language for a new server, see the configuration documentation for your platform.

Back up all databases in SAP ASE before and after you change character sets or sort orders. After you back up your databases, use **bcp** to copy the data in and out of your databases if:

- A database contains character data and you want to convert the data to a new character set. Do not load a database dump of the data into a server that uses the new default character set. SAP ASE assumes the loaded data is in the new character set, and corrupts the data.
- You are changing the default sort order only and not the default character set. You cannot load a database from a dump performed prior to changing the sort order—if you attempt to, an error message appears, and SAP ASE aborts the load.
- You change the default character set, and either the old or the new sort order is not binary. You cannot load a database dump that was made before you changed the character set.

You cannot reload your data from a database dump once you have reconfigured the default character set and sort order (unless both old and new character sets use a binary sort order and no conversion is required between the old and new character sets).

Example: Converting a Unicode Database to UTF-8

A fictitious database named `xpubs` is modified to use `univarchar` columns.

Prerequisites

Assume a database was created using the following script on a server that has all the installation defaults, namely character set “iso_1” and default sort order ID 50, “binary_iso_1”:

```
create database xpubs
go
use xpubs
go
create table authors (au_id int, au_lname varchar(255), au_fname
varchar(255))
go
create index au_idx on authors(au_lname, au_fname)
go
```

Then the data was loaded into the server using a series of inserts and updates.

To convert the data to UTF-8:

Task

1. Extract the data and convert it to UTF-8 form (the conversion occurs with the **-J** parameter):

   ```
   % bcp xpubs..authors out authors.utf8.bcp -c -Jutf8 -Usa -P
   ```

2. Install UTF-8 as the default character set:

   ```
   charset -Usa -P binary.srt utf8
   isql -Usa -P
   sp_configure 'default sortorder id', 50, 'utf8'
   ```

3. Shutdown the server.
4. Restart the server and modify the default character set and re-create indexes on the system tables. :

   ```
   isql -Usa -P
   sp_dboption xpubs, 'select into', true
   go
   use xpubs
   go
   checkpoint
   go
   delete from authors
   go
   quit
   ```

5. Restart the server.
6. Reload the data:

```
bcp xpubs..authors in authors.utf8.bcp -c -Jutf8 -Usa -P
```

Migrating Selected Columns to unichar

With a working database running with UTF-8 as the default character set, you can convert select columns to `univarchar`.

The columns are modified to the new datatypes, the data is converted in place, and the index is re-created.

For example, to migrate a select of columns from the xpubs database to unichar:

```
% isql -Usa -P
> use xpubs
> go
> alter table authors modify au_lname univarchar(255), au_fname
univarchar(255)
> go
```

Note: Currently, the **alter table modify** command does not support `text`, `image`, or unitext columns. To migrate from a `text` to a unitext column, you must first use **bcp**, create a table with unitext columns, and then use **bcp** again to place data into the new table. This migration path only works when you invoke **bcp** with `–Jutf8` option.

Migrating to or from unitext

The **alter table modify** command does not support `text`, `image`, or unitext columns. To migrate from a `text` to a unitext column, you must first use **bcp**, create a table with unitext columns, and then use **bcp** again to place data into the new table. This migration path only works when you invoke **bcp** with `–Jutf8` option.

Before Changing the Character Set or Sort Order

You must perform some preliminary steps before you change the character set or sort order.

1. Dump all user databases and the `master` database. If you have made changes to `model` or `sybsystemprocs`, dump them also.
2. Load the Language Module if it is not already loaded (see the configuration documentation for your platform for complete instructions).
3. If you are changing the SAP ASE default character set, and your current databases contain non ASCII-7 data, use **bcp** with the necessary flags to copy the existing data out of your databases.

Once you have loaded the Language Module, you can run the SAP ASE installation program, which allows you to:

- Install or remove message languages and character sets included with SAP ASE
- Change the default message language or character set
- Select a different sort order

See the configuration documentation for your platform for instructions on using the installation program

Note: Before you change the character set or sort order, SAP ASE must have as many open databases as there are databases managed by the server. If SAP ASE does not have a sufficient number of open databases when it is re-started after a change in sort order, SAP ASE prints this message to the error log and the server will revert to the former sort order:

```
The configuration parameter 'number of open databases' must be at
least as large as the number of databases,
in order to change the character set or sort order." Re-start
Adaptive Server, use sp_configure to increase
'number of open databases' to at least %d, then re-configure the
character set or sort order
```

To reconfigure the language, character set, or sort order, use the **sqlloc** utility, described in *Utility Guide*. If you are using Windows, use the **Server Config** utility, described in *Configuration Guide > Default SAP ASE Configuration*.

If you installed additional languages but did not change the SAP ASE character set or sort order, you have completed the reconfiguration process.

Set the User's Default Language

If you install an additional language, users running client programs can run **create login** to set that language as their default language, or set the LANG variable on the client machine, with the appropriate entries in **locales.dat**.

Manage Suspect Partitions

Partitions are marked suspect because of a sort order or character set change on a range-partitioned table, or because of a cross-platform dump and load with a hash-partitioned table.

If the table is marked with suspect partitions:

- All updates and cursor activities are suspended on this table.
- No **alter table** commands, except **partition by**, are allowed. **create index** and **drop index** are not allowed on a table with suspect partitions.
- The **select** command is allowed on tables containing suspect partitions. However, the optimizer treats such tables as round-robin partitioned tables, to avoid using the possibly corrupt **partition** condition.

To fix table with suspect partitions:

- If the **partition** condition needs fixing after a sort-order change, you can use **alter table** with the **partition by** option to repartition a table that has suspect partitions.
- If the partition condition does not need fixing, you can use the **reorg rebuild table** command to rebuild the table, redistributing only the data rows among the partitions.

- If the indexes as well as the partitions on a table are marked suspect, use **partition by** or **reorg rebuild** to fix both the suspect index and suspect partitions.

To manage suspect partitions in cross-platform dump and load operations:

- During the first **online database** command, after you execute **load database** across two platforms with different endian types, the hash partition is marked suspect.
- Any global clustered index on a round-robin partition, which has an internally generated partition condition with a `unichar` or `varchar` partition key, is marked suspect.
- After the database is online, use **sp_post_xpload** to fix the suspect partitions and indexes.

Installing Date Strings for Unsupported Languages

Use **sp_addlanguage** to install names for the days of the week and months of the year for languages that do not have language modules.

sp_addlanguage lets you define:

- A language name and (optionally) an alias for the name
- A list of the full names of months and a list of abbreviations for the month names
- A list of the full names of the days of the week
- The date format for entering dates (such as month/day/year)
- The number of the first day of the week

To add the information for Italian:

```
sp_addlanguage italian, italiano,
"gennaio,febbraio,marzo,aprile,maggio,giugno,luglio,agosto,settembr
e,ottobre, novembre,dicembre",
"genn,feb,mar,apr,mag,giu,lug,ago,sett,ott,nov,dic",
"lunedi,martedi,mercoledi,giovedi,venerdi,sabato,domenica", dmy, 1
```

Next

sp_addlanguage enforces strict data entry rules. The lists of month names, month abbreviations, and days of the week must be comma-separated lists with no spaces or line feeds (returns). Also, they must contain the correct number of elements (12 for month strings, 7 for day-of-the-week strings.)

Valid values for the date formats are: `mdy`, `dmy`, `ymd`, `ydm`, `myd`, and `dym`. The `dmy` value indicates that the dates are in day/month/year order. This format affects only data entry; to change output format, you must use the **convert** function.

Server Versus Client Date Interpretation

Generally, date values are resolved on the client. When a user selects date values, SAP ASE sends them to the client in an internal format.

The client uses the `common.loc` file and other localization files in the default language subdirectory of the `locales` directory on the client to convert the internal format to character data. For example, if the user's default language is Spanish, SAP ASE looks for the `common.loc` file in `$SYBASE/$SYBASE_ASE/locales/spanish/char_set`. It uses the information in the file to display, for example, `12 febrero 2007`.

Assume that the user's default language is set to Italian, a language for which SAP ASE does not provide a language module, and that the date values in Italian have been added. When the client connects to the server and looks for the `common.loc` file for Italian, it does not find the file. The client prints an error message and connects to the server. If the user then selects date values, the dates are displayed in U.S. English format.To display the date values added with **sp_addlanguage**, use the **convert** function to force the dates to be converted to character data at the server.

