



REFERENCE MANUAL | PUBLIC

SAP Adaptive Server Enterprise 16.0 SP03

Document Version: 1.0 – 2020-03-04

Reference Manual: Configuration Parameters

Content

- 1 Configuration Parameters. 13**
- 1.1 abstract plan cache. 13
- 1.2 abstract plan dump. 14
- 1.3 abstract plan dynamic replace. 14
- 1.4 abstract plan load. 15
- 1.5 abstract plan replace. 16
- 1.6 abstract plan sharing. 16
- 1.7 additional network memory. 17
- 1.8 aggressive task stealing. 19
- 1.9 allocate max shared memory. 19
- 1.10 allow backward scans. 20
- 1.11 allow kerberos null password. 21
- 1.12 allow memory grow at startup. 21
- 1.13 allow nested triggers. 22
- 1.14 allow procedure grouping. 23
- 1.15 allow remote access. 23
- 1.16 allow resource limits. 24
- 1.17 allow sendmsg. 24
- 1.18 allow sql server async i/o. 25
- 1.19 allow statement rollback. 25
- 1.20 allow updates to system tables. 26
- 1.21 async poll timeout. 27
- 1.22 average cap size. 27
- 1.23 audit queue size. 28
- 1.24 auditing. 29
- 1.25 automatic cluster takeover. 29
- 1.26 automatic master key access. 30
- 1.27 autotune rule history depth. 30
- 1.28 bind backupserver address. 31
- 1.29 builtin date strings. 31
- 1.30 caps per ccb. 32
- 1.31 capture compression statistics. 32
- 1.32 capture missing statistics. 33
- 1.33 check password for digit. 33
- 1.34 cipc large message pool size. 34
- 1.35 cipc regular message pool size. 34

1.36	cis bulk insert array size.	35
1.37	cis bulk insert batch size.	35
1.38	cis connect timeout.	36
1.39	cis cursor rows.	36
1.40	cis idle connection timeout.	37
1.41	cis packet size.	37
1.42	cis pushdown for HANA.	38
1.43	cis rpc handling.	38
1.44	cluster heartbeat interval.	39
1.45	cluster heartbeat retries.	40
1.46	cluster redundancy level.	40
1.47	cluster vote timeout.	41
1.48	column default cache size.	42
1.49	compression info pool size.	42
1.50	compression memory size.	43
1.51	configuration file.	43
1.52	cost of a logical io.	44
1.53	cost of a physical io.	44
1.54	cost of a cpu unit.	45
1.55	cpu accounting flush interval.	46
1.56	cpu grace time.	47
1.57	current audit table.	47
1.58	deadlock checking period.	49
1.59	deadlock pipe active.	50
1.60	deadlock pipe max messages.	50
1.61	deadlock retries.	51
1.62	decompression row threshold.	52
1.63	default character set id.	52
1.64	default database size.	53
1.65	default exp_row_size percent.	53
1.66	default fill factor percent.	54
1.67	default language id.	55
1.68	default network packet size.	55
	Requesting a Larger Packet Size at Login.	56
1.69	default sortorder id.	57
1.70	default unicode sortorder.	57
1.71	default xml sortorder.	58
1.72	deferred name resolution.	58
1.73	disable character set conversions.	59
1.74	disable disk mirroring.	59
1.75	disable jsagent core dump.	60

1.76	disable varbinary truncation.	60
1.77	disk i/o structures.	61
1.78	dma object pool size.	61
1.79	dtm detach timeout period.	62
1.80	dtm lock timeout period.	62
1.81	dump history filename.	64
1.82	dump on conditions.	64
1.83	dynamic allocation on demand.	65
1.84	dynamic sql plan pinning.	65
1.85	early row send increment.	66
1.86	enable async database init.	67
1.87	enable backupserver ha.	67
1.88	enable buffered io for load.	68
1.89	enable bulk inserts.	68
1.90	enable cis.	69
1.91	enable compression.	69
1.92	enable concurrent dump tran.	70
1.93	enable console logging.	71
1.94	enable deferred parallel.	71
1.95	enable delta dump tran.	72
1.96	enable dtm.	72
1.97	enforce dump configuration.	73
1.98	enable dump history.	73
1.99	enable encrypted columns.	74
1.100	enable file access.	74
1.101	enable fips.	75
1.102	enable full-text search.	75
1.103	enable functionality group.	76
1.104	enable granular permissions.	77
1.105	enable ha.	77
1.106	enable HCB index.	78
1.107	enable housekeeper gc.	78
1.108	enable hp posix async i/o.	80
1.109	enable hugepages.	81
1.110	enable i/o fencing.	81
1.111	enable inline default sharing.	82
1.112	enable in-memory row storage.	82
1.113	enable ISM (Solaris only).	83
1.114	enable java.	84
1.115	enable job scheduler.	85
1.116	enable js restart logging.	85

1.117	enable large chunk etc.	86
1.118	enable large pool for load.	86
1.119	enable ldap user auth.	87
1.120	enable LFB index.	87
1.121	enable lightweight rvm.	88
1.122	enable literal autoparam.	88
1.123	enable lock remastering.	89
1.124	enable logins during recovery.	89
1.125	enable mem scale.	90
1.126	enable merge join.	90
1.127	enable metrics capture.	91
1.128	enable monitoring.	92
1.129	enable pam user auth.	92
1.130	enable pci.	93
1.131	enable permissive unicode.	93
1.132	enable predicated privileges.	94
1.133	enable query tuning mem limit.	95
1.134	enable query tuning time limit.	95
1.135	enable rapidlog.	96
1.136	enable rapidtimer.	96
1.137	enable real time messaging.	97
1.138	enable rep agent threads.	97
1.139	enable resolve as owner.	98
1.140	enable row level access control.	98
1.141	enable select into in tran.	99
1.142	enable semantic partitioning.	99
1.143	enable sort-merge join and jtc.	100
1.144	enable sql debugger.	100
1.145	enable ssl.	101
1.146	enable sticky statistics.	101
1.147	enable stmt cache monitoring.	102
1.148	enable streamlined parallel.	102
1.149	enable surrogate processing.	103
1.150	enable transactional memory.	103
1.151	enable unicode conversion.	104
1.152	enable unicode normalization.	104
1.153	enable utility lvl 0 scan wait.	105
1.154	enable webservice.	106
1.155	enable workload analyzer.	106
1.156	enable xact coordination.	107
1.157	enable xml.	108

1.158	engine local cache percent.	108
1.159	engine memory log size.	109
1.160	errorlog pipe active.	109
1.161	errorlog pipe max messages.	110
1.162	errorlog size.	111
1.163	esp execution priority.	111
1.164	esp execution stacksize.	112
1.165	esp unload dll.	112
1.166	event buffers per engine.	113
1.167	event log computer name (Windows Only).	113
1.168	event logging (Windows only).	114
1.169	executable codesize + overhead.	114
1.170	extended cache size.	115
1.171	extend implicit conversion.	115
1.172	external keystore.	116
1.173	fips login password encryption.	117
1.174	FM enabled.	117
1.175	global async prefetch limit.	118
1.176	global cache partition number.	118
1.177	HADR connect timeout.	119
1.178	HADR distinct server name.	119
1.179	HADR login stall time.	120
1.180	HADR mode.	120
1.181	HADR primary check frequency.	121
1.182	HADR primary wait time.	121
1.183	HADR remote query timeout.	122
1.184	hash table hit rate threshold.	122
1.185	HCB index auto tuning.	123
1.186	HCB index memory pool size.	123
1.187	HCB index tuning interval.	124
1.188	heap memory per user.	125
1.189	histogram tuning factor.	125
1.190	housekeeper free write percent.	126
1.191	i/o accounting flush interval.	128
1.192	i/o batch size.	128
1.193	i/o polling process count.	129
1.194	identity burning set factor.	130
1.195	identity grab size.	131
1.196	identity reservation size.	132
1.197	idle migration timeout.	132
1.198	imrs cache utilization.	133

1.199	inline table functions.	134
1.200	job scheduler interval.	135
1.201	job scheduler memory.	135
1.202	job scheduler tasks.	136
1.203	js heartbeat interval.	137
1.204	js job output width.	137
1.205	js restart delay.	138
1.206	kernel mode.	138
1.207	kernel resource memory.	139
1.208	large allocation auto tune.	140
1.209	LFB memory size.	140
1.210	license information.	141
1.211	lock address spinlock ratio.	141
1.212	lock hashtable size.	142
1.213	lock scheme.	143
1.214	lock shared memory.	143
1.215	lock spinlock ratio.	144
1.216	lock table spinlock ratio.	145
1.217	lock timeout pipe active.	145
1.218	lock timeout pipe max messages.	146
1.219	lock wait period.	147
1.220	log audit logon failure.	147
1.221	log audit logon success.	148
1.222	max async i/os per engine.	148
	On the Linux Platform.	149
1.223	max async i/os per server.	149
1.224	max buffers per lava operator.	151
	Differences Between number of sort buffers and max data buffers per lava sort.	152
1.225	max cis remote connections.	152
1.226	max concurrently recovered db.	153
1.227	max js restart attempts.	154
1.228	max memory.	154
	If SAP ASE Cannot Start.	155
1.229	max native threads per engine.	155
1.230	max nesting level.	156
1.231	max network packet size.	156
	Choosing Packet Sizes.	157
1.232	max network peek depth (UNIX Only).	158
1.233	max number network listeners.	159
1.234	max number of IN elements.	160
1.235	max online engines.	160

1.236	max online q engines.	161
1.237	max parallel degree.	162
1.238	max pci slots.	163
1.239	max query parallel degree.	164
1.240	max repartition degree.	164
1.241	max resource granularity.	165
1.242	max scan parallel degree.	165
1.243	max sql text monitored.	166
1.244	max transfer history.	167
1.245	max utility parallel degree.	168
1.246	maximum dump conditions.	168
1.247	maximum failed logins.	169
1.248	maximum job output.	169
1.249	memory alignment boundary.	170
1.250	memory dump compression level.	170
1.251	memory per worker process.	171
1.252	messaging memory.	171
1.253	metrics elap max.	172
1.254	metrics exec max.	172
1.255	metrics lio max.	173
1.256	metrics pio max.	173
1.257	min pages for parallel scan.	174
1.258	minimum password length.	174
1.259	mnc_full_index_filter.	175
1.260	msg confidentiality reqd.	175
1.261	msg integrity reqd.	176
1.262	net password encryption reqd.	177
1.263	network polling mode.	178
	In-Line Network Polling.	178
1.264	nonpushdown pipe active.	179
1.265	nonpushdown pipe max messages.	179
1.266	number of alarms.	180
1.267	number of aux scan descriptors.	181
	Monitoring and Estimating Scan Descriptor Usage.	182
1.268	number of backup connections.	183
1.269	number of ccbs.	184
1.270	number of checkpoint tasks.	184
1.271	number of devices.	185
	Setting the number of devices on UNIX.	185
1.272	number of disk tasks.	186
1.273	number of dtx participants.	187

Optimizing the Number of DTX Participants for Your System.	188
1.274 number of dump threads.	188
1.275 number of early send rows.	189
1.276 number of engines at startup.	190
1.277 number of hcb gc tasks per db.	191
1.278 number of histogram steps.	191
1.279 number of imrs gc tasks per db.	192
1.280 number of index trips.	192
1.281 number of java sockets.	193
1.282 number of large i/o buffers.	194
1.283 number of lob gc tasks per db.	194
1.284 number of locks.	195
1.285 number of mailboxes.	196
1.286 number of messages.	196
1.287 number of network tasks.	197
1.288 number of oam trips.	197
1.289 number of open databases.	198
Optimizing the number of open databases.	199
1.290 number of open indexes.	200
Optimizing number of open indexes.	200
1.291 number of open objects.	201
Optimizing number of open objects.	202
1.292 number of open partitions.	203
Optimizing the number of open partitions Parameter for Your System.	203
1.293 number of pack tasks per db.	204
1.294 number of pre-allocated extents.	205
Using a Value of 32 for the number of pre-allocated extents.	205
1.295 number of q engines at startup.	206
1.296 number of reexecutions.	207
1.297 number of remote connections.	207
1.298 number of remote logins.	208
1.299 number of remote sites.	208
1.300 number of sort buffers.	209
1.301 number of user connections.	210
Upper Limit to the Maximum Number of User Connections.	210
Reserved Connections.	211
Optimizing the Maximum Number of User Connections.	211
1.302 number of worker processes.	212
1.303 NVCache Lazy Cleaner Pool Size.	213
1.304 o/s file descriptors.	213
1.305 object lockwait timing.	214