The following query generates a result set with the dates in U.S. English format:

```
select pubdate from titles
```

The query below, however, returns the date with the month names in Italian:

```
select convert(char(19),pubdate) from titles
```

Internationalization and localization files

The files that support data processing in a particular language are called internationalization files. Several types of internationalization files come with SAP ASE.

Table 15. Internationalization files

File	Location	Purpose and contents
`charset.loc`	In each character set subdirectory of the `charsets` directory	Character set definition files that define the lexical properties of each character, such as alphanumeric, punctuation, operand, and uppercase or lowercase. Used by SAP ASE to correctly process data.
`*.srt`	In each character set subdirectory of the `charsets` directory	Defines the sort order for alphanumeric and special characters, including ligatures, diacritics, and other language-specific considerations.
`*.xlt`	In each character set subdirectory of the `charsets` directory	Terminal-specific character translation files for use with utilities such as **bcp** and **isql**. For more information about how the `.xlt` files are used, see the *Utility Guide*.

Warning! Do not alter any of the internationalization files. If you need to install a new terminal definition or sort order, contact your local SAP office or distributor.

Character sets directory structure

The `charsets` directory is located in `$SYBASE/`. There is a separate subdirectory for each character set in the `charsets` directory.

Within the subdirectory for each character set (for example, `cp850`) are the character set and sort order definition files and terminal-specific files.

Figure 11: Structure of the charsets directory

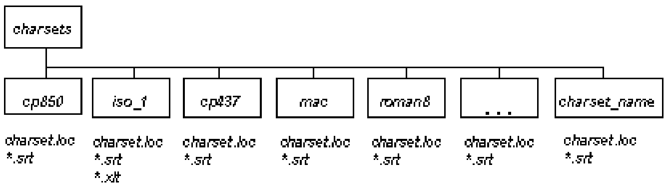

If you load additional character sets, they also appear in the `charsets` directory.

Types of localization files

SAP ASE includes several localization files for each Language Module.

Note: All SAP ASE-related locales files (used by **dataserver**, **sqlloc**, **syconfig**, and so on) are in `$SYBASE/$SYBASE_ASE/locales`, except `locales.dat`, which is in `$SYBASE/locales`. All Open Client/Server-related locales files (**ctlib**, **ctisql**, **ctbcp**, **optdiag**, **installjava**, and so on) are located in `$SYBASE/locales`.

SAP ASE localization files include:

- `locales.dat` – in the `locales` directory. Used by client applications to identify the default message language and character set.
- `server.loc` – in the character set subdirectories under each language subdirectory in the `$SYBASE/$SYBASE_ASE/locales` directory. Software messages translated into the local language. SAP products have product-specific `*.loc` files. If an entry is not translated, the software message or string appears in U.S. English instead of the local language.
- `common.loc` – in each language and character set directory of the `locales` directory. `common.loc` contains the local names of the months of the year and their abbreviations, and information about the local date, time, and money formats.

Warning! Do not alter any of the localization files. If you need to alter any information in those files, contact your local SAP office or distributor.

Software Messages Directory Structure

Within the `$SYBASE/$SYBASE_ASE/locales` directory is a subdirectory for each language installed.

There is always a `us_english` subdirectory (`english` on Windows).)

During installation, when you are prompted to select the languages you want installed on SAP ASE, the installation program lists the supported software message languages. If you install language modules for additional languages, you see subdirectories for those languages. Within each language subdirectory are subdirectories for the supported character sets; for example, `cp850` is a supported character set for `us_english`. Software message files for each SAP product reside in the character set subdirectories.

Global variables for Languages and Character Sets

SAP ASE includes global variables that display information about languages and character sets.

The following global variables contain information about languages:

- *@@langid* – contains the local language ID of the language currently in use (specified in `syslanguages.langid`)
- *@@language* – contains the name of the language currently in use (specified in `syslanguages.name`)

The following global variables contain information about character sets:

- *@@char_convert* – contains 0 if character set conversion is not in effect. Contains 1 if character set conversion is in effect.
- *@@client_csname* – the client's character set name. Set to NULL if client character set has never been initialized; otherwise, contains the name of the character set for the connection.
- *@@client_csid* – the client's character set ID. Set to -1 if client character set has never been initialized; otherwise, contains the client character set ID from `syscharsets` for the connection.
- *@@client_csexpansion* – returns the expansion factor used when converting from server's character set to client's character set.
- *@@maxcharlen* – the maximum length, in bytes, of a character in the SAP ASE default character set.
- *@@ncharsize* – the maximum length, in bytes, of a character set in the current server default character set.
- *@@unicharsize* – equals 2.

See *Reference Manual: Building Blocks > Global Variables* for a list of all global variables.

CHAPTER 11 Configuring Client/Server Character Set Conversions

In a heterogeneous environment, SAP ASE may need to communicate with clients running on different platforms using different character sets.

Although different character sets may support the same language group (for example, ISO 8858-1 and CP 850 support the group 1 languages), they may encode the same characters differently. For example, in ISO 8859-1, the character *à* is encoded as *0xE0* in hexadecimal. However, in CP 850 the same character is encoded as *0x85* in hexadecimal.

To maintain data integrity between your clients and servers, data must be converted between the character sets. The goal is to ensure that an “a” remains an “a” even when crossing between machine and character set boundaries. This process is known as *character set conversion*.

Supported Character Set Conversions

Character set conversion occurs between a pair of character sets. The supported conversions in any particular client/server system depend on the character sets used by the server and its clients. One type of character set conversion occurs if the server uses a native character set as the default; a different type of conversion is used if the server default is Unicode UTF-8.

Conversion for Native Character Sets

SAP ASE supports character set conversion between native character sets that belong to the same language group.

If the server has a native character set as its default, the clients’ character sets must belong to the same language group. In figure below, the clients’ character sets and the SAP ASE default character set all belong to language group 1. Data is correctly converted between the client character sets and the server default character set. Since they all belong to the same language group, the clients can view all data on the server, no matter which client submitted the data.

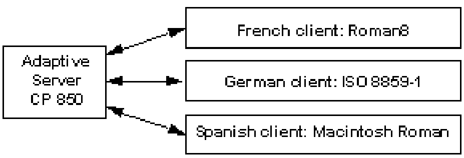

Conversion in a Unicode System

SAP ASE supports character set conversion between UTF-8 and any native character set that SAP supports.

In a Unicode system, since the server default character set is UTF-8, the client character set may be a native character set from any language group. Therefore, a Japanese client (group 101), a French client (group 1), and an Arabic client (group 6) can all send and receive data from the same server. Data from each client is correctly converted as it passes between each client and the server.

Figure 12: Character set conversion in a Unicode system

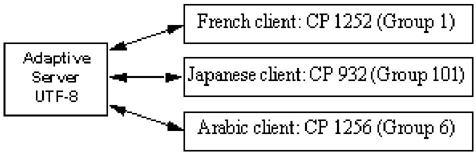

Each client can view data only in the language supported by its character set. Therefore, the Japanese client can view any Japanese data on the server, but it cannot view Arabic or French data. Likewise, the French client can view French or any other Western European language supported by its character set, but not Japanese or Arabic.

Figure 13: Viewing Unicode data

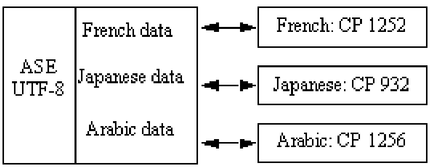

An additional character set, ASCII-7, is a subset of *every* character set, including Unicode, and is therefore compatible with **all** character sets in **all** language groups. If either the SAP ASE or the client's character set is ASCII-7, any 7-bit ASCII character can pass between the client and server unaltered and without conversion.

SAP recommends that you do not configure a server for ASCII-7. You can achieve the same benefits of compatibility by restricting each client to use only the first 128 characters of each native character set.

SAP ASE Direct Conversions

SAP ASE direct conversions occur between two native character sets of the same language group.

For example, SAP ASE supports conversion between CP 437 and CP 850, because both belong to language group 1. SAP ASE direct conversions exist between many, but not all, native character sets of a language group.

Unicode Conversions

Unicode conversions exists for all native character sets. When converting between two native character sets, Unicode conversion uses Unicode as an intermediate character set.

For example, to convert between the server default character set (CP 437), and the client character set (CP 860), CP 437 is first converted to Unicode; Unicode is then converted to CP 860.

Unicode conversions may be used either when the default character set of the server is UTF-8, or a native character set. You must specifically configure your server to use Unicode conversions (unless the server's default character set is UTF-8).

Earlier versions of SAP ASE used direct conversions, and it is the default method for character set conversions. However, Unicode conversions allow easier and less complex character set conversion. While SAP ASE direct conversions are still supported, SAP now also uses Unicode conversions to provide complete conversion support for all character sets and has no plans to add new direct conversions.

Allowing Unicode noncharacters

In versions of SAP ASE earlier than 15.7, the `unichar`, `univarchar`, `unitext`, `char`, `varchar`, and `text` datatypes under the `utf-8` default character set did not accept Unicode noncharacters (code points are permanently reserved for internal use).

SAP ASE versions 15.7 and later allows you to ignore Unicode noncharacters by enabling the **enable functionaliry group** or the **enable permissive unicode** configuration parameters.

If you do not enable this feature, SAP ASE rejects these noncharacters as in earlier versions.

For more information about the Unicode standard, see *the Unicode Consortium Web site*.

When you enable this feature, Unicode noncharacters are not detected in:

- Parameters:
  - Presented as `univarchar` and `unitext` (UTF-16) datatypes
  - Presented as `varchar` and `text` (UTF-8) datatypes
  - As parameters to dynamic SQL statements
  - As input to parameterized language statements
- String literals when the server's character set is UTF-8

- Escaped string literals (those prefixed with U&), regardless of the server's character set
- Conversion processes between `unichar` (UTF-16) and `varchar` (UTF-8) in either direction

In addition, Unicode noncharacters are acceptable in simple expressions such as comparisons, where they sort higher than legal Unicode characters.

You can use Unicode noncharacters as parameters to these functions:

ascii()	**lower()**	**sortkey()**
char_length()	**ltrim()**	**soundex()**
charindex()	**patindex()**	**str_replace()**
compare()	**replicate()**	**stuff()**
datalength()	**reverse()**	**substring()**
difference()	**right()**	**upper()**
left()	**rtrim()**	**uscalar()**
len()		

Note: This feature does not affect UTF-16 surrogate handling enabled with the **enable surrogate handling** configuration parameter. See *System Administration Guide > Setting Configuration Parameters*.

Choosing a Conversion Type

By default, SAP ASE uses direct conversions to convert data between different character sets.

Set the **enable unicode conversions** option to either 1 or 2 to use the Unicode conversions:

- Set to 1 – uses SAP ASE direct conversions or Unicode conversions. SAP ASE first checks to see if an SAP ASE direct conversion exists for the server and client character set. If a direct conversion is used; if no direct conversion exists, the Unicode conversion is used.

  Use this setting if the character sets in your client/server system fall into both columns 1 and 2 in the table below.
- Set to 2 – uses Unicode conversions only. SAP ASE uses Unicode conversions, without attempting to find an SAP ASE direct conversion. Use this setting if the client/server conversions result in a change in the data length.

If all character sets fall into column 2 in the table below, set **enable unicode conversions** to 2 to always use Unicode conversions.

For SAP ASE versions 15.0 and later, the default value for **enable unicode conversions** is 1.

If the server default is UTF-8, the server automatically uses Unicode conversions only.

Non-Unicode Client/Server Systems

In a non-Unicode system, the character sets of the server and clients are native character sets; therefore, you can use the SAP ASE direct conversions.

However, there are some character sets for which there is no SAP ASE direct conversion; in this situation, you must use Unicode conversions.

- If all character sets in your client/server system are column 1 in the table below, use the SAP ASE direct conversions. The character sets must all belong to the same language group.
- If the character sets in your client/server system are in column 2 in the table below, or some combination of columns 1 and 2, configure your server to use Unicode conversions. Again, the character sets must all belong to the same language group.

For example, assume the server default character set is CP 850 and the clients' character sets are either ISO 8859-1 or ROMAN 8. The table below shows that direct conversions exist between CP 850 and the client character sets. Now, suppose you add a client using CP 1252. Since there is no direct conversion between CP 1252 and CP 850, (the default server character set), you must use Unicode conversions to convert between CP 1252 and CP 850. When you have a mixture of character sets—some where you can use SAP ASE direct conversions and others where you must use Unicode conversions—you can specify that a combination of SAP ASE direct conversion and Unicode conversion be used.

Unicode Client/Server Systems

If your server default is Unicode UTF-8, then all conversions are between UTF-8 and the native character set being used on the client systems. In a Unicode system, Unicode conversions are used exclusively.

Table 16. Conversion methods for character sets

Language group	Column 1 – SAP ASE direct conversions and Unicode conversions	Column 2 – Unicode conversions only
Group 1	CP 437, CP 850, ISO 8859-1, Macintosh Roman	CP 860, CP 1252, ISO 8859-15, CP 863
Group 2	CP 852, CP 1250, CP 8859-1, Macintosh Central European	ISO 8859-2
Group 4	No conversions needed (only one character set supported)	
Group 5	CP 855, CP 866, CP 1251, ISO 8859-5, Koi8, Macintosh Cyrillic	
Group 6		CP 864, CP 1256, ISO 8859-6

Language group	Column 1 – SAP ASE direct conversions and Unicode conversions	Column 2 – Unicode conversions only
Group 7	CP 869, CP 1253, GREEK8, ISO 8859-7, Macintosh Greek	
Group 8		CP 1255, ISO 8859-8
Group 9	CP 857, CP 1254, ISO 8859-9, Macintosh Turkish, TURKISH8	
Group 101	DEC Kanjii, EUC-JIS, Shift-JIS	CP 932
Group 102		CP 936, EUG-GB, GB18303
Group 103		Big 5, CP 950, EUC-CNS
Group 104		EUCKSC, CP 949
Group 105		CP 874, TIS 620
Group 106	No conversions needed (only one character set supported)	
Unicode	No conversions needed (only one character set supported)	

Enabling and Disabling Character Set Conversion

A client that requests a connection identifies its character set to SAP ASE. SAP ASE compares the client character set with its default character set, and if the two names are identical, no conversion is required.

If the names differ, SAP ASE determines whether it supports conversion between its default and the client's character set. If it does not, it sends an error message to the client and continues with the login process. If it does, character set conversion is automatically enabled. If the default character set of the server is UTF-8, Unicode conversions are automatically used. If the default is a native character set, the server uses SAP ASE direct conversions, unless the user requests Unicode conversions.

You can disable character set conversion at the server level. You may want to do this if:

- All of your clients are using the same character set as the server default, and therefore, no conversion is required.
- Conversion between the client character set and the server default is not supported.
- You want to store data in the server without changing the encoding.

To disable character set conversion at the server level, set the **disable character set conversion** parameter to 1.

Next

You can control character set conversion at the connection level using the **set char_convert** command from within a client session. **set char_convert off** turns conversion off between a particular client and the server. You may want to **set char_convert off** if the client and the server use the same character set, which makes conversion unnecessary. **set char_convert on** turns conversion back on.

Characters That Cannot Be Converted

You cannot convert all character sets.

Characters may not be converted if:

- The character exists (is encoded) in the source character set, but does not exist in the target character set. For example, the OE ligature is part of the Macintosh character set (code point 0xCE). This character does not exist in the ISO 8859-1 character set. If the OE ligature exists in data that is being converted from the Macintosh to the ISO 8859-1 character set, it causes a conversion error.
- The character exists in both the source and the target character set, but in the target character set, the character is represented by a different number of bytes than in the source character set.

  For example, 1-byte accented characters (such as á, è) are 2-byte characters in UTF-8; 2-byte Thai characters are 3-byte characters in UTF-8. Avoid this limitation by configuring the **enable unicode conversion** option to 1 or 2.

Error handling in character set conversion

The SAP ASE character set conversion reports errors when a character exists in the client's character set but not in the server's character set, or vice versa.

SAP ASE must guarantee that data successfully converted on input to the server can be successfully converted back to the client's character set when the client retrieves that data. To do this effectively, SAP ASE must avoid putting suspect data into the database.