1.306	open index hash spinlock ratio.	214
1.307	open index spinlock ratio.	215
1.308	open object spinlock ratio.	215
1.309	optimization goal.	216
1.310	optimize temp table resolution.	217
1.311	optimization timeout limit.	217
1.312	optimize dump for fast load.	218
1.313	optimizer level.	218
1.314	page lock promotion hwm.	219
1.315	page lock promotion lwm.	220
1.316	page lock promotion pct.	221
1.317	page utilization percent.	222
1.318	partition groups.	223
1.319	partition spinlock ratio.	224
1.320	pci memory size.	224
1.321	per object statistics active.	225
1.322	percent database for history.	225
1.323	percent database for output.	226
1.324	percent history free.	226
1.325	percent output free.	227
1.326	performance monitoring option.	227
1.327	permission cache entries.	228
1.328	physical lock cushion.	229
1.329	plan text pipe active.	229
1.330	plan text pipe max messages.	230
1.331	point query rate threshold.	230
1.332	prevent automatic upgrade.	231
1.333	print deadlock information.	232
1.334	print recovery information.	232
1.335	procedure cache size.	233
1.336	procedure deferred compilation.	234
1.337	process wait events.	234
1.338	prod-consumer overlap factor.	235
1.339	quorum heartbeat interval.	235
1.340	quorum heartbeat retries.	236
1.341	quoted identifier enhancements.	236
1.342	rapidlog buffer size.	237
1.343	rapidlog max files.	238
1.344	read committed with lock.	238
1.345	recovery interval in minutes.	239
1.346	recovery prefetch size.	240

1.347	remote server pre-read packets.	241
1.348	replication agent memory size.	241
1.349	restore database options.	242
1.350	restricted decrypt permission.	243
1.351	row lock promotion hwm.	244
1.352	row lock promotion lwm.	245
1.353	row lock promotion pct.	246
1.354	rtm thread idle wait period.	247
1.355	rules file.	247
1.356	runnable process search count.	247
	runnable process search count versus idle timeout.	248
1.357	sampling percent.	249
1.358	scavenge temp objects.	249
1.359	secure default login.	250
1.360	select for update.	251
1.361	select on syscomments.text.	252
1.362	send doneinproc tokens.	252
1.363	session migration timeout.	253
1.364	session tempdb log cache size.	253
1.365	shared memory starting address.	254
1.366	show deferred compilation text.	254
1.367	sigstack csmd min size.	255
1.368	sigstack min size.	256
1.369	simplified native access plan.	256
1.370	size of auto identity column.	257
1.371	size of global fixed heap.	258
1.372	size of process object heap.	259
1.373	size of shared class heap.	259
1.374	size of unilib cache.	260
1.375	solaris async i/o mode (Solaris Only).	261
1.376	sproc optimize timeout limit.	261
1.377	sql batch capture.	262
1.378	sql perfmon integration (Windows Only).	263
1.379	sql server clock tick length.	263
1.380	sql text pipe active.	264
1.381	sql text pipe max messages.	265
1.382	stack guard size.	265
1.383	stack size.	267
1.384	start xp server during boot.	268
1.385	startup delay.	268
1.386	statement cache size.	269

1.387	statement pipe active.	269
1.388	statement pipe max messages.	270
1.389	statement statistics active.	270
1.390	streamlined dynamic sql.	271
1.391	strict dtm enforcement.	272
1.392	suppress js max task message.	272
1.393	suspend audit when device full.	273
1.394	syb_sendmsg port number.	274
1.395	sysstatistics flush interval.	275
1.396	systemwide password expiration.	275
1.397	tape retention in days.	276
1.398	tcp no delay.	277
1.399	text prefetch size.	277
1.400	threshold event max messages.	278
1.401	threshold event monitoring.	278
1.402	time slice.	279
1.403	total data cache size.	279
1.404	total logical memory.	280
1.405	total physical memory.	281
1.406	transfer utility memory size.	282
1.407	txn to pss ratio.	282
	Optimizing the txn to pss ratio for your system.	283
1.408	unified login required.	284
1.409	update statistics hashing.	284
1.410	upgrade version.	285
1.411	use security services.	286
1.412	user log cache queue size.	286
1.413	user log cache size.	287
1.414	user log cache spinlock ratio.	288
	utility statistics hashing.	289
1.415	wait event timing.	289
1.416	wait on uncommitted insert.	290
1.417	workload manager cache size.	293
1.418	xact coordination interval.	293
1.419	xp_cmdshell context.	295

1 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.

For information about setting configuration parameters, or permission requirements, see *System Administration Guide: Volume 1 > Setting Configuration Parameters*.

1.1 abstract plan cache

Enables caching of abstract plan hash keys.

Summary	Description
Default value	0 (off)
Valid 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.

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

Related Information

[abstract plan load \[page 15\]](#)

1.2 abstract plan dump

Enables the saving of abstract plans to the `ap_stdout` abstract plans group.

Summary	Description
Default value	0 (off)
Valid 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*.

1.3 abstract plan dynamic replace

The abstract plan dynamic replace configuration parameter will ignore the plan hint in the plan clause, and force the query to use the plan hint in `sysqueryplans`.

The `abstract plan dynamic replace` configuration parameter supports the use of parameterized SQL text.

For example, `sysqueryplans` may have a plan hint `'(i_scan ind tab)'` that is used for for query `select count(*) from tab`.

In the following example, `'(t_scan tab)'` is the abstract plan hint in plan clause:

```
select count(*) from tab plan '( t_scan tab )'
```

When `abstract plan dynamic replace` is not enabled, the query uses the `'(t_scan tab)'` plant hin in plan clause.

When `abstract plan dynamic replace` is enabled, the query uses the `'(i_scan ind tab)'` plan hint in `sysqueryplans`, thereby ignoring the plant hint in plan clause, and forcing the query to use the plan hint in `sysqueryplans`.

See *Performance and Tuning Series: Query Processing and Abstract Plans > Introduction to Abstract Plans > Parameterized SQL Text* for details on how to use `abstract plan dynamic replace` with parameterized SQL text.

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

`abstract plan replace` is ignored when `abstract plan dynamic replace` is enabled.

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

Related Information

[abstract plan replace \[page 16\]](#)

[abstract plan load \[page 15\]](#)

1.4 abstract plan load

Enables association of queries with abstract plans in the `ap_stdin` abstract plans group.

Summary	Description
Default value	0 (off)
Valid 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*.

Related Information

[abstract plan replace \[page 16\]](#)

[abstract plan dynamic replace \[page 14\]](#)

1.5 abstract plan replace

Enables plan replacement for abstract plans in the `ap_stdout` abstract plans group.

Summary	Description
Default value	0 (off)
Valid 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 replace mode to take effect.

`abstract plan replace` is ignored when `abstract plan dynamic replace` is enabled.

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

Related Information

[abstract plan load \[page 15\]](#)

[abstract plan dynamic replace \[page 14\]](#)

1.6 abstract plan sharing

Enables abstract plan sharing between different users.

Summary	Description
Default value	0 (off)

Summary	Description
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.

1.7 additional network memory

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

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Memory Use, Network Communication, Physical Memory

Usage

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`.

i 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	
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.

Related Information

[default network packet size \[page 55\]](#)

1.8 aggressive task stealing

Sets the SAP ASE scheduler task stealing policy to aggressive.

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

1.9 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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 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.

Related Information

[max memory \[page 154\]](#)

1.10 allow backward scans

Controls how the optimizer performs `select` queries that contain the `order by . . . desc` command.

Summary	Description
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*.

Related Information

[print deadlock information \[page 232\]](#)

1.11 allow kerberos null password

Enables Kerberos administrators to pass the caller's password as null so that they can reset the login password with the `alter login` command.

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

Note

Instead of resetting the login password to null, `allow kerberos null password` allows Kerberos administrators to pass the caller password as null so that they can reset login passwords.

1.12 allow memory grow at startup

Determines if the value for `max memory` increases when the server starts.

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

When the server starts, enabling `allow memory grow at startup` allows the server to automatically increase the value for `max memory`, if needed, and prints the following message to the error log:

```
The current 'max memory' value '%d' is less than the 'total logical memory' value '%d'. ASE has increased the value of configuration parameter 'max memory' to '%d' to enable boot.
```

If `allow memory grow at startup` is disabled, the server does not increase the value for `max memory`. If the amount of memory is insufficient, the start up fails with error message 5899.

1.13 allow nested triggers

Controls the use of nested triggers.

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

i Note

The multiple triggers feature does not change the behavior of the `allow nested triggers` configuration parameter.

1.14 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	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System security officer
Configuration group	Security Related

1.15 allow remote access

Controls logins from remote SAP ASE servers.

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

The default value of 1 allows SAP ASE to communicate with Backup Server.

i 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.

1.16 allow resource limits

Controls the use of resource limits.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.17 allow sendmsg

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

Summary	Description
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`.

i Note

Sending messages to UDP ports is not supported on Windows.

1.18 allow sql server async i/o

Enables SAP ASE to run with asynchronous disk I/O.

Summary	Description
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.

1.19 allow statement rollback

Enables the server to perform a rollback, even when the query includes a `convert` function (which can cause errors based on arithmetic overflow).

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

1.20 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	Description
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

⚠ Caution

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*.

i 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.

1.21 async poll timeout

When running in threaded kernel mode, a task may hang on the function `kctCheckAsync()`. The task can be killed using the `async poll timeout` configuration parameter.

⚠ Caution

You can also use the `kill` command to kill the task when the trace flag 7866 is enabled. Using these methods forcefully kills the remote connection that is hanging.

Summary	Description
Default value	0 (disabled)
Valid values	0 to 2147483647 (in minutes)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

The `async poll timeout` configuration parameter forcefully closes the remote server connection. The configuration parameter can be configured in minutes. The value should not be set to a very small value, otherwise a working connection might be forcefully terminated. The connection will be closed irrespective of any pending results or not.

To determine if the task which is hanging is the `kctCheckAsync()` function, use `dbcc stacktrace (<spid>)`.

1.22 average cap size

Reserved for future use.

Summary	Description
Default value	200
Valid values	100 – 10000
Status	Static
Display level	Diagnostics

Summary	Description
Required role	
Configuration group	Diagnostics

1.23 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	Description
Default value	100
Valid values	1 – 65535
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Memory Use, Security Related

When you configure the queue size, 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.

1.24 auditing

Enables or disables auditing for SAP ASE.

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

1.25 automatic cluster takeover

Determines whether an instance automatically recovers from a cluster failure.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
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 `automatic 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.

1.26 automatic master key access

Allows SAP ASE to read the key encryption keys from the master key start-up file.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Security Related

SAP ASE creates a master key start-up file if it does not already exist, but does not write the key encryption key values to the file until the `automatic_startup` key copies either of the master, or until the dual master keys are created.

When `automatic master key access` is disabled, SAP ASE drops the key encryption keys for master and dual master keys from the server memory. SAP ASE does not erase the key encryption key values from the master key start-up file.

1.27 autotune rule history depth

Determines the amount of history SAP ASE retains about autotuning. SAP ASE internal. Do not change.

Summary	Description
Default value	10

Summary	Description
Valid values	0 – 255
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

1.28 bind backupserver address

Binds the Backup Server listener address to the connection socket when performing a remote dump or load.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Backup/Recovery

1.29 builtin date strings

Determines whether 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic

Summary	Description
Display level	
Required role	
Configuration group	Query Tuning

Values are:

- 0 - if a string is given as an argument to a chronological system function, the server interprets it as a datetime value regardless of its apparent precision. This is the default.
- 1 - makes the server interpret the argument string as bigdatetime. This affects the result of chronological system functions.

You can also change the values using the `set` option `builtin_date_strings`.

1.30 caps per ccb

Reserved for future use.

Summary	Description
Default value	50
Valid values	5 – 50
Status	Static
Display level	
Required role	
Configuration group	Diagnostics

1.31 capture compression statistics

Enables the `monTableCompression` monitoring table to begin capturing compression statistics.