When SAP ASE encounters a conversion error in the data being entered, it generates this message:

```
Msg 2402, Severity 16 (EX_USER):
Error converting client characters into server’s
character set. Some character(s) could not be converted.
```

A conversion error prevents query execution on insert and update statements. If this occurs, review your data for problem characters and replace them.

When SAP ASE encounters a conversion error while sending data to the client, it replaces the bytes of the suspect characters with ASCII question marks (?). The query batch continues to completion. When the statement is complete, SAP ASE sends the following message:

```
Msg 2403, Severity 16 (EX_INFO):
WARNING! Some character(s) could not be converted into
client’s character set. Unconverted bytes were changed
to question marks ('?').
```

Conversions and Changes to Data Lengths

In some cases, converting data between the server’s character set and the client’s character set results in a change to the length of the data. For example, when the character set on one system uses one byte to represent each character and the character set on the other system requires two bytes per character.

When character set conversion results in a change in data length, there are two possibilities:

- The data length decreases, as in the following examples:
  - Greek or Russian in multibyte UTF-8 to a single-byte Greek or Russian character set
  - Japanese two-byte Hankaku Katakana characters in EUC-JIS to single-byte characters in Shift-JIS
- The data length increases, as in the following examples:
  - Single-byte Thai to multibyte Thai in UTF-8
  - Single-byte Japanese characters in Shift-JIS to two-byte Hankaku Katakana in EUC-JIS

To configure your system or application:

1. Configure the server to use Unicode conversions. If the data length increases between the server and the client, you must also complete steps 2 and 3.
2. The client must be using Open Client 11.1 or later. It must inform the server that it can handle CS_LONGCHAR data at connection time, using the Open Client **ct_capability** function.

   The *capability* parameter must be set to CS_DATA_LCHAR and the *value* parameter must be set to CS_TRUE, where *connection* is a pointer to a CS_CONNECTION structure:

   ```
   CS_INT capval = CS_TRUE
   ct_capability(connection,CS_SET,CS_CAP_RESPONS,
         CS_DATA_LCHAR,&capval)
   ```

3. When conversions result in an increase in data length, `char` and `varchar` data are converted to the client's character set and are sent to the client as CS_LONGCHAR data. The client application must be coded to extract the data received as CS_LONGCHAR.

Specify the Character Set for Utility Programs

The SAP utility programs assume that the default character set of the client platform is the same character set the client is using.

However, occasionally the client character set differs from the character set for the platform. For this reason, you may need to specify the client character set at the command line. A command line option for the **isql**, **bcp**, and **defncopy** utilities specifies the client's character set, and temporarily overrides settings of the LANG variable or settings in `locales.dat`.

`-J` *charset_name* (UNIX and PC) sets the client's character set to the *charset_name*.

If you omit the client character set's command line flag, the platform's default character set is used. See the *Utility Guide*.

Display and file character set command line options

You might require character set conversion between the client and a terminal, and between the client and a file system.

This figure illustrates the paths and command line options that are available in the standalone utilities **isql**, **bcp**, and **defncopy**:

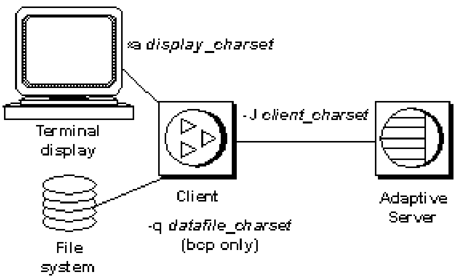

Use:

- `-J` or **/clientcharset** command line option to specify the character set used by the client when it sends and receives character data to and from SAP ASE.
- `-a` command line option if you are running the client from a terminal with a character set that differs from the client character set. In the figure above, the `-a` option and the `-J` option are used together to identify the character set translation file (`.xlt` file) needed for the conversion.

- `-a` without the `-J` parameter only if the client character set is the same as the default character set.
- `-q` command line option if you are running **bcp** to copy character data to or from a file system that uses a character set that differs from the client character set. In the figure above, use the `-q` or**/filecharset** option and the `-J` or **/clientcharset** option together to identify the character set translation file (`.xlt` file) needed for the conversion.

CHAPTER 12 Diagnosing System Problems

Diagnosing system problems involves, among other duties, reviewing error messages, managing processes, performing dumps, starting and stopping servers.

How SAP ASE Uses Error Messages

SAP ASE displays an error message when it encounters a problem.

The error message includes:

- A *message number*, which uniquely identifies the error message
- A *severity level number* between 10 and 24, which indicates the type and severity of the problem
- An *error state number*, which allows unique identification of the line of SAP ASE code at which the error was raised
- An *error message*, which tells you what the problem is, and may suggest how to fix it

See the *Configuration Guide* for your platform for a description of the error log format.

For example, if you try to access a table that does not exist, you see:

```
select * from publisher
```

```
Msg 208, Level 16, State 1:
publisher not found. Specify owner.objectname or use sp_help to check
whether the object exists (sp_help may produce lots of output).
```

There may be more than one error message for a single query. If there is more than one error in a batch or query, SAP ASE usually reports only the first one. Subsequent errors are reported the next time you execute the batch or query.

Error messages are stored in `master..sysmessages`, which is updated with each new version of SAP ASE (and has thousands of rows). Here are the first few rows (from an SAP ASE that uses us_english as the default language):

```
select error, severity, description
from sysmessages
where error >=101 and error <=106
and langid is null
```

```
error severity description
----- -------- --------------------------------------------------
  101       15 Line %d: SQL syntax error.
  102       15 Incorrect syntax near '%.*s'.
  103       15 The %S_MSG that starts with ’%.*s’ is too long.
               Maximum length is %d.
  104       15 Order-by items must appear in the select-list if
```

```
                the statement contains set operators.
 105          15 Unclosed quote before the character string '%.*s'.
 106          16 Too many table names in the query. The maximum
                allowable is %d.
```

You can query `sysmessages` to generate a custom list of error messages:

- If your server supports more than one language, `sysmessages` stores each message in each language. The column `langid` is NULL for us_english and matches the `syslanguages.langid` for other languages installed on the server.
- The `sqlstate` column stores the SQLSTATE value for error conditions and exceptions defined in ANSI SQL92.
- Message numbers 17000 and higher are system procedure error messages and message strings.

Error Log Format

All SAP ASE error messages use the same format.

```
instance_ID:thread_ID:family_ID:spid date time {server | kernel}
message
```

where:

- *instance_ID* – the instance running in a clustered environment. This value is 00 if SAP ASE is running in a nonclustered environment.
- *thread_ID* – the thread on which the error occured. Always apears as a four-digit number (for example, 10 is shown as 0010).

  Note: In process mode, SAP ASE displays the engine involved for each log entry. The engine number is indicated by a four-digit number. If only one engine is online, it is indicated by 00.
- *family_ID* – family ID of the originating thread:
  - In serial processing, you see 00000.
  - In parallel processing, you see the server process ID number of the parent of the originating thread.
- *spid* – the server process ID of the originating thread:
  - In serial processing, this is the server process ID number of the thread that generated the message. If the thread is a system task, the spid value is 00000.
  - In parallel processing, you see the server process ID number of the originating thread.
- *date* – appeared in the format `yyyy/mm/dd`, which allows you to sort error messages by date.
- *time* – shown in 24-hour format, which includes seconds and hundredths of a second.
- *server* | *kernel* – this entry is for use only by SAP Technical Support.
- *message* – the error message

For example:

```
00:0024:00000:00000:2010/04/27 10:28:07.82 kernel Thread 24 (LWP
24390) of Threadpool syb_default_pool online as engine 0
```

- *instance_ID* = 0 – the server is not configured for a clustered environment
- *thread_ID* = 0024 – SAP ASE brought engine 0 online as thread number 24.
- *family_ID* = 00000 – the process is running in serial mode.
- *spid* = 00000 – this is a system task.
- *date* and *time* = 2010/04/27 10:28:07.82.
- *server* | *kernel* = kernel.

Error Messages and Message Numbers

A combination of message number (`error`) and language ID (`langid`) uniquely identifies each error message. Messages that share the same message number but have different language IDs.

```
select error, description, langid
from sysmessages
where error = 101
```

```
error description                            langid
----- ------------------------------------- ------
  101 Line %d: SQL syntax error.               NULL
  101 Ligne %1!: erreur de syntaxe SQL.           1
  101 Zeile %1!: SQL Syntaxfehler.                2

(3 rows affected)
```

The error message text describes the problem. The descriptions often include a line number, a reference to a type of database object (a table, column, stored procedure, and so forth), or the name of a particular database object.