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

Summary	Description
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	SQL Server Administration

1.32 capture missing statistics

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

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

i Note

Enabling `capture missing statistics` can result in poor performance and should be used in a production environment for no more than a short period of time. This functionality is intended to be used in a test environment for debugging.

1.33 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	Description
Default value	0 (off)

Summary	Description
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	10
Required role	System security office
Configuration group	Security Related

1.34 cipc large message pool size

Specifies the number of large message buffers allocated by CIPC at start-up time.

Summary	Description
Default value	512
Valid values	512 – 2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

1.35 cipc regular message pool size

Specifies the number of regular message buffers allocated by CIPC at start-up time.

Summary	Description
Default value	8192
Valid values	2048 – 2147483647
Status	Static
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	Shared Disk Cluster

1.36 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	Description
Default value	50
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

1.37 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	Description
Default value	0
Valid 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.

1.38 `cis connect timeout`

Determines the wait time, in seconds, for a successful Client-Library connection.

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

1.39 `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	Description
Default value	50
Valid values	1 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

1.40 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	Description
Default value	0
Valid 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.

1.41 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	Description
Default value	512
Valid 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`, `untext`, 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 network memory
- max 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`.

Related Information

[additional network memory \[page 17\]](#)

[max network packet size \[page 156\]](#)

1.42 cis pushdown for HANA

(SAP HANA accelerator for SAP ASE only) Enables pushdowns for SQL functions and temporary tables to a SAP HANA server.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

See the [SAP HANA accelerator for SAP ASE](#) documentation for more information.

1.43 cis rpc handling

Specifies the default method for remote procedural call (RPC) handling.

Summary	Description
Default value	0 (off); 1 (on for the Cluster Edition)
Valid values	0 (off), 1 (on)

Summary	Description
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.
- 1 forces RPC handling to use Component Integration Service access methods.

See `set_cis_rpc_handling` in the *Component Integration Services Users Guide*.

1.44 cluster heartbeat interval

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

Summary	Description
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.

1.45 cluster heartbeat retries

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

Summary	Description
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.

1.46 cluster redundancy level

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	Description
Default value	1)
Valid values	1 to $\langle n \rangle - 1$, where n is the maximum number of instances specified in <code>cluster.cfg</code> or the quorum file.
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

Usage

If the number of failed instances exceeds the maximum number specified by the cluster redundancy level (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` or for various cache sizes 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.

Set `cluster_redundancy_level` to 1 if the failure of multiple instances is expected to be uncommon.

Related Information

[max_memory \[page 154\]](#)

[number_of_locks \[page 195\]](#)

1.47 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	Description
Default value	60
Valid values	1 – 127
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
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.

1.48 column default cache size

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

Summary	Description
Default value	1024 pages (2 MB)
Valid values	128 – 8192 (pages)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

1.49 compression info pool size

Determines the size of the memory pool used for compression.

Summary	Description
Default value	4096
Valid values	0 – 2147483647
Status	Dynamic
Display level	10
Required role	System administrator

Summary	Description
Configuration group	SQL Server Administration

1.50 compression memory size

Determines the size (in 2 KB pages) of the memory pool that SAP ASE uses to decompress a compressed dump.

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

SAP ASE uses the value for `compression memory size` while loading a compressed dump into an archive database.

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*.

1.51 configuration file

Specifies the location of the configuration file currently in use.

Summary	Description
Default value	0 (off)
Valid values	One of: 0, <code>verify</code> , <code>read</code> , <code>write</code> , or <code>restore</code>
Status	Dynamic
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	General Information

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

1.52 cost of a logical io

Specifies the cost of a single logical I/O.

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

1.53 cost of a physical io

Specifies the cost of a single physical I/O.

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

1.54 cost of a cpu unit

Specifies the cost of a single CPU operation.

Summary	Description
Default value	0 – 254
Valid values	1 – 65534
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

Usage

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 an 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.

Related Information

[cost of a physical io \[page 44\]](#)

1.55 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	Description
Default value	200
Valid values	1 – 2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

Usage

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 `select` statements 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.

1.56 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	Description
Default value	500
Valid 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 information on task scheduling.

Related Information

[sql server clock tick length \[page 263\]](#)

[time slice \[page 279\]](#)

1.57 current audit table

Establishes the table where SAP ASE writes audit rows.

Summary	Description
Default value	1

Summary	Description
Valid values	0 – 8
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Security Related

Usage

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.

1.58 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	Description
Default value	500
Valid values	0 – 2147483
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Lock Manager

Usage

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*.

1.59 deadlock pipe active

Controls whether SAP ASE collects deadlock messages.

Summary	Description
Default value	0
Valid values	0 – 1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Memory Use, Monitoring

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

1.60 deadlock pipe max messages

Determines the number of deadlock messages SAP ASE stores per engine.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

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

Related Information

[sql text pipe max messages \[page 265\]](#)

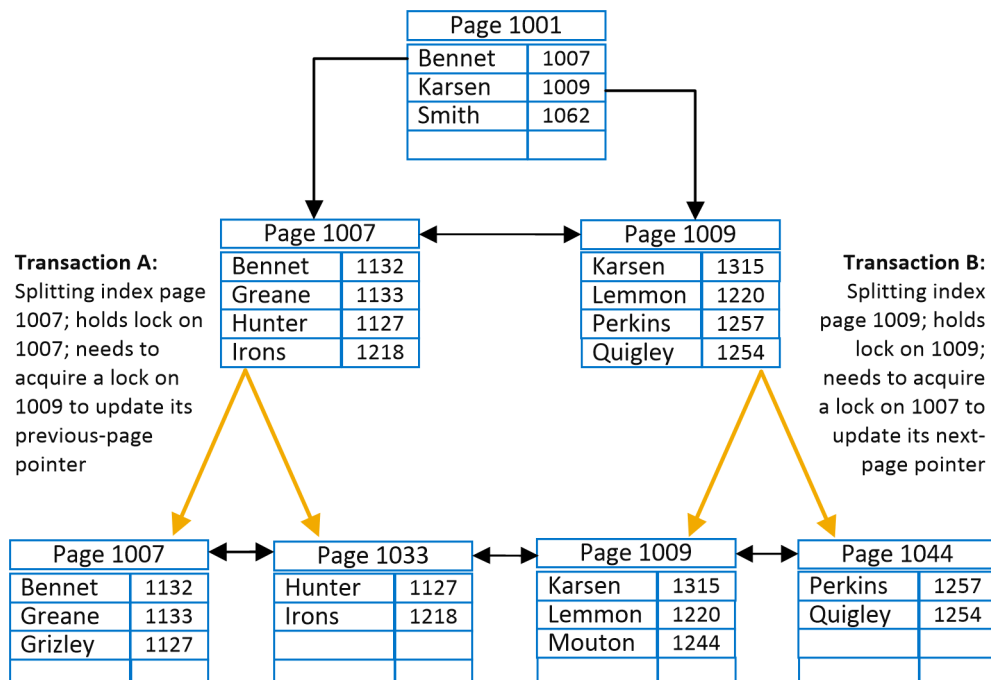
1.61 deadlock retries

Specifies the number of times a transaction can attempt to acquire a lock when deadlocking occurs during an index `page split` or shrink.

Summary	Description
Default value	5
Valid values	0 – 2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	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.



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*.

1.62 decompression row threshold

Determines the maximum number of columns in a table that are uncompressed using column decompression. The server uses row decompression if the number of columns is larger than this value.

Summary	Description
Default value	20
Valid values	5 – 100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

1.63 default character set id

Specifies the number of the default character set used by the server.

Summary	Description
Default value	1
Valid 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.

1.64 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	Description
Default value	3 MB
Valid 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.

i Note

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

See `create database` in *Reference Manual: Commands*.

1.65 default exp_row_size percent

Reserves space for expanding updates in data-only-locked tables, to reduce row forwarding.

Summary	Description
Default value	5
Valid values	0 – 100
Status	Dynamic
Display level	Intermediate

Summary	Description
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*, `create table` in *Reference Manual: Commands*, and `sp_chattribute` in *Reference Manual: Procedures*.

1.66 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	Description
Default value	0
Valid values	0–100
Status	Dynamic
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*.

1.67 default language id

Specifies 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	Description
Default value	0
Valid 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.

1.68 default network packet size

Configures the default packet size for all SAP ASE users.

Summary	Description
Default value	2048
Valid values	512– 65024
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	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*.

Related Information

[additional network memory \[page 17\]](#)

[max network packet size \[page 156\]](#)

[total logical memory \[page 280\]](#)

1.68.1 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
```

1.69 default sortorder id

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

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

1.70 default unicode sortorder

Specifies a string parameter that uniquely defines the default Unicode sort order installed on the server.

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

1.71 default xml sortorder

Specifies 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	Description
Default value	binary
Valid 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*.

1.72 deferred name resolution

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

Summary	Description
Default value	0 (disabled)
Valid values	0 to 1
Status	Dynamic
Display level	Comprehensive
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*.

1.73 disable character set conversions

Enables or disables character set conversion for data moving between clients and SAP ASE.

Summary	Description
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.

1.74 disable disk mirroring

Enables or disables disk mirroring for SAP ASE.

Summary	Description
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.

i Note

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

1.75 disable jsagent core dump

Enables or disables JS Agent core dumps.

Summary	Description
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 core dump` to on (1) disables core dumps and is not recommended.

i Note

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

1.76 disable varbinary truncation

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

Summary	Description
Default value	0 (on)
Valid values	0 (on), 1 (off)
Status	Dynamic
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	SQL Server Administration

1.77 disk i/o structures

Specifies the initial number of disk I/O control blocks SAP ASE allocates at start-up.

Summary	Description
Default value	256
Valid 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`.

1.78 dma object pool size

Specifies the number of DMA (direct memory access) objects allocated by CIPC at start-up time.

Summary	Description
Default value	4096
Valid values	2048 – 2147483647

Summary	Description
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

1.79 dtm detach timeout period

Sets the amount of time, in minutes, that a distributed transaction branch can remain in the detached state.

Summary	Description
Default value	5 (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.

1.80 dtm lock timeout period

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

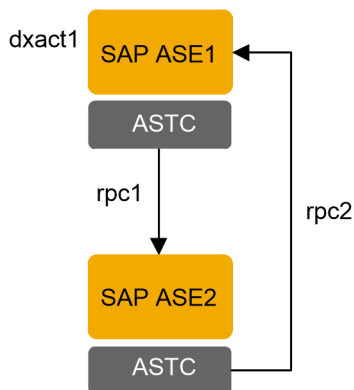
Summary	Description
Default value	300 (seconds)
Valid values	1 – 2147483647 (seconds)

Summary	Description
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	

Usage

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.



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.

i Note

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

1.81 dump history filename

Specifies the path and name of your dump history file.

Summary	Description
Default value	dumphist
Valid values	
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Backup/Recovery

1.82 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Group Diagnostics

i 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.

1.83 dynamic allocation on demand

Determines when memory is allocated for changes to dynamic memory configuration parameters.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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.
- 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.

1.84 dynamic sql plan pinning

Improves performance by reducing the time spent by server connections waiting for access to the query plan manager.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Unicode

Usage

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 reuse.

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.

1.85 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	Description
Default value	2147483647

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

1.86 enable async database init

Ensures that all `create database` and `alter database` commands initialize databases asynchronously by default.

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

1.87 enable backupserver ha

Enables or disables the high availability Backup Server for the cluster.

Summary	Description
Default value	1
Valid values	1 (enabled), 0 (disabled)
Status	Dynamic

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

1.88 enable buffered io for load

Enables Backup Server to open the database devices using buffered I/O at load time, regardless of the settings for devices being configured to use DSYNC or DIRECTIO.

Summary	Description
Default value	0
Valid values	1 (enabled), 0 (disabled)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator or manage server configuration
Configuration group	Backup/Recovery

By default, Backup Server uses the same settings as SAP ASE when opening a device. When `enable buffered io for load` is enabled, SAP ASE will request that Backup Server open the database devices using buffered I/O at load time, independently of the devices being configured to use synchronous I/O or direct I/O. Once the `load` command completes, all the in memory buffers will be flushed to disk. This configuration parameter will result in a heavy usage of the operating system file cache, and, in some cases, it can speed up the load operations.

1.89 enable bulk inserts

Allows you to perform bulk inserts with the `merge` command.

Summary	Description
Default value	0

Summary	Description
Valid values	0, 1
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

Bulk insert is turned off when it is set to 0. When `enable_bulk_inserts` is set to 1, it attempts bulk insert by caching all inserted rows if all other conditions are satisfied. When more than the eight predefined threshold pages are inserted, it will go through bulk insert. Otherwise, normal non-bulk insert is used. The eight-page limit cannot be reconfigured.

1.90 enable cis

Enables or disables Component Integration Service.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Component Integration Services

1.91 enable compression

Enables or disables data compression.

Summary	Description
Default value	0 (off)

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

See the *Compression Users Guide*.

1.92 enable concurrent dump tran

Enables or disables SAP ASE to use concurrent dumps.

Summary	Description
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.

1.93 enable console logging

Determines whether SAP ASE sends messages to the console.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	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.

1.94 enable deferred parallel

Configures the optimizer to generate parallel plans based on the best serial plan.

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

When `enable deferred parallel` is disabled (set to 0), the optimizer uses native parallel optimization. When `enable deferred parallel` is enabled (set to 1), the optimization time and memory resources are reduced.

1.95 enable delta dump tran

Ensures that two sequential transaction dumps do not contain duplicate pages.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System Administrator or manage server configuration privileges
Configuration group	Backup/Recovery

If you execute `dump transaction` on a database with long running open transactions, enabling `enable delta dump tran` may reduce the transaction dumps size and the amount of time required to perform the dump.

If you issue `dump transaction` concurrently with a full or cumulative dump with `enable delta dump tran` disabled, each transaction dump contains all log pages, cumulatively, since the first concurrent dump started. However, if `enable delta dump tran` is enabled, duplicate log pages are not included in the dump, resulting in significantly smaller transaction dump sizes.

1.96 enable dtm

Enables or disables the SAP ASE distributed transaction management (DTM) feature.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Static
Display level	10
Required role	System administrator
Configuration group	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.

1.97 enforce dump configuration

Enables dump operations to use a dump configuration.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Backup/Recovery

1.98 enable dump history

Determines whether there are updates to the dump history file at the end of the database dump operation.

Summary	Description
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.

1.99 enable encrypted columns

Enables and disables encrypted columns.

Summary	Description
Default value	0 (off)
Valid 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'
```

i Note

Using encrypted columns increases the logical memory used by 8198 kilobytes.

1.100 enable file access

Enables access through proxy tables to the external file system.

Summary	Description
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.

1.101 enable fips

Enables FIPS 140-2 cryptographic module support for encrypting passwords in the backup server.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System security officer
Configuration group	Security Related

The backup server uses the FIPS certified Common Crypto Library to perform login password encryption when this parameter is set to 1, otherwise it uses the Common Crypto Library. You need to reboot the backup server for this parameter to take effect.

1.102 enable full-text search

Enables Enhanced Full-Text Search services.

Summary	Description
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.

1.103 enable functionality group

Enables or disables the changes available for specific features from SAP ASE version 15.7 and later.

Summary	Description
Default value	0
Valid 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.

1.104 enable granular permissions

Enables or disables grantable system privileges that allow you to enforce separation of duties.

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

The granular permissions feature requires the ASE_PRIVACY license. You must have `sso_role` privileges to turn on granular permissions, and the `manage security configuration` system privilege to turn off granular permissions.

1.105 enable ha

Enables or disables SAP ASE as a companion server in an active-active high availability subsystem.

Summary	Description
Default value	0 (off)
Valid 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.

i 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).

1.106 enable HCB index

Enables or disables SAP ASE for hash-cached BTree indexes.

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

1.107 enable housekeeper gc

Configures the housekeeper task.

Summary	Description
Default value	1 (on)
Valid values	<ul style="list-style-type: none">0 – disables the housekeeper garbage collection task, but enables the <code>delete</code> command's lazy garbage

Summary

Description

collection. You must use `reorg_reclaim_space` to deallocate empty pages.

Note

You should not use this value. Although this is the cheapest option with the lowest performance impact, using this value may cause performance problems if many empty pages accumulate.

- 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 – general garbage collection is disabled, except for DOL row deletes where the server uses aggressive garbage collection (not recommended).
- 3 – general lazy garbage collection is enabled, and the server uses aggressive DOL deletes (not recommended).
- 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

Status	Dynamic
Display level	Comprehensive
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*.

`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*.

1.108 enable hp posix async i/o

Enables or disables asynchronous I/O on database devices created on HP-UX 11.31 and later file systems.

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

<code>enable hp posix async i/o</code>	<code>allow sql server async i/o = 0</code>	<code>allow sql server async i/o = 1</code>
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
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

i 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.

1.109 enable hugepages

Enables and disables the use of huge pages on Linux platforms that support huge pages.

Summary	Description
Default value	0
Valid values	<ul style="list-style-type: none">• 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.
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Physical Memory

1.110 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

1.111 enable inline default sharing

Enables SAP ASE to share inline defaults.

Summary	Description
Default value	0 (off)
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*.

i 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`.

1.112 enable in-memory row storage

Enables SAP ASE to use in-memory row storage.

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

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

Enabling this configuration parameter does not consume extra memory. However, creating row storage caches incurs memory consumption.

The `manage server configuration` privilege is required to set or change the `enable in-memory row storage` configuration parameter if you have enabled granular permissions.

Interactions when restarting the server:

- Device startup – Any devices of type `imrslog` are activated, even if `enable in-memory row storage` is disabled, allowing for a complete server recovery.
- Cache manager – row storage caches are initialized during a restart, even if `enable in-memory row storage` is disabled, and all memory pools associated with these caches are created and initialized, allowing for a complete recovery in the IMRS.
- online database – The database to which an in-memory row storage was assigned may be fully recovered but will fail to come online if `enable in-memory row storage` is disabled.

Interactions with DMLs:

- Typically, issuing selects or DMLs does not depend on row caching or snapshot isolation, allowing the system and business transactions to continue during changes to the IMRS license or the `enable in-memory row storage` configuration option.
- `dump database` and `dump transaction` – issuing these commands does not rely on `enable in-memory row storage` being enabled, allowing the database administrator to migrate the database's contents to another server configured and licensed for the IMRS.
- `load database` – you cannot load a database dump or transaction log from a database enabled with row storage caches if the target database does not have `enable in-memory row storage` enabled.

1.113 enable ISM (Solaris only)

Enables and disables SAP ASE to use intimate shared memory (ISM) on the Solaris platform.

Summary	Description
Default value	0
Valid values	<ul style="list-style-type: none"> • 0 – use ISM if possible; if ISM is not available, use regular shared memory. • 1 – use only regular shared memory.

Summary	Description
	<ul style="list-style-type: none"> 2 – use only ISM.
Status	Static
Display level	
Required role	System administrator
Configuration group	Physical Memory

1.114 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Java Services

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).

Related Information

[runnable process search count \[page 247\]](#)

1.115 enable job scheduler

Determines whether Job Scheduler starts when the SAP ASE server starts.

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

1.116 enable js restart logging

Enables or disables diagnostics logging after the restart of Job Scheduler.

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

1.117 enable large chunk elc

Enables or disables large allocation in the engine local cache.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Meta-Data Caches

1.118 enable large pool for load

Configures the use of large buffer pools during the recovery phase for load database and load transaction commands.

Summary	Description
Default value	<p>The default value for <code>enable large pool for load</code> depends on the value to which <code>enable functionality group</code> is set. If you set <code>enable functionality group</code> to:</p> <ul style="list-style-type: none">• 0 – the default value for <code>enable large pool for load</code> is 0• 1 – the default value for <code>enable large pool for load</code> is 1 <p>However, if you set <code>enable large pool for load</code> to 1, it uses a value of 1 regardless of what you set <code>enable functionality group</code> to.</p>
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
Configuration group	SQL Server Administration

Related Information

[enable functionality group \[page 76\]](#)

1.119 enable ldap user auth

Enables or disables SAP ASE to authenticate each user on the LDAP server.

Summary	Description
Default value	0 (off)
Valid values	<ul style="list-style-type: none"> 0 (off) – allows only <code>syslogins</code> authentication. 1 (on) – allows both LDAP and <code>syslogins</code> authentication. 2 (on) – allows only LDAP authentication.
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
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.

1.120 enable LFB index

Enables or disables Latch-Free Indexes feature (part of the Mem Scale licensable option).

i Note

Before enabling LFB index, set the `enable mem scale` configuration parameter.

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

1.121 enable lightweight rvm

Enables and disables (RVM) Reference Validation Mechanism checks to determine if there are no permission changes between executions of the same compiled plans, streamlining the query execution path.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System Administrator
Configuration group	Query Tuning

Typically, the query processor performs lightweight permission checks only for in-memory database queries. Enable `enable lightweight rvm` to allow lightweight permission checks for database queries.

1.122 enable literal autoparam

Enables and disables literal server-wide parameterization.

Summary	Description
Default value	0 (off)

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

1.123 enable lock remastering

Allows an SAP ASE Cluster Edition background process to locally manage locks once failed-over client connections move back to the original instance.

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

1.124 enable logins during recovery

Determines whether non-system administrator logins are allowed during database recovery.

Summary	Description
Default value	1
Valid values	<ul style="list-style-type: none"> 0 (enabled) – indicates that logins are not allowed during recovery, that is, only the system administrator can log in to SAP ASE

Summary	Description
	<ul style="list-style-type: none"> 1 (disabled) – indicates that logins are allowed during recovery
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Security Related

1.125 enable mem scale

Enables or disables the Mem Scale licensable option.

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

1.126 enable merge join

Enables or disables merge join at the server level.

Summary	Description
Default value	2
Valid values	<ul style="list-style-type: none"> 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.

Summary	Description
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
<code>allows_mix</code>	on
<code>allows_dss</code>	on
<code>allows_oltp</code>	off

1.127 enable metrics capture

Enables SAP ASE to capture metrics at the server level.

Summary	Description
Default value	0 (off)
Valid 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.

1.128 enable_monitoring

Controls whether SAP ASE collects the monitoring table data.

Summary	Description
Default value	0 (off)
Valid 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.

i Note

Any data that populates the monitoring tables without enabling `enable_monitoring` should not be considered valid.

1.129 enable_pam_user_auth

Controls the ability to authenticate users using pluggable authentication modules (PAM).

Summary	Description
Default value	0 (off)
Valid values	<ul style="list-style-type: none">0 (off) – allows only <code>syslogins</code> authentication.1 (on) – allows both PAM and <code>syslogins</code> authentication. When <code>enable_pam_user_auth</code> is set to 1, SAP ASE uses the PAM provider to authenticate each user. If the PAM authentication fails, SAP ASE searches <code>syslogins</code> to authenticate the user. Use level 1 when you are migrating users from SAP ASE authentication to PAM authentication.2 (on) – allows only PAM authentication.

Summary	Description
Status	Dynamic
Display level	Intermediate
Required role	System security officer
Configuration group	Security Related

1.130 enable pci

Enables or disables the Java PCI Bridge for SAP ASE.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on), 2 (on with operating system override) 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.
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	User Environment

1.131 enable permissive unicode

Enables or disables SAP ASE to ignore Unicode noncharacters.

Summary	Description
Default value	The default value for <code>enable permissive unicode</code> depends on the value to which <code>enable</code>

Summary	Description
	<p>functionality group is set. If you set enable functionality group to:</p> <ul style="list-style-type: none"> • 0 – the default value for enable permissive unicode is 0 • 1 – the default value for enable permissive unicode is 1 <p>However, if you set enable permissive unicode to 1, it uses a value of 1 regardless of what you set enable functionality group to.</p>
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.

Related Information

[enable functionality group \[page 76\]](#)

1.132 enable predicated privileges

Enables or disables predicated privileges.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1(on)
Status	Dynamic

Summary	Description
Display level	Comprehensive
Required role	System administrator
Configuration group	Security Related

1.133 enable query tuning mem limit

Enables or disables the query tuning memory limit.

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

1.134 enable query tuning time limit

Enables or disables the query tuning time limit.

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

1.135 enable rapidlog

Enables the capture of diagnostic for `Proc Cache Header` memory pool messages.