In the `description` field of `sysmessages`, a percent sign (%) followed by a character or character string serves as a placeholder for these pieces of data, which SAP ASE supplies when it encounters the problem and generates the error message. “%d” is a placeholder for a number; “%S_MSG” is a placeholder for a type of database object; “%.*s”—all within quotes—is a placeholder for the name of a particular database object.

For example, the `description` field for message number 103 is:

```
The %S_MSG that starts with '%.*s' is too long. Maximum length is
%d.
```

The actual error message that appears to a user might be:

```
The column that starts with 'title' is too long. Maximum length is
80.
```

For errors that you report to Technical Support, include the numbers, object types, and object names.

Variables in Error Message Text

Error messages include variables to stand for characters, numbers, database structures, and so on.

Symbol	Stands for
%d, %D	Decimal number
%x,%X,%.*x,%lx, %04x, %08lx	Hexadecimal number
%s	Null-terminated string
%.*s, %*s, %*.s	String, usually the name of a particular database object
%S_*type*	SAP ASE-defined structure
%c	Single character
%f	Floating-point number
%ld	Long decimal
%lf	Double floating-point number

SAP ASE error logging

Error messages from SAP ASE are sent to the user's screen and to the error log file.

The stack trace from fatal error messages (severity levels 19 and higher) and error messages from the kernel are sent to an error log file. The name of this file varies; see the configuration documentation for your platform or the *Utility Guide*.

Note: The error log file is owned by the user who installed SAP ASE (or the person who started SAP ASE after an error log was removed). Permissions or ownership problems with the error log at the operating system level can block successful start-up of SAP ASE.

SAP ASE creates an error log for you if one does not already exist. Specify the location of the error log at start-up with the *errorlogfile* parameter in the runserver file or at the command line. The SAP installer utility configures the runserver file with `$SYBASE/$SYBASE_ASE/install` as the location of the error log if you do not choose an alternate location. If you do not specify the location in the runserver file or at the command line, the location of the error log is the directory from which you start SAP ASE. For more information about specifying the location of the error log, see **dataserver** in the *Utility Guide*.

Note: Always start SAP ASE from the same directory, or with the runserver file or the error log flag, so that you can locate your error log.

Each time you start a server, messages in the error log provide information on the success (or failure) of the start and the recovery of each database on the server. Subsequent fatal error messages and all kernel error messages are appended to the error log file. To reduce the size of the error log by deleting old or unneeded messages, “prune” the log while SAP ASE is shut down.

Reporting Errors

When you report an error to SAP Technical Support, include:

- The message number, level number, and state number.
- Any numbers, database object types, or database object names that are included in the error message.
- The context in which the message was generated, that is, the command that was running at the time. You can help by providing a hard copy of the error log.

Severity Levels

The severity level of a message indicates the type and severity of the problem that SAP ASE has encountered.

For maximum integrity, when SAP ASE responds to error conditions, it displays messages from `sysmessages`, but takes action according to an internal table. A few corresponding messages differ in severity levels, so you may occasionally notice a difference in expected behavior if you are developing applications or procedures that refer to SAP ASE messages and severity levels.

Warning! You can create your own error numbers and messages based on SAP ASE error numbers (for example, by adding 20,000 to the SAP ASE value). However, you cannot alter the SAP ASE-supplied system messages in the `sysmessages` system table.

You can add user-defined error messages to `sysusermessages` with **sp_addmessage**. See the *Reference Manual: Procedures*.

Users should inform the system administrator whenever problems that generate severity levels of 17 and higher occur. The system administrator is responsible for resolving them and tracking their frequency.

If the problem has affected an entire database, the system administrator may have to use the database consistency checker (**dbcc**) to determine the extent of the damage. The **dbcc** may identify some objects that have to be removed. It can repair some damage, but you may have to reload the database.

For more information, see the following chapters in the *System Administration Guide: Volume 2*:

- *Checking Database Consistency* – **dbcc**
- *Backing Upa and Restoring User Databases* – loading a user database
- *Restoring the System Databases* – loading system databases

Severity Levels 10 – 18

Error messages with severity levels 10–16 are generated by problems that are caused by user errors. These problems can be corrected by the user. Severity levels 17 and 18 do not terminate the user's session.

Error messages with severity levels 17 and higher should be reported to the system administrator or database owner.

Level 10: Status Information

Messages with severity level 10 provide additional information after certain commands have been executed and, typically, do not display the message number or severity level.

For example, after a **create database** command, SAP ASE displays a message telling the user how much of the requested space has been allocated for the new database.

Level 11: Specified Database Object Not Found

Messages with severity level 11 indicate that SAP ASE cannot find an object that is referenced in a command.

This is often because the user has made a mistake in typing the name of a database object, because the user did not specify the object owner's name, or because of confusion about which database is current. Check the spelling of object names, use the owner names if the object is not owned by the user or "dbo," and make sure you are in the correct database.

Level 12: Wrong Datatype Encountered

Messages with severity level 12 indicate a problem with datatypes. For example, the user may have tried to enter a value of the wrong datatype in a column or to compare columns of different and incompatible datatypes.

To correct comparison problems, use the **convert** function with **select**. See the *Reference Manual: Building Blocks* or the *Transact-SQL Users Guide*.

Level 13: User Transaction Syntax Error

Messages with severity level 13 indicate that something is wrong with the current user-defined transaction.

For example, the user may have issued a **commit transaction** command without having issued a **begin transaction**, or they may have tried to roll back a transaction to a savepoint that has not been defined (sometimes there may be a typing or spelling mistake in the name of the savepoint).

Severity level 13 can also indicate a deadlock, in which case the deadlock victim's process is rolled back. The user must restart his or her command.

Level 14: Insufficient Permission to Execute Command

Messages with severity level 14 mean that the user does not have the necessary permission to execute the command or access the database object. Users can ask the owner of the database object, the owner of the database, or the system administrator to grant them permission to use the command or object in question.

Level 15: Syntax Error in SQL Statement

Messages with severity level 15 indicate that the user has made a mistake in the syntax of the command. The text of these error messages includes the line numbers on which the mistake occurs and the specific word near which it occurs.

Level 16: Miscellaneous User Error

Most error messages with severity level 16 reflect that the user has made a nonfatal mistake that does not fall into any of the other categories. Severity level 16 and higher might also indicate software or hardware errors.

For example, the user may have tried to update a view in a way that violates the restrictions. Another error that falls into this category is unqualified column names in a command that includes more than one table with that column name. SAP ASE has no way to determine which one the user intends. Check the command syntax and working database context.

Messages that ordinarily have severities greater than 16 show severity 16 when they are raised by **dbcc checktable** or **dbcc checkalloc** so that checks can continue to the next object. When you are running the **dbcc** utility, check the *Error Messages and Troubleshooting Guide* for information about error messages between 2500 and 2599 with a severity level of 16.

Note: Levels 17 and 18 are usually not reported in the error log. Users should be instructed to notify the system administrator when level 17 and 18 errors occur.

Level 17: Insufficient Resources

Error messages with severity level 17 mean that the command has caused SAP ASE to run out of resources or to exceed some limit set by the system administrator. The user can continue, although he or she might not be able to execute a particular command

These system limits include the number of databases that can be open at the same time and the number of connections allowed to SAP ASE. They are stored in system tables and can be checked with **sp_configure**.

The database owner can correct the level 17 error messages indicating that the user has run out of space. Other level 17 error messages should be corrected by the system administrator.

Level 18: Nonfatal Internal Error Detected

Error messages with severity level 18 indicate an internal software bug. However, the command runs to completion, and the connection to SAP ASE is maintained.

The user can continue with the work they are doing, although they may not be able to execute a particular command. An example of a situation that generates severity level 18 is SAP ASE detecting that a decision about the access path for a particular query has been made without a valid reason.

Since problems that generate such messages do not keep users from their work, users tend not to report them. However, users should be instructed to inform the system administrator every time an error message with this severity level (or higher) occurs so that the system administrator can report them.

Severity Levels 19 – 26

Fatal problems generate error messages with severity levels 19 and higher. They break the user's connection to SAP ASE (some of the higher severity levels shut down SAP ASE). To continue working, the user must restart the client program.