Summary	Description
Default value	0 (off)
Valid values	0 – 255
Status	Dynamic
Display level	
Required role	System administrator
Configuration group	Error Log

i Note

Use diagnostic for `Proc Cache Header` memory pool only under the guidance of the SAP support team.

See *Performance and Turning Series: Basics > Diagnostic for 'Proc Cache Header' Memory Pool*.

1.136 enable rapidtimer

Enables the collection of SAP ASE diagnostic information. Instructions for `enable rapidtimer` are provided by SAP Technical Support.

Summary	Description
Default value	0 (off)
Valid values	0 – 65536
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Error Log

1.137 enable real time messaging

Enables or disables the real time messaging services.

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

1.138 enable rep agent threads

Enables or disables the RepAgent thread within SAP ASE.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Memory Use, Rep Agent Thread Administration

Other steps are also required to enable replication. For more information, see the Replication Server documentation.

1.139 enable resolve as owner

Enables the resolve as owner functionality, which allows users to view data from objects they do not own, and without including the owner name prefix with the object.

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

1.140 enable row level access control

Enables or disables row-level access control.

Summary	Description
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`.

1.141 enable select into in tran

Enables and disables the ability to use the `select into` command in a multistatement transaction.

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 > Transactions: Maintain Data Consistency and Recovery > Using select into in Multistatement Transactions*.

Note

The default value for `enable select into in tran` depends on the value to which `enable functionality group` is set. If you set `enable functionality group` to:

- 0 – the default value for `enable select into in tran` is 0.
- 1 – the default value for `enable select into in tran` is 1.

However, if you set `enable select into in tran` to 1, it uses a value of 1 regardless of what you set `enable functionality group` to.

See `enable functionality group`.

1.142 enable semantic partitioning

Enables or disables partitioning other than round-robin (for example list, hash, and range partitioning) in SAP ASE.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic

Summary	Description
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.

1.143 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	Description
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 versions of SAP ASE earlier than 15.0 to use additional `join` clauses.

For more information about compatibility mode, see the *Migration Technology Guide*.

1.144 enable sql debugger

Enables or disables the SAP ASE SQL debugger, which allows you to step through your T-SQL code.

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

Summary	Description
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

1.145 enable ssl

Enables or disables Secure Sockets Layer session-based security.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Security Related

1.146 enable sticky statistics

Allows you to disable stickiness for `update_statistics`.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
Configuration group	Query Tuning

1.147 enable stmt cache monitoring

Enables or disables SAP ASE to collect monitoring information about the statement cache.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

Once enabled, `monStatementCache` and `monCachedStatement` display valid data.

1.148 enable streamlined parallel

Reduces the default parallel plan search space, avoids bad parallel plans, and focuses the parallel optimizer on a more promising parallel plan search space.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

When `enable streamlined parallel` is enabled, the query processor:

- Avoids repartitioning on both sides of a join.
- Repartitions only on the smaller side of a join.
- Avoids parallel reformatting.
- Avoids putting an `N->1 xchg` enforcement over any subplan other than two phase aggregate pop or immediately under `Emit`.

1.149 enable surrogate processing

Enables or disables the processing and maintains the integrity of surrogate pairs in Unicode data.

Summary	Description
Default value	1 (on)
Valid 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.

1.150 enable transactional memory

Enables or disables the Transactional Memory feature (part of Mem Scale licensable option).

i Note

Before enabling transactional memory, set the `enable mem scale` configuration parameter.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic

Summary	Description
Display level	Comprehensive
Required role	System administrator
Configuration group	Processors

1.151 enable unicode conversion

Enables or disables character conversion using Unilib for the `char`, `varchar`, and `text` datatypes.

Summary	Description
Default value	1
Valid values	<ul style="list-style-type: none"> 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 group	Languages, Unicode

1.152 enable unicode normalization

Enables or disables Unilib character normalization.

Summary	Description
Default value	1 (on)
Valid values	<ul style="list-style-type: none"> 0 (off) – 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—

Summary	Description
	<p>but only if <i>all</i> clients present Unicode data to the server using the same representation.</p> <ul style="list-style-type: none"> 1 (on) – set <code>enable unicode normalization</code> to 1 to use the built-in process that enforces normalization on all incoming Unicode .data.
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.

Note

Once disabled, you cannot turn normalization on again. This limitation prevents non-normalized data from entering the database.

1.153 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	Description
Default value	<p>The default value for <code>enable utility lvl 0 scan wait</code> depends on the value to which <code>enable functionality group</code> is set. If you set <code>enable functionality group</code> to:</p> <ul style="list-style-type: none"> 0 – the default value for <code>enable utility lvl 0 scan wait</code> is 0. 1 – the default value for <code>enable utility lvl 0 scan wait</code> is 1. <p>However, if you set <code>enable utility lvl 0 scan wait</code> to 1, it uses a value of 1 regardless of what you set <code>enable functionality group</code> to. See <code>enable functionality group</code>.</p>

Summary	Description
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

1.154 enable webservices

Enables or disables Webservices.

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

1.155 enable workload analyzer

Enables and disables the workload analyzer.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	SQL Server Administration

1.156 enable xact coordination

Enables or disables SAP ASE transaction coordination services.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on) <ul style="list-style-type: none"> 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.
Status	Static
Display level	10
Required role	System administrator
Configuration group	DTM Administration

1.157 enable xml

Enables or disables XML services.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on - not external), 2 (on - external)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

Setting `enable xml` to:

- 0 – disables XML services.
- 1 – enables XML services, but not with an external entity reference.
- 2 – enables XML services with an external entity reference.

1.158 engine local cache percent

Enables or disables modification of the Engine Local Cache (ELC) as a percentage of procedure cache.

Summary	Description
Default value	50
Valid 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 percent of procedure cache size.

This configurations option can be used only when large ELC is enabled (using `enable large chunk elc` configuration parameter).