When a fatal error occurs, the process freezes its state before it stops, recording information about what has happened. The process is then killed and disappears.

When the user's connection is broken, he or she may or may not be able to reconnect and resume working. Some problems with severity levels in this range affect only one user and one process. Others affect all the processes in the database. In some cases, the system administrator must restart SAP ASE. These problems do not necessarily damage a database or its objects, but they can. They may also result from earlier damage to a database or its objects. Other problems are caused by hardware malfunctions.

Error messages from the kernel are directed to the error log file.

Level 19: SAP ASE Fatal Error in Resource

Error messages with severity level 19 indicate that some nonconfigurable internal limit has been exceeded and that SAP ASE cannot recover gracefully. You must reconnect to SAP ASE.

Level 20: SAP ASE Fatal Error in Current Process

Error messages with severity level 20 indicate that SAP ASE has encountered a bug in a command. The problem has affected only the current process, and the database is unlikely to have been damaged. Run **dbcc** diagnostics. The user must reconnect to SAP ASE.

Level 21: SAP ASE Fatal Error in Database Processes

Error messages with severity level 21 indicate that SAP ASE has encountered a bug that affects all the processes in the current database. However, it is unlikely that the database itself has been damaged. Restart SAP ASE and run **dbcc** diagnostics. The user must reconnect to SAP ASE.

Level 22: SAP ASE Fatal Error: Table Integrity Suspect

Error messages with severity level 22 indicate that the table or index specified in the message has been previously damaged by a software or hardware problem.

The first step is to restart SAP ASE and run **dbcc** to determine whether other objects in the database are also damaged. Whatever the report from **dbcc** may be, The problem may be only in the cache, and not on the disk itself. If so, restarting SAP ASE fixes the problem.

If restarting does not help, then the problem is on the disk as well. Sometimes, the problem can be solved by dropping the object specified in the error message. For example, if the message tells you that SAP ASE has found a row with length 0 in a nonclustered index, the table owner can drop the index and re-create it.

SAP ASE takes any pages or indexes offline that it finds to be suspect during recovery. Use **sp_setsuspect_granularity** to determine whether recovery marks an entire database or only individual pages as suspect. See **sp_setsuspect_granularity** in the *Reference Manual: Procedures*.

The user must reconnect to SAP ASE.

Level 23: Fatal Error: Database Integrity Suspect

Error messages with severity level 23 indicate that the integrity of the entire database is suspect due to previous damage caused by a software or hardware problem. Restart SAP ASE and run **dbcc** diagnostics.

Even when a level 23 error indicates that the entire database is suspect, the damage may be confined to the cache, and the disk itself may be fine. If so, restarting SAP ASE with **startserver** fixes the problem.

Level 24: Hardware Error or System Table Corruption

Error messages with severity level 24 reflect a media failure or (in rare cases) the corruption of `sysusages`. The system administrator may have to reload the database. You may need to call your hardware vendor.

Level 25: SAP ASE Internal Error

Users do not see level 25 errors, as this level is used only for SAP ASE internal errors.

Level 26: Rule Error

Error messages with severity level 26 reflect that an internal locking or synchronization rule has been broken. You must shut down and restart SAP ASE.

Backup Server Error Logging

Like SAP ASE, Backup Server creates an error log if one does not already exist.

Specify the location of the error log at start-up with the *error_log_file* parameter in the runserver file or at the command line. The SAP installer configures the runserver file with `$SYBASE/install` as the location of the error log if you do not choose an alternate location during installation. If you do not specify the location in the runserver file or at the command line, the location of the error log is the directory from which you start Backup Server. Use the **backupserver -V** option (**bcksvr -V** on Windows) to limit the messages printed to the error log. See the sections describing Backup Server in the *Utility Guide*.

Backup Server error messages are in this form:

```
MMM DD YYY: Backup Server:N.N.N.N: Message Text
```

Backup Server message numbers consist of four integers separated by periods, in the form N.N.N.N. Messages in the form N.N.N are sent by Open Server.

The four components of a Backup Server error message are *major.minor.severity.state*:

- The *major* component generally indicates the functional area of the Backup Server code where the error occurred:
  - 1 – system errors.
  - 2 – Open Server event errors.
  - 3 – Backup Server remote procedure call errors.
  - 4 – I/O service layer errors.
  - 5 – network data transfer errors.
  - 6 – volume-handling errors.
  - 7 – option-parsing errors.

Major error categories 1– 6 may result from Backup Server internal errors or a variety of system problems. Major errors in category 7 are almost always due to problems in the options you specified in your dump or load command.

- *minor* numbers are assigned in order within a major category.
- *severity* is:
  - 1 – informational, no user action necessary.
  - 2, 3 – an unexpected condition, possibly fatal to the session, has occurred. The error may have occurred with usage, environment, or internal logic, or any combination of these factors.
  - 4 – an unexpected condition, fatal to the execution of the Backup Server, has occurred. The Backup Server must exit immediately.
- *state* codes have a one-to-one mapping to instances of the error report within the code. If you need to contact Technical Support about Backup Server errors, the state code helps determine the exact cause of the error.

Killing Processes

A process is a unit of execution carried out by SAP ASE. The **kill** command removes ongoing processes.

Each process is assigned a unique process identification number when it starts. This number is called a `spid`. These numbers are stored, along with other information about each process, in `master..sysprocesses`. Processes running in a parallel-processes environment create child processes, each of which has its own `spid`s. Several processes create and assign `spids`: starting SAP ASE, login tasks, checkpoints, the housekeeper tasks, and so on. You can see most of the information by running **sp_who**.

Running **sp_who** on a single-engine server shows the **sp_who** process running and all other processes that are “runnable” or in one of the sleep states. In multi-engine servers, there can be a process running for each engine.

The most frequent reason for killing a process is that it interferes with other users, and the person responsible for running it is not available. The process may hold locks that block access to database objects, or there may be many sleeping processes occupying the available user connections. A system administrator can kill most running or "runnable" processes, including those that are waiting for:

- An alarm, such as a **waitfor** command
- Network sends or receives
- A lock
- Synchronization messages from another process in a family

SAP ASE allows you to kill processes only if it can cleanly roll back any uncompleted transactions and release all system resources that are used by the process. For processes that are part of a family, killing any of the child processes also kills all other processes in the family.

However, it is easiest to kill the parent process. For a family of processes, the **kill** command is detected more quickly if the status of the child processes is `sync sleep`.

This table shows the status values that **sp_who** reports and when the **kill** command takes effect:

Status	Indicates	Effects of kill command
`recv sleep`	Waiting on a network read.	Immediate.
`send sleep`	Waiting on a network send.	Immediate.
`alarm sleep`	Waiting on an alarm such as: `waitfor delay "10:00"`	Immediate.
`lock sleep`	Waiting on a lock acquisition.	Immediate.
`sync sleep`	Waiting on a synchronization message from another process in the family.	Immediate. Other processes in the family must also be brought to state in which they can be killed.
`sleeping`	Waiting on a disk I/O, or some other resource. Probably indicates a process that is running, but doing extensive disk I/O	Killed when it "wakes up," usually immediate; a few sleeping processes do not wake up and require a server restart to clear.
`runnable`	In the queue of runnable processes.	Immediate.
`running`	Actively running on one of the server engines.	Immediate.
`infected`	Server has detected serious error condition; extremely rare.	**kill** command not recommended. Server restart probably required to clear process.
`background`	A process, such as a threshold procedure, run by SAP ASE rather than by a user process.	Immediate; use **kill** with extreme care. Recommend a careful check of `sysprocesses` before killing a background process.
`log suspend`	Processes suspended by reaching the last-chance threshold on the log.	Immediate.

Only system administrators can issue the **kill** command; permission to use it cannot be transferred.