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
```

1.159 engine memory log size

For diagnostic use only and has no relevance in a production environment. Keep all default settings unless otherwise requested by SAP Technical Support.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	
Required role	
Configuration group	Physical Memory

1.160 errorlog pipe active

Controls whether SAP ASE collects error log messages.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic

Summary	Description
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 the `monErrorLog` monitoring table to retrieve these error log messages.

Related Information

[errorlog pipe max messages \[page 110\]](#)

1.161 errorlog pipe max messages

Determines the number of error log messages SAP ASE stores per engine.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

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

Related Information

[errorlog pipe active \[page 109\]](#)

[sql text pipe max messages \[page 265\]](#)

1.162 errorlog size

Sets the threshold for the size of the error log. Once this threshold is reached, SAP ASE dynamically closes the current error log and opens a new one .

Summary	Description
Default value	0 (off)
Valid values	0 – 200GB
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

The error log rollover feature is enabled once you set the threshold to a size between 1MB and 200GB. Once enabled, SAP ASE checks the size of its error log every minute and, if it exceeds the threshold, rolls over the error log.

1.163 esp execution priority

Sets the priority of the XP Server thread for ESP execution.

Summary	Description
Default value	8
Valid 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.

1.164 esp execution stacksize

Sets the size of the stack, in bytes, to be allocated for ESP execution.

Summary	Description
Default value	Platform dependent: <ul style="list-style-type: none">• 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
Valid 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.

1.165 esp unload dll

Specifies whether DLLs that support ESPs should be automatically unloaded from XP Server memory after the ESP call has completed.

Summary	Description
Default value	0 (off)
Valid 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. You can still unload individual DLLs explicitly at runtime, using `sp_freedll`.
- 1 – DLLs are automatically unloaded.

1.166 event buffers per engine

Controls the number of monitor event buffers. Not used in the current version of SAP ASE.

1.167 event log computer name (Windows Only)

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	Description
Default value	LocalSystem
Valid values	<ul style="list-style-type: none"> • Name of a 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*.

1.168 event logging (Windows only)

Enables and disables the logging of SAP ASE messages in the Windows Event Log.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Error Log

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.

1.169 executable codesize + overhead

A calculated value that is not user-configurable, this parameter exports the combined size, in kilobytes, of the SAP ASE executable and overhead.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Calculated
Display level	Basic
Required role	System administrator
Configuration group	Memory Use

1.170 extended cache size

Specifies the size of the secondary cache.

Summary	Description
Default value	0
Valid values	0 – 31457280
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	Cache Manager

1.171 extend implicit conversion

Controls the use of implicit conversion.

Summary	Description
Default value	0
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	Server Configuration Management
Configuration group	Query Tuning

When `extend implicit conversion` configuration is enabled, implicit conversion between character types and arithmetic types is allowed.

The supported character types are `char` and `varchar`, and arithmetic types are `smallint`, `int`, `bigint`, `numeric` and `decimal`.

Behavior When Off

When casting between a character type and a numeric type, the size of the result is the size of the numeric expression:

Example: "11.1" + 111.11 = 122.21

The result type is numeric(5,2) (or 2 decimal places and a precision of 5 digits) and this is based on the numeric argument of 111.11

Example: "111.11" + 1.1 = Overflow

The result type that of the numeric argument is numeric (2,1).

Behavior When On

When the value of the character string is known during parsing, the numeric precision and length are determined by examining the string literal. When the value of the character string is not known during parsing, for example the character string is stored in a local variable, the following formula applies, however, the precision and scale are not allowed to exceed their respective maximums, which are 77 and 38.

Numeric([string length -1] x 2], string length -1)

1.172 external keystore

SAP ASE can use the external keystore to encrypt or decrypt master encryption keys. `external keystore` enables SAP ASE to store key externally.

Summary	Description
Default value	Null
Valid values	Null, HSM, or hsm
Status	Dynamic
Display level	Comprehensive
Required role	<ul style="list-style-type: none">When granular permission is disabled, you need to have the <code>sso_role</code> permissions.When granular permission is enabled, you need to have the <code>manage security configuration</code> permissions.
Configuration group	Security Related

1.173 fips login password encryption

Provides FIPS 140-2 cryptographic module support for encrypting passwords in transmission, in memory, and on disk.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
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. If this configuration is enabled, SAP ASE uses the FIPS certified Common Crypto Library to perform login password encryption, otherwise it uses the Common Crypto Library.

1.174 FM enabled

Set or reset by the Fault Manager to indicate to SAP ASE that it, and the HADR setup, is being monitored. Do not change the value of `FM enabled` unless you are explicitly told to by the HADR product documentation or by SAP Technical Support.

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

1.175 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	Description
Default value	10
Valid 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*.

1.176 global cache partition number

Sets the default number of cache partitions for all data caches.

Summary	Description
Default value	1
Valid 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*.

1.177 HADR connect timeout

Determines the wait time for the successful connection between primary SAP ASE, standby SAP ASE, and SAP ASE on the DR site in an HADR environment.

Summary	Description
Default value	60 (sec)
Valid values	1 – 32767 (sec)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	SQL Server Administration

1.178 HADR distinct server name

Enables the server to append "DR" to the server name for the `HADR_MEMBER` server class entries in `sys.servers`. For example, a server named "PARIS" is listed as "PARISDR" in `sys.servers`.

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

i Note

Do not change the value for `HADR distinct server name` unless instructed to do so by technical support. Changing it can lead to misconfiguration of the HADR system.

1.179 HADR login stall time

The amount of time (in seconds) the standby server delays before sending a redirect list when it detects that the other server is not in `primary active` mode.

Summary	Description
Default value	2
Valid values	1 – 32767
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	SQL Server Administration

1.180 HADR mode

Determines the mode of the instance.

Summary	Description
Default value	-1 (HADR is disabled)
Valid values	HADR mode is one of: <ul style="list-style-type: none">• -1 – HADR is disabled.• 0 – HADR is enabled. This is a standby server.• 1 – HADR is enabled. This is a primary server• 2 – Used internally by SAP ASE.• 3 – Used internally by SAP ASE.• 4 – Used internally by SAP ASE.• 5 – HADR is enabled. This is a disaster recovery (DR) standby node.
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	SQL Server Administration

i Note

Do not set HADR mode to values 2, 3, and 4. These values are used internally by SAP ASE.

1.181 HADR primary check frequency

Determines how often, in seconds, the standby server checks the primary server's HADR mode and state.

Summary	Description
Default value	10
Valid values	1 – 32767
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	SQL Server Administration

1.182 HADR primary wait time

The amount of time, in seconds, the standby server continues to send the redirect list to the clients in absence of primary server before failing the connection.

Summary	Description
Default value	300
Valid values	1 – 32767
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	SQL Server Administration

Setting a value for `HADR primary wait time` is effective during failover since it defines the amount of time the standby server sends the redirect list to the client when there is no primary server present. After this time passes, the standby server fails the connection.

1.183 HADR remote query timeout

Determines the wait time for the successful query while retrieving the HADR state of the primary SAP ASE from standby SAP ASE or SAP ASE on the DR node.

Summary	Description
Default value	0 (sec)
Valid values	<ul style="list-style-type: none">0 (sec): no timeout1 – 32767 (sec)
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	SQL Server Administration

1.184 hash table hit rate threshold

Sets the percentage threshold for HCB auto tuning hash table scan hit rate.

Summary	Description
Default value	60
Valid values	0 – 100
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

A hash table scan hit rate that falls below this threshold indicates that index hash caching may not be effective, and may result in the HCB auto tuning task disabling index hash caching on these indexes.

1.185 HCB index auto tuning

Enables and disables hash-cache BTree index auto tuning.

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

`HCB index auto tuning` allows SAP ASE to find, and configure, indexes that are suitable for enabling hash caching. Later, if the server determines that the hash is not used efficiently, it disables hash caching on those indexes and re-enables them if the workload of point query increases. Since the server performs these tasks automatically, you can more readily use index hash caching and with low overhead, especially for large database systems with hundreds of tables and indexes.

The server creates the HCB auto tuning system task when you enable `HCB index auto tuning`. This task wakes up periodically to evaluate whether hash caching is efficient on the active indexes. If it determines hash caching is not efficient, it disables hash caching on those indexes. If hash caching is disabled and the workload of point query has increased, it re-enables hash caching on those indexes.

When you disable `HCB index auto tuning`, the HCB auto-tuning system task disables index hash caching on the indexes it manipulated, and then exits.

Consider the following when you enable `HCB index auto tuning`:

- `HCB index auto tuning` only works on unique indexes on IMRS-enabled tables
- `HCB index auto tuning` does not work on small tables
- All tuning decisions (that is, enabling or disabling hash caching) is lost when you restart the server

1.186 HCB index memory pool size

Determines the size of the memory pool available for hash-cached BTree indexes.

Summary	Description
Default value	4096
Valid values	0 – 2147483647

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

The memory pool is not created by default until you enable the `enable HCB index` configuration parameter. This memory pool is created with the default size, but users can change the memory size later with the `HCB index memory pool size` configuration parameter.

The default value is listed in 2KB bytes, so the default size of the memory pool is $4096 \times 2\text{KB} = 8192\text{KB}$. That is, the default size is 8 MB.

1.187 HCB index tuning interval

Determines the interval (in minutes) after which the HCB index tuning task wakes and tunes the hash-cached B-tree index.

Summary	Description
Default value	10
Valid values	0 – 32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

A value of 0 means the auto tuning thread does not wake up.

1.188 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	Description
Default value	4 KB
Valid values	0 – 2147483647 bytes
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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. You should set `heap memory per user` to three times the size of your logical page.

1.189 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	Description
Default value	20
Valid 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`.

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).

1.190 housekeeper free write percent

Specifies the maximum percentage by which the housekeeper wash task can increase database writes.

Summary	Description
Default value	1
Valid values	0 – 100
Status	Dynamic
Display level	Intermediate

Summary	Description
Required role	System administrator
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 `sysstat` 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.

1.191 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	Description
Default value	1000
Valid 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.

1.192 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	Description
Default value	100

Summary	Description
Valid 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.

1.193 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	Description
Default value	10
Valid values	1 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

i 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 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*.

Related Information

[cpu grace time \[page 47\]](#)

[sql server clock tick length \[page 263\]](#)

[time slice \[page 279\]](#)

1.194 identity burning set factor

Changes the percentage of potential column values that is made available in a block of column values.

Summary	Description
Default value	5000
Valid 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`.

1.195 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	Description
Default value	1
Valid values	1 – 2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

`identity grab size` is useful when you are performing inserts and you want all the inserted data to have contiguous IDENTITY numbers. For example, 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.

You should set `identity grab size` to a value large enough to accommodate the largest group of records you want to insert into contiguous rows.

1.196 identity reservation size

Sets a limit for the number of identity values.

Summary	Description
Default value	1
Valid values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

1.197 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	Description
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.

1.198 imrs cache utilization

Specifies a threshold, as a percentage, for the size of row_storage cache bound to an IMRS- or on disk MVCC-enabled database. The server employs heuristics internally to maintain the row storage cache to the value specified by `imrs cache utilization`.

Summary	Description
Default value	70
Valid values	10 – 90
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	IMRS Background System Tasks

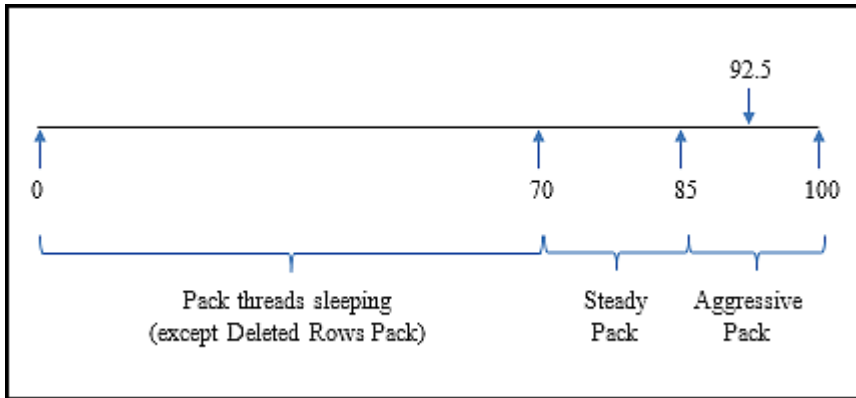
Generally, the transaction workload increases the cache utilization in the row storage and background pack task by packing the cold rows from the row storage cache. The server uses this threshold to determine when, and how aggressively, it wants to perform the pack operation to bring the row storage cache utilization to the specified value.

The appropriate value for `imrs cache utilization` depends on the workload and the size of the row storage. However these are some general guidelines to help determine a value:

- Specifying a smaller value for `imrs cache utilization` triggers pack early and more aggressively and helps to quickly reduce row storage cache utilization. However, a very small value may pack some hot or warm data from the row storage cache early, causing performance degradation when subsequently accessing this data.
- Specifying a larger value helps retain more hot data in the row storage cache. However very large values for `imrs cache utilization` may cause out-of-memory errors (error number 728), since these values make the pack subsystem less aggressive.
Reducing the value for `imrs cache utilization` may help to avoid out-of-memory problems.
Specifically, use a smaller number if you are using `snapshot isolation`.

In some workloads, if pack does not help to bring the cache utilization to the value of `imrs cache utilization`, having a highly utilized row storage cache may cause the server to temporarily disable new data storage in the row storage cache, reducing the load on the cache and reducing cache utilization to the value for `imrs cache utilization`. However, temporarily disabling new data storage is done only if you have not enabled `snapshot isolation` on the table. If `snapshot isolation` is enabled, every DML operation must create a version in the row storage cache to provide `snapshot isolation`.

Based on the value of `imrs cache utilization`, the pack operation decides to operate in different modes to maintain cache utilization to the requested value. This example sets the `imrs cache utilization` to 70:



The modes are:

- Steady state pack – pack runs in this mode when utilization of IMRS cache goes beyond the value set for `imrs cache utilization` (70 in the example above). It continues in this mode until the threshold for aggressive pack is reached. In this mode, pack attempts to pack rows by applying ILM heuristics. Based on partition level counters and timestamps, ILM heuristics indicate whether row is likely to be cold. Only those rows that ILM indicates are cold are packed in this mode.
- Aggressive pack – pack runs in this mode when the percentage utilization of IMRS cache goes beyond the midway point between `imrs cache utilization` and 100 (85 in the example above).
- Early pack of deleted rows – pack runs in this mode only when percentage utilization of IMRS cache is lower than the value of `imrs cache utilization` (70 in the example above) and there are large number of deleted rows in at least one of the IMRS enabled partitions in a database.

1.199 inline table functions

Inlines those inline table UDFs by tranforming them into the parameterized views internally during the execution.

Summary	Description
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 `inline table functions` is set to 1, all inline table UDFs created after it will be changed to parameterized views internally.

1.200 job scheduler interval

Sets the interval when the Job Scheduler checks which scheduled jobs are due to be executed.

Summary	Description
Default value	1 (in seconds)
Valid values	1 – 600
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

1.201 job scheduler memory

Determines the size of the memory pool (which is of type `bucketpool`) assigned to the Job Scheduler.

Summary	Description
Default value	8 MB
Valid values	4 – 1024
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

1.202 job scheduler tasks

Sets the maximum number of jobs that can run simultaneously through Job Scheduler.

Summary	Description
Default value	4
Valid 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, 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.

Related Information

[max memory \[page 154\]](#)

[number of user connections \[page 210\]](#)

1.203 js heartbeat interval

Specifies the intervals between two JS Agent heartbeat checks, in minutes.

Summary	Description
Default value	1
Valid values	1 – 1440
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	SQL Server Administration

1.204 js job output width

Determines the line width the output uses for jobs stored in the `js_output` table.

Summary	Description
Default value	80
Valid values	1 – 32768
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

1.205 js restart delay

Sets the delay period between two Job Scheduler auto restart attempts after abnormal shutdown of Job Scheduler.

Summary	Description
Default value	60
Valid values	0 – 1440
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

1.206 kernel mode

Determines the mode the SAP ASE kernel uses, threaded or process.

Summary	Description
Default value	<code>threaded</code>
Valid values	<code>process</code>
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 compatibility. Process mode may not support features for SAP ASE 15.7 and later.

See *System Administration Guide: Volume 2 > Managing Multiprocessor Servers*.

1.207 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	Description
Default value	<ul style="list-style-type: none">(Solaris) 8384(All other platforms) 6396
Valid values	75 – 2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Physical 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.

SAP ASE adjusts the required value of `kernel resource memory` automatically. Modify this configuration parameter only under the guidance of technical support.

Related Information

[max memory \[page 154\]](#)

[max online engines \[page 160\]](#)

[number of user connections \[page 210\]](#)

1.208 large allocation auto tune

Configures SAP ASE preallocate large amounts of memory for query execution, which reduces procedure cache contention.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Meta-Data Caches

1.209 LFB memory size

Determines the size of the Latch-Free B-tree (LFB) bucket pool.

Summary	Description
Default value	4096
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

The LFB bucket pool is designed to help resolve the following types of spinlock contention:

- Page level – this occurs when queries and DMLs block each other on same page.
- Spin level – acquiring an SH_LATCH on the same page can be slow because of the large number of concurrent queries. You can avoid page level contention using a skip share latch on index pages to avoid spinlock contention caused by allocating and deallocating memory.

The LFB bucket pool is used for the LFB mapping table and for delta updates to resolve these spinlock contentions.

1.210 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	Description
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*.

1.211 lock address 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	Description
Default value	100
Valid values	1 – 2147483647
Status	Static
Display level	Comprehensive

Summary	Description
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 `lock address 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 `lock address spinlock ratio`, SAP ASE uses only 1 spinlock for the entire table.

1.212 lock hashtable size

Specifies the number of hash buckets in the lock hash table.

Summary	Description
Default value	2048
Valid values	1 – 2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	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`.

Related Information

[number of locks \[page 195\]](#)

1.213 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	Description
Default value	allpages
Valid 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
```

1.214 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)

Summary	Description
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.

1.215 lock spinlock ratio

For SAP ASEs running with multiple engines, 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	Description
Default value	85
Valid values	1 – 2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	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 Multiprocessor 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*.

Related Information

[lock hashtable size \[page 142\]](#)

1.216 lock table spinlock ratio

For SAP ASEs running with multiple engines, sets the number of rows in the internal table locks hash table that are protected by one spinlock.

Summary	Description
Default value	20
Valid 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 `lock table spinlock ratio` is 20, which defines six spinlocks for the table locks hash table. The first five spinlocks protect 20 rows each; the sixth spinlock protects the last row. If you specify a value of 101 or greater for `lock table spinlock ratio`, SAP ASE uses only one spinlock for the entire table.