The syntax is:

```
kill spid
```

You can kill only one process at a time, but you can perform a series of kill commands in a batch. For example:

```
1> kill 7
2> kill 8
3> kill 9
4> go
```

A **kill** command is irreversible and cannot be included in a user-defined transaction. `spid` must be a numeric constant; you cannot use a variable. Here is some sample output from **sp_who**:

```
sp_who
```

```
fid  spid  status      loginame  origname  hostname        blk_spid
  dbname
        tempdbname    cmd                block_xloid   threadpool
---  ----  ----------  --------  --------  ----------------  --------
-  ------
        ----------    -----------------  -----------   -------------
------
0       1  recv sleep    howard    howard          svr30eng
 0  master
            tempdb    AWAITING COMMAND
 0      syb_default_pool
0       2  sleeping        NULL      NULL
 0  master
            tempdb    NETWORK HANDLER             0      syb_defaul
t_pool
0       3  sleeping        NULL      NULL
 0  master
          tempdb   DEADLOCK TUNE              0     syb_default_pool
0       4  sleeping        NULL      NULL
 0  master
          tempdb   MIRROR HANDLER             0     syb_default_pool
0       5  sleeping        NULL      NULL
 0  master
            tempdb    CHECKPOINT SLEEP            0      syb_defaul
t_pool
0       6  sleeping        NULL      NULL
 0  master
          tempdb   HOUSEKEEPER                0     syb_default_pool
0       7  recv sleep      bill      bill           bigblue
 0  master
            tempdb    AWAITING COMMAND
 0      syb_default_pool
0       8  recv sleep    wilbur    wilbur             hazel
 0  master
            tempdb    AWAITING COMMAND
 0      syb_default_pool
0       9  recv sleep      joan      joan          luv2work
 0  master
            tempdb    AWAITING COMMAND
 0      syb_default_pool
0      10  running        foote     foote          svr47hum
 0  master
            tempdb    SELECT                      0      syb_defaul
t_master
(10 rows affected, return status = 0)
```

In the example above, processes 2–6 cannot be killed: they are system processes. The login name NULL and the lack of a host name identify processes them as system processes. NETWORK HANDLER, MIRROR HANDLER, HOUSEKEEPER, and CHECKPOINT SLEEP (or, rarely, CHECKPOINT) always appear in **sp_who** output. AUDIT PROCESS appears if auditing is available.

Processes 1, 8, 9, and 10 can be killed, since they have the status values “recv sleep,” “send sleep,” “alarm sleep,” and “lock sleep.”

In **sp_who** output, you cannot tell whether a is “recv sleep” belongs to a user who is using SAP ASE and may be pausing to examine the results of a command, or whether a user has restarted a PC or other terminal, and left a stranded process. Query the `sysprocesses` table to learn more about questionable processes. For example, this query shows the host process ID and client software used by process 8:

```
select hostprocess, program_name
    from sysprocesses
where spid = 8
```

```
hostprocess program_name
----------- ----------------
3993        isql
```

This query, plus the information about the user and host from the **sp_who** results, provides additional information for tracking down the process from the operating system level.

Using kill with statusonly

The **kill ...statusonly** command reports on the progress of a server process ID (spid) in rollback status, but does not terminate the spid.

The **statusonly** report displays:

- The percent of rollback completed
- The estimated length of time, in seconds, before the rollback completes
- Information about the amount of log space scanned, and to be scanned, for a full transaction rollback
- The different types of rollbacks, and reports a subset of information in cases where the rollback is not a full transaction rollback

To track the progress of a rollback, you must run **kill...with statusonly** multiple times:

```
kill spid with statusonly
```

Where *spid* is the number of the process you are terminating.

This example reports on the process of the rollback of spid number 13:

```
kill 13 with statusonly
```

```
spid: 13 Transaction rollback in progress. Estimated rollback
completion: 17% Estimated time left: 13 seconds
```

If the rollback of the spid has completed when you issue **kill...statusonly** or if SAP ASE cannot roll back the specified spid, **kill...statusonly** returns the following message:

```
Status report cannot be obtained. KILL spid:nn is not in progress.
```

Using sp_lock to Examine Blocking Processes

In addition to **sp_who**, **sp_lock** can help identify processes that are blocking other processes.

If the `blk_spid` column in the **sp_who** report indicates that another process has been blocked while waiting to acquire locks, **sp_lock** can display information about the blocking process.

For example, say process 10 in the **sp_who** output above is blocked by process 7. To see information about process 7, execute:

```
sp_lock 7
```

For more information about locking in SAP ASE, see the *Performance and Tuning Series: Locking and Concurrency Control*.

Using Shared Memory Dumps

You can configure SAP ASE to dump shared memory to files when specific conditions occur (for example, segmentation faults).

The shared memory dump file provides a snapshot of the server at the time the specified condition occurred, eliminating the difficulty of using a debugger or specialized instrumented code to trap an error at the moment it occurs. Technical support engineers can use this dump file to diagnose the cause of the shared memory dump, and determine how to configure Adaptiver Server to avoid the problem in the future.

Configuring Shared Memory Dump Conditions

Use **sp_shmdumpconfig** to define or modify the shared memory dump conditions and system-wide attributes of those conditions.

The syntax for **sp_shmdumpconfig** is:

```
sp_shmdumpconfig "action", type, value, max_dumps, dump_dir,
dump_file, option1, option2, option3, option4, option5
```

Note: sp_shmdumpconfig uses positional parameters. When setting a parameter that falls to the right of parameters you do not want to set, specify null values for the unset parameters.

sp_shmdumpconfig allows you to:

- Create new dump conditions
- Modify existing dump conditions
- Drop existing dump conditions
- Reset the dump count for a dump condition

- Display the current shared memory dump configuration
- Set and modify system wide attributes for a dump condition.

Set these parameters to use **sp_shmdumpconfig**:

- **dump on conditions** – determines whether SAP ASE generates a dump of data in shared memory
- **maximum dump conditions** – sets the maximum number of dump conditions that can be configured
  See *System Administration Guide: Volume 1 > Setting Configuration Parameters*.

See the *Reference Manual: Procedures*.

This example describes a typical shared memory dump:

1. Enable SAP ASE to generate a shared memory dump:
   ```
   sp_configure "dump on conditions", 1
   ```
2. Specify the dump condition using **sp_shmdumpconfig**. This example request a one-time memory dump on a signal 11:
   ```
   sp_shmdumpconfig "add", signal, 11, 1 ,"/sybase"
   ```
3. You can turn off collection by deleting the dump condition after collecting the desired data. For example, to drop the condition for error 631 and disable shared memory dumps:
   ```
   sp_shmdumpconfig "drop", "error", 631
   ```
4. Disable shared memory dumps on SAP ASE:
   ```
   sp_configure "dump on conditions", 0
   ```

System-Wide Default Settings

Each dump condition can inherit attributes from the system-wide default settings. Specifying an attribute value for a condition overrides the system-wide defaults. To update the system-wide default settings, use the **sp_shmdumpconfig** stored procedure with a **type** value of *defaults*.

It is typically more convenient to specify any attribute values in the system-wide default settings, and to allow all dump conditions to inherit these values.

SAP recommends that all dump conditions have an explicit dump directory (the *dump_dir* parameter). To ensure this, specify a dump directory in the system-wide defaults and allow all dump conditions to inherit this value.

The default attribute values for the system-wide settings are:

```
sp_shmdumpconfig
```

```
Configured Shared Memory Dump Conditions
----------------------------------------

Defaults   ---
Maximum Dumps:          1
```

```
Halt Engines:            Halt
Cluster:                 Local
Page Cache:              Omit
Procedure Cache:         Include
Unused Space:            Omit
Dump Directory:          $SYBASE
Dump File Name:          Generated File Name
Estimated File Size:     101 MB

Current number of conditions: 0
Maximum number of conditions: 10

Configurable Shared Memory Dump Configuration Settings
------------------------------------------------------
Dump on conditions: 0

You must run 'sp_configure 'dump on conditions',1' to
enable the shared memory dump facility.

Number of dump threads: 1
Include errorlog in dump file: 1
Merge parallel files after dump: 1

Server Memory Allocation
Procedure Cache Data Caches Server Memory Total Memory
-------------- ----------- ------------- -----------
         16 MB        9 MB         86 MB      109 MB
```

Housekeeper functionality

The housekeeper task provides important background tasks.