1.217 lock timeout pipe active

Controls whether SAP ASE collects lock timeout messages.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic

Summary	Description
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` monitoring table.

Related Information

[lock timeout pipe max messages \[page 146\]](#)

1.218 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	Description
Default value	0
Valid values	0 – 2147483648
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

1.219 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	Description
Default value	2147483647
Valid 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"
```

1.220 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on), 2 (on - allows event logging when auditing is not enabled)
Status	Dynamic
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	Error Log

An administrator should consider the security implications of setting value 2. This allows the reason for the logon failure into the errorlog. Considerations should include how well the errorlog is protected from hackers attempting to crack passwords and learning details of a login account, such as: account name is valid, or account is locked, or a password is incorrect.

1.221 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Error Log

1.222 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	Description
Default value	Platform-dependent
Valid values	1 – platform-dependent value
Status	Static
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	O/S Resources

1.222.1 On the Linux Platform

On the Linux platform, `max_async_i/os_per_engine` controls the number of asynchronous I/Os 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 I/Os 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 I/O 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.

For more information, see *Enable Asynchronous Disk I/O* and `solaris_async_i/o_mode` (Solaris Only).

1.223 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	Description
Default value	Platform-dependent
Valid values	1– platform-dependent value
Status	Static
Display level	Comprehensive

Summary	Description
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:

Table 1: Distribution of Asynchronous I/O Pending

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 async 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*.

Related Information

[disk i/o structures \[page 61\]](#)

[max async i/os per engine \[page 148\]](#)

1.224 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	Description
Default value	2048
Valid 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,
- $(\text{max resource granularity}) \times (\text{the number of data buffers in tempdb's pagesize pool})$

Related Information

[max resource granularity \[page 165\]](#)

1.224.1 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:
 $(\text{number of sort buffers}) \times (\text{rows_per_page}) \times 100$
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:
 $\text{max resource granularity} \times (\text{the number of cache buffers in the tempdb's single-page pool})$

i Note

The differences between `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.

1.225 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	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic

Summary	Description
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.

1.226 max concurrently recovered db

Determines the degree of parallelism.

Summary	Description
Default value	0
Valid values	1 – number of engines at start-up minus 1 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.
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Backup/Recovery

1.227 max js restart attempts

Restricts the number of restart attempts and prevents the Job Scheduler restart feature from going into an infinite loop.

Summary	Description
Default value	3
Valid values	0 – 10. The value 0 indicates that the Job Scheduler Auto restart feature is disabled.
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	SQL Server Administration

1.228 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	Description
Default value	Platform-dependent
Valid values	Platform-dependent minimum – 2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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.

Use the `allow memory grow at startup` configuration parameter to enable the server to automatically increase the value for `max memory`, if needed, when it starts. See `allow memory grow at startup` [here \[page 21\]](#).

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

1.228.1 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.

1.229 max native threads per engine

Defines the maximum number of native threads the server spawns per engine.

Summary	Description
Default value	50
Valid 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.

i Note

`max native threads per engine` is ignored in threaded mode.

1.230 max 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	Description
Default value	16
Valid values	16 – 100
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

1.231 max network packet size

Specifies the maximum network packet size that can be requested by clients communicating with SAP ASE.

Summary	Description
Default value	2048
Valid 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 `bcpl` 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.

Related Information

[additional network memory \[page 17\]](#)

[default network packet size \[page 55\]](#)

1.231.1 Choosing Packet Sizes

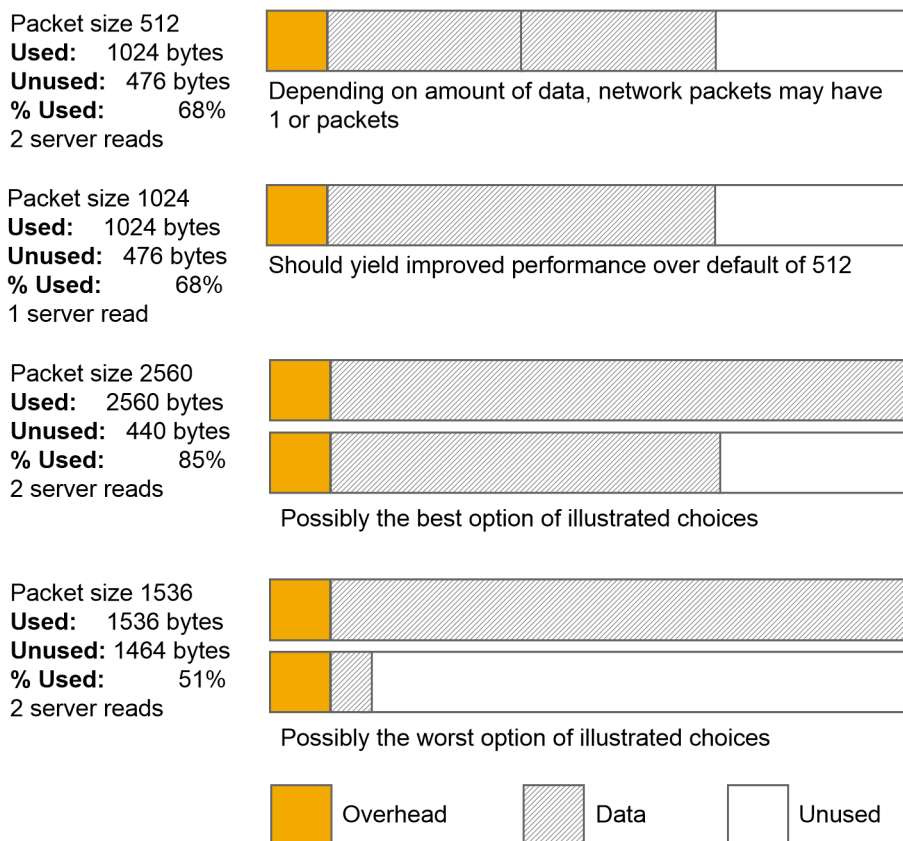
For best performance, choose a server packet size that works efficiently with the underlying packet size on your network.

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.

Underlying network packets: 1500 bytes after overhead



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*.

1.232 max network peek depth (UNIX Only)

Specifies how many levels deep SAP ASE peeks into a connections operating system receive buffer for a pending cancel.

Summary	Description
Default value	0
Valid values	0–2147483647
Status	Dynamic

Summary	Description
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.233 max number network listeners

Specifies the maximum number of network listeners allowed by SAP ASE at one time.

Summary	Description
Default value	5
Valid values	0–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.234 max number of IN elements

Limits the number of elements in the largest `in` clause in a query. If the limit is crossed, SAP ASE reports an error.

Summary	Description
Default value	0 (off)
Valid values	256–128000
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

1.235 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	Description
Default value	1
Valid values	1 – 1024 (threaded mode) 1 – 128 (process mode)
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	Memory Use, Processors

See *System Administration Guide: Volume 2 > Managing Multiprocessor 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*.

1.236 max online q engines

(Process mode only) Specifies the maximum number of Q engines you can have online, and is required for MQ.

Summary	Description
Default value	0
Valid 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, you should 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.

See the *Active Messaging Users Guide*.

Related Information

[max online engines \[page 160\]](#)

[number of engines at startup \[page 190\]](#)

1.237 max parallel degree

Specifies the server-wide maximum number of worker processes allowed per query. This is called the “maximum degree of parallelism.”

Summary	Description
Default value	1
Valid 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*

Related Information

[number of worker processes \[page 212\]](#)

1.238 max_pci_slots

Sets the maximum number of PCI slots SAP ASE allows.

Summary	Description
Default value	0
Valid values	0 – 256
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

Available PCI slots by default:

- Slot number 1 – JVM
 - Module – JVM Adapter
 - Vendor – Sybase an SAP Company
 - Logical name – PCA/JVM
- Slot number 2 – ODBC
 - Module – ODBC Adapter
 - Vendor – Sybase an SAP Company
 - Logical name – PCA/ODBC

For more information about PCI slots, see *Java in SAP Adaptive Server Enterprise*.

1.239 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	Description
Default value	1
Valid 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*.

Related Information

[number of worker processes \[page 212\]](#)

1.240 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	Description
Default value	1
Valid values	1 – value of <code>max parallel degree</code>
Status	Dynamic

Summary	Description
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.

1.241 max resource granularity

Indicates the maximum percentage of the system's resources a query can use.

Summary	Description
Default value	10
Valid 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.

1.242 max scan parallel degree

Specifies the server-wide maximum degree of parallelism for hash-based scans.

Summary	Description
Default value	1
Valid values	1–255

Summary	Description
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 6 worker processes (as long as `max parallel degree` is set to 6 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.

1.243 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	Description
Default value	0

Summary	Description
Valid values	0 – 2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Memory Use, Monitoring

If you do not allocate enough memory for the batch statements, the text you want to view may be truncated. For this reason, you should 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.

1.244 max transfer history

Controls how many transfer history entries SAP ASE retains in the `spt_TableTransfer` table in each database.

Summary	Description
Default value	10
Valid 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.

1.245 max utility 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	Description
Default value	1
Valid values	1– 255
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Query Tuning

1.246 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	Description
Default value	10
Valid values	10–100
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	Group Diagnostics

i 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.

1.247 maximum failed logins

Allows you to set the server-wide maximum number of failed login attempts for logins and roles.

Summary	Description
Default value	0
Valid 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.

1.248 maximum job output

Sets limit, in bytes, on the maximum output a single job can produce.

Summary	Description
Default value	32768
Valid 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.

1.249 memory alignment boundary

Determines the memory address boundary on which data caches are aligned.

Summary	Description
Default value	Logical page size
Valid 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.

i 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.

1.250 memory dump compression level

Controls the compression level for shared memory dumps.

Summary	Description
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.

1.251 memory per worker process

Specifies the amount of memory, in bytes, used by worker processes.

Summary	Description
Default value	1024
Valid 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*.

Related Information

[number of worker processes \[page 212\]](#)

1.252 messaging memory

Configures the amount of memory available for SAP messaging.

Summary	Description
Default value	400

Summary	Description
Valid values	60 – 2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Memory Use, Physical Memory

1.253 metrics elap max

Configures maximum elapsed time and thresholds for QP metrics.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

1.254 metrics exec max

Configures maximum execution time and thresholds for QP metrics.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	Query Tuning

1.255 metrics lio max

Configures maximum logical I/O and thresholds for QP metrics.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

1.256 metrics pio max

Configures maximum physical I/O and thresholds for QP metrics.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

1.257 min pages for parallel scan

Controls the number of tables and indexes that SAP ASE can access in parallel.

Summary	Description
Default value	200
Valid 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.

1.258 minimum password length

Allows you to customize the length of server-wide password values or per-login or per-role password values.

Summary	Description
Default value	6
Valid values	0 – 30
Status	Dynamic
Display level	10
Required role	System security officer
Configuration group	Security Related

See *Security Administration Guide > Set or Change the Minimum Password Length*.

1.259 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	Description
Default value	2
Valid values	0 – 2 <ul style="list-style-type: none">• 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.

1.260 msg_confidentiality_reqd

Requires that all messages into and out of SAP ASE be encrypted.

Summary	Description
Default value	0 (off)
Valid 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 for messages to be encrypted.

Related Information

[use security services \[page 286\]](#)

1.261 msg integrity reqd

Requires that all messages be checked for data integrity.

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

`use security services` must be 1 for message integrity checks to occur.

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`.

Related Information

[use security services \[page 286\]](#)

1.262 net password encryption reqd

Restricts login authentication to use only RSA encryption algorithm or the SAP proprietary algorithm.

Summary	Description
Default value	0
Valid values	0 – 3
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Security Related

Table 2: 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.
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 re-connect 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.

To connect to an SAP ASE server using the `ddlgen` utility when setting `net password encryption reqd` to 1 or 2, use the parameter `ddlgen -C` parameter to specify your connection properties. See `ddlgen` in *Utility Guide*.

i 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.

1.263 network polling mode

Configures the SAP ASE network polling mode.

Summary	Description
Default value	threaded
Valid 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. You should set `network polling mode` to `compact` when there are multiple engine groups, and the load is distributed across the engines.

1.263.1 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.

i 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.

1.264 nonpushdown pipe active

(SAP HANA accelerator for SAP ASE only) Determines if accelerator for SAP ASE collects historical, nonpushdown statement information for SQL statements sent to the SAP HANA server.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

See the [SAP HANA accelerator for SAP ASE](#) documentation for more information.

1.265 nonpushdown pipe max messages

(SAP HANA accelerator for SAP ASE only) Determines the maximum number of messages that can be stored in the `monHANANonPushdown` monitoring table for historical nonpushdown statements for SQL statements sent to the SAP HANA server.

Summary	Description
Default value	0
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

See the [SAP HANA accelerator for SAP ASE](#) documentation for more information.

1.266 number of alarms

Specifies the number of alarm structures allocated by SAP ASE.

Summary	Description
Default value	40
Valid values	40 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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 value of `number of alarms` significantly, adjust `max memory` accordingly.

i Note

In a synchronous replication environment, you must increase the value for `number of alarms` by adding the value to which you set the match the value for `number of user connections` configuration parameter. For example, if the `number of alarms` is set to 40 (the default) and the `number of user connections` is set to 2000, increase the `number of alarms` to 2040.

Related Information

[max memory \[page 154\]](#)

[number of user connections \[page 210\]](#)

1.267 number of aux scan descriptors

Sets the number of auxiliary scan descriptors available in a pool shared by all users on a server.

Summary	Description
Default value	200
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.267.1 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.

Context

`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:

Procedure

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.

1.268 number of backup connections

Sets the maximum number of user connections Backup Server establishes to dump or load in-memory databases.

Summary	Description
Default value	0
Valid 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.

1.269 number of ccbs

Reserved for future use.

Summary	Description
Default value	0
Valid values	0 – 100
Status	Static
Display level	
Required role	
Configuration group	Diagnostics

1.270 number of checkpoint tasks

Configures parallel checkpoints.

Summary	Description
Default value	1
Valid values	1– 8
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Backup/Recovery

The value of number of checkpoint tasks must be less than or equal to the value of number of engines at start-up.

The maximum value is limited by the value of the configuration parameters number of engines at startup and number of open databases, with a maximum of 8. The default value sets serial checkpoints as the default behavior.

Related Information

[number of engines at startup \[page 190\]](#)

1.271 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	Description
Default value	10
Valid values	1–2,147,483,647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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.

1.271.1 Setting the number of devices on UNIX

If you are using a large number of devices on UNIX platforms, you should 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

i 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.
```