These tasks include:

- The housekeeper wash, housekeeper garbage collection, and housekeeper chores tasks. **sp_who** displays all three tasks in its output as `HK WASH`, `HK GC`, and `HK CHORES`:

```
sp_who
```

```
fid    spid   status       loginame     origname     hostname
blk_spid  dbname  tempdbname  cmd        block_xloid  threadpool
----  -----  ----------    ----------    ---------   ----------
--------  -------  ---------  -------------  --------- ----------
---
   0      8    sleeping         NULL         NULL     luv2work
       0   master     tempdb        HK
WASH         0  syb_default_pool
   0      9    sleeping         NULL         NULL         NULL
     0   master    tempdb        HK GC        0  syb_default_pool
   0     10    sleeping         NULL         NULL         NULL
       0   master     tempdb      HK
CHORES         0  syb_default_pool
```

- The general automatic restart of housekeeper-related system tasks: you need not restart the server if these system tasks quit unexpectedly.

A system administrator can change all housekeeper task priorities.

sp_showpsexe, as well as **sp_who**, recognizes all three housekeeper names.

For more information about **sp_who** and **sp_showpsexe**, see the *Reference Manual: Procedures*.

Housekeeper wash

Washing buffers is an optional task that, if enabled, runs only during idle times.

You can turn off this task using the configuration parameter **housekeeper free write percent**. The housekeeper wash task is the only housekeeper task for which you use this configuration parameter.

Housekeeper chores

The housekeeper chores task runs only at idle times, and does not use a common configuration parameter.

It manages miscellaneous chores, such as:

- Flushing table statistics.
- Flushing account statistics.
- Handling timeout of detached transactions. You can turn off this task using the configuration parameter **dtm detach timeout period**.
- Checking licence usage. You can turn off this task using the configuration parameter **license information**.

Housekeeper Garbage Collection

There are two forms of garbage collection, lazy and aggressive.

These terms describe two distinct tests for finding empty pages:

- Lazy garbage collection refers to an inexpensive test to find empty pages. This test may not be effective during long-running transactions, and empty pages may accumulate. Lazy garbage collection is inexpensive to use, but can lower performance, which is affected by the fragmentation of allocated table space, and by the accumulation of empty pages that must be evaluated during queries.
- Aggressive garbage collection refers to a sophisticated test for empty pages. This test is more expensive than the lazy garbage collection test, because it checks each deleted row in a page to determine whether the deleted transactions are committed.

Use the **enable housekeeper GC** configuration parameter to configure the **delete** command and the housekeeper garbage collection task for aggressive or lazy garbage collection.

The aggressive housekeeper garbage collection self-tunes the frequency with which the housekeeper garbage collection task examines the housekeeper list, so that the frequency of examination matches the rate at which the application generates empty pages.

Configure SAP ASE Priority Level

The housekeeper garbage collection task typically operates at the priority level of an ordinary user, competing for CPU time with ordinary user tasks, preventing the list of empty pages from growing faster than the housekeeper can delete them.

However, if SAP ASE is configured for threaded mode, use the **sp_bindexeclass 'sv'** object type to change the server-wide priority setting for the housekeeper wash task. Set the priority to the `EC1`, `EC2, EC3` level or define a new user-created execution class. This example sets the housekeeper wash task to the highest, EC1, priority:

```
sp_bindexeclass 'HK WASH', 'sv', NULL, 'EC1'
```

Use **sp_setpsexe** to set the priority level for the housekeeper task for the session if SAP ASE is configured for threaded or process mode.

This example sets the priority level for the housekeeper wash task (with a spid of 8) to `HIGH` for the current session:

```
sp_setpsexe 8, 'priority', 'HIGH'
```

See the *Reference Manual: Procedures*.

Configuring enable housekeeper GC

You can configure the level for SAP ASE garbage collection.

Configure SAP ASE for garbage collection task using:

```
sp_configure "enable housekeeper GC", value
```

The valid values for the **enable housekeeper GC** configuration parameter are:

- 0 – disables the housekeeper garbage collection task, but enables lazy garbage collection by the **delete** command. You must use `reorg reclaim_space` to deallocate empty pages. This is the cheapest option with the lowest performance impact, but it may cause performance problems if many empty pages accumulate. SAP recommends that you do not use this value.
- 1 – enables lazy garbage collection, by both the housekeeper garbage collection task and the **delete** command. This is the default value. If more empty pages accumulate than your application allows, consider options 4 or 5. You can use the **optdiag** utility to obtain statistics of empty pages.
- 2 – reserved for future use.
- 3 – reserved for future use.
- 4 – enables aggressive garbage collection for both the housekeeper garbage collection task and the **delete** command. This option is the most effective, but the **delete** command is the

most expensive. This option is ideal if the deletes on your dataonly locked tables are in a batch.

- 5 – enables aggressive garbage collection for the housekeeper, and lazy garbage collection by **delete**. This option is less expensive for deletes than option 4. This option is suitable when deletes are caused by concurrent transactions.

For example, enter:

```
sp_configure "enable housekeeper GC", 4
```

Using The reorg Command

Garbage collection is most effective when you set **enable housekeeper GC** to 4 or 5.

SAP recommends that you set the parameter value to 5. However, if performance considerations prevent setting this parameter to 4 or 5, and you have an accumulation of empty pages, run **reorg** on the affected tables. You can obtain statistics on empty pages through the **optdiag** utility.

When the server is shut down or crashes, requests to deallocate pages that the housekeeper garbage collection task has not yet serviced are lost. These pages, empty but not deallocated by the housekeeper garbage collection task, remain allocated until you remove them by running **reorg**.

See *System Administration Guide: Volume 2 > Using the reorg Command*.

Shutting Down Servers

A system administrator can shut down SAP ASE or Backup Server using the **shutdown** command.

The syntax is:

```
shutdown [backup_server_name] [with {wait|nowait}]
```

The default for the **shutdown** command is **with wait**. That is, **shutdown** and **shutdown with wait** do exactly the same thing.

Shutting Down SAP ASE

shutdown minimizes the amount of work that automatic recovery must do when you restart SAP ASE.

If you do not provide a server name, **shutdown** shuts down the SAP ASE you are using. When you issue a **shutdown** command, SAP ASE:

1. Disables logins, except for system administrators
2. Performs a checkpoint in each database, flushing pages that have changed from memory to disk
3. Waits for currently executing SQL statements or procedures to finish

The **with nowait** option shuts down SAP ASE immediately. User processes are aborted, and recovery may take longer after a **shutdown with nowait**. You can help minimize recovery time by issuing a **checkpoint** command before you issue a **shutdown with nowait** command.

Shutting down a Backup Server

Include the Backup Server name when you using the **shutdown** command.

The default is **with wait**, so any dumps or loads in progress complete before the Backup Server process halts. After you issue a **shutdown** command, no new dump or load sessions can be started on the Backup Server.

To see the names of the Backup Servers that are accessible from your SAP ASE, execute **sp_helpserver**. Use the value in the `name` column in the **shutdown** command. You can shut down a Backup Server only if it is:

- Listed in `sysservers` on your SAP ASE, and
- Listed in your local `interfaces` file.

Use **sp_addserver** to add a Backup Server to `sysservers`.

1. Run run **sp_who** to check for active dumps and load currenly running on your Backup Server:

```
SYB_BACKUP...sp_who

spid    status     loginame   hostname     blk      cmd
-----   --------   --------   ----------   ---     --------------
    1   sleeping       NULL         NULL     0      CONNECT HANDLER
    2   sleeping      NULL         NULL     0      DEFERRED HANDLER
    3   runnable       NULL          NULL     0       SCHEDULER
    4   runnable       NULL          NULL     0       SITE HANDLER
    5   running           sa   heliotrope     0       NULL
```

2. Shut down a Backup Server:

```
shutdown SYB_BACKUP
```

Using nowait on a Backup Server

The **shutdown** ***backup_server*** **with nowait** command shuts down the Backup Server, regardless of current activity. Use it only in severe circumstances. It can leave your dumps or loads in incomplete or inconsistent states.

If you use **shutdown with nowait** during a log or database dump, check for the message indicating that the dump completed. If you did not receive this message, or if you are not sure whether the dump completed, your next dump should be a **dump database**, not a transaction dump. This guarantees that you are not relying on possibly inconsistent dumps.

If you use **shutdown with nowait** during a load of any kind, and you did not receive the message indicating that the load completed, you may not be able to issue further **load transaction** commands on the database. Run a full database consistency check (**dbcc**) on the database before you use it. You may have to reissue the full set of load commands, starting with **load database**.

Learning about known problems

The release bulletin is a valuable resource for learning about known problems or incompatibilities with SAP ASE and Backup Server. Reading the release bulletin in advance can save you the time and guesswork of troubleshooting known problems.