1.272 number of disk tasks

Controls the number of tasks dedicated to polling and completing disk I/Os.

Summary	Description
Default value	1
Valid values	UNIX: 1- 64 Windows: 1
Status	Static
	<div data-bbox="820 1769 914 1803" data-label="Section-Header"> <h3>i Note</h3> </div> <div data-bbox="820 1821 1388 1888" data-label="Text"> <p>When increasing the value, the status is dynamic. When decreasing the value, the status is static.</p> </div>
Display level	Basic

Summary	Description
Required role	System administrator
Configuration group	Disk I/O, Processors

You only need to increase the value for `number of disk tasks` above 1 when monitoring indicates the change is necessary.

Internally, when SAP ASE creates tasks and assigns them to `syb_system_pool`, it spawns a thread to run that task. You can configure the number of disk tasks that causes SAP ASE to create more threads.

You should configure multiple disk tasks only if the single disk task becomes CPU bound. When this occurs, you will see a very high (85% or higher) thread utilization for the disk task in `sp_sysmon`. At this level, I/O latency may increase because the single disk task cannot process the completions fast enough. At this point, it is appropriate to add additional task.

However, do not add additional disk tasks because I/O busy in engine utilization is high without first checking the thread utilization for the disk task. A high value for `I/O busy` without a high value for `Thread Utilization` for the disk tasks indicates that the bottleneck is within the host disk subsystem and not within SAP ASE. Adding additional disk tasks does not help.

Unnecessary disk tasks may actually hurt performance, just as having more engines than necessary can hurt performance, due to added spinlock contention, CPU context switching, and CPU cache hit issues that accompany additional threads.

1.273 number of dtx participants

Sets the total number of remote transactions that the SAP ASE transaction coordination service can propagate and coordinate simultaneously.

Summary	Description
Default value	500
Valid values	100 – 2147483647
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	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.

1.273.1 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`.

1.274 number of dump threads

Controls the number of threads that SAP ASE spawns to perform a memory dump.

Summary	Description
Default value	Disabled
Valid values	1 (disabled, no parallelism) – 8 (fully parallel)
Status	Dynamic
Display level	Comprehensive
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.

1.275 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	Description
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.

Related Information

[early row send increment \[page 66\]](#)

1.276 number of engines at startup

Is used exclusively during start-up to set the number of engines brought online.

Summary	Description
Default value	1
Valid values	1 – number of CPUs on machine
Status	Static
Display level	Basic
Required role	System administrator
Configuration group	Java Services, Memory Use, Processors

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`.

i 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.

Related Information

[max online engines \[page 160\]](#)

1.277 number of hcb gc tasks per db

Determines the number of garbage collection tasks assigned to reclaim unused memory from the hash nodes.

Summary	Description
Default value	1
Valid values	1 – 128
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	IMRS Background System Tasks

1.278 number of histogram steps

Specifies the number of steps in a histogram.

Summary	Description
Default value	20
Valid values	3 – 2147483647
Status	Dynamic
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	Query Tuning

1.279 number of imrs gc tasks per db

Determines the default number of IMRS garbage collector tasks for an IMRS-enabled databases or when on-disk MVCC-enabled databases are brought online.

Summary	Description
Default value	2
Valid values	0–32767
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	IMRS Background System Tasks

Increasing the value increases number of tasks in all the databases. However, decreasing the value takes effect when you restart the server.

1.280 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	Description
Default value	0
Valid values	0–65535
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
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*.

i 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`.

1.281 number of java sockets

Enables the Java VM and the `java.net` classes that SAP supports.

Summary	Description
Default value	0
Valid values	0 – 32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Java Services, Memory Use

1.282 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	Description
Default value	6
Valid values	1–256
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

i 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.

1.283 number of lob gc tasks per db

Determines the default number of LOB garbage collector tasks for an IMRS- or on-disk MVCC-enabled databases.

Summary	Description
Default value	2
Valid values	0–32767

Summary	Description
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	IMRS Background System Tasks

Increasing the value increases number of tasks in all the databases. However, decreasing the value takes effect when you restart the server.

1.284 number of locks

Sets the total number of available locks for all users on SAP ASE.

Summary	Description
Default value	10000
Valid values	1000–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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.

i Note

Datarows locking may require that you change the value for `number of locks`. See the *Performance and Tuning Series: Locking and Concurrency Control*.

1.285 number of mailboxes

Specifies the number of mailbox structures allocated by SAP ASE.

Summary	Description
Default value	30
Valid values	30–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.286 number of messages

Specifies the number of message structures allocated by SAP ASE.

Summary	Description
Default value	64
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.287 number of network tasks

Controls the number of tasks dedicated to polling and completing network I/Os.

Summary	Description
Default value	1
Valid values	<ul style="list-style-type: none">(UNIX) 1– 64(Windows) 1
Status	Static
Display level	Basic
Required role	System administrator
Configuration group	Disk I/O, Processors

`number of network tasks` functions only when `network polling mode` is set to `threaded`.

Related Information

[network polling mode \[page 178\]](#)

1.288 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, the longer that aged OAM pages stay in cache.

Summary	Description
Default value	0
Valid values	0–65535
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	

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.

i 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`.

1.289 number of open databases

Sets the maximum number of databases that can be open simultaneously on SAP ASE.

Summary	Description
Default value	12
Valid values	6 –2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Memory Use, Meta-Data Caches, SQL Server Administration

When you calculate a value, include the system databases `master`, `model`, `sybssystemprocs`, 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.

1.289.1 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.

Procedure

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.

1.290 number of open indexes

Sets the maximum number of indexes that can be used simultaneously on SAP ASE.

Summary	Description
Default value	500
Valid values	100–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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.

1.290.1 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.

Procedure

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.

1.291 number of open objects

Sets the maximum number of objects that can be open simultaneously on SAP ASE.

Summary	Description
Default value	500
Valid values	100-2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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.

1.291.1 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.

Procedure

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  Reuse_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
```

Results

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*.

1.292 number of open partitions

Specifies the number of partitions that SAP ASE can access at one time.

Summary	Description
Default value	500
Valid values	100 – 2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Memory Use, Meta-Data Caches

1.292.1 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	Reuse_cnt	Instance_Name	Num_free	Num_active	Pct_act	Max_Used
number of open partitions	0		27 NULL	57	51.8	83

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*.

1.293 number of pack tasks per db

Determines the number of `imrs_pack` threads per IMRS- or on-disk MVCC-enabled database.

Summary	Description
Default value	2
Valid values	0 – 32767
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	IMRS Background System Tasks

Increasing the value for `number of pack tasks per db` also increases the amount of memory the server uses. `number of pack tasks per db` is dynamic when it is increased. That is, if you increase its value, all existing IMRS-enabled databases are immediately assigned more pack tasks. However, decreasing its value is a static action, and requires restarting the server.

1.294 number of pre-allocated extents

Specifies the number of extents (eight pages) allocated in a single trip to the page manager.

Summary	Description
Default value	2
Valid 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.

1.294.1 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.

1.295 number of q engines at startup

Specifies the number of Q engines that are online when the server starts, a requirement for MQ.

Summary	Description
Default value	0
Valid values	0 – 127
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Processors

You may need to increase `max online engines` to accommodate the number of `max online q engines`.

Related Information

[max online engines \[page 160\]](#)

[max online q engines \[page 161\]](#)

1.296 number of reexecutions

Specifies the maximum number of internal re-executions of the DMLs the query processor can attempt when a statement running at `statement snapshot` isolation level receives `write` conflicts.

Summary	Description
Default value	-1
Valid values	-1-2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Query Tuning

Setting the `number of reexecutions` to -1 indicates that the number of reexecutions is unlimited.

The server aborts the statement with error 16873 when it receives more than specified number of `write` conflicts:

```
Msg 16873, Level 16, State 1: Reached maximum internal reexecution threshold (10). Configure 'number of reexecutions' with increased number or DEFAULT (-1).
```

1.297 number of remote connections

Specifies the number of logical connections that can simultaneously be open to and from an SAP ASE.

Summary	Description
Default value	20
Valid values	5-32767
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Memory Use, Network Communication

Each simultaneous connection to XP Server for ESP execution uses up to one remote connection each.

1.298 number of remote logins

Controls the number of active user connections from SAP ASE to remote servers.

Summary	Description
Default value	20
Valid values	0–32767
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	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`.

Related Information

[number of remote connections \[page 207\]](#)

1.299 number of remote sites

Determines the maximum number of remote sites that can simultaneously access SAP ASE.

Summary	Description
Default value	10
Valid values	0–32767
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	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.

1.300 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	Description
Default value	500
Valid values	0–32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

The server uses sort buffers for sort operations like `update statistics`, `create index`, `order by`, `sort merge join`, reformatting (`store_index`) and so on. Each operation can use a number of sort buffers, up to value configured for `number of sort buffers`.

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.

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):

$$(<\text{Number of sort buffers}>) \times (<\text{rows per page}>) \times 100$$

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.

1.301 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	Description
Default value	25
Valid values	5–32767
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Memory Use, User Environment

1.301.1 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 Multiprocessor 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>`.

1.301.2 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 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 the '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
```

1.301.3 Optimizing the Maximum 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`.

i 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.

1.302 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	Description
Default value	0
Valid values	0–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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.

Related Information

[memory per worker process \[page 171\]](#)

1.303 NVCache Lazy Cleaner Pool Size

Determines the size of the NV Cache pool cleaner.

Summary	Description
Default value	512
Valid values	512–2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Non-Volatile Caches

1.304 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	Description
Default value	0
Valid 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>`.

1.305 object lockwait timing

Controls whether SAP ASE collects timing statistics for requests of locks on objects.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

1.306 open index hash spinlock ratio

(Multiprocessing systems only) Sets the number of index metadata descriptor hash tables that are protected by one spinlock.

Summary	Description
Default value	100
Valid values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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*.

1.307 open index spinlock ratio

Specifies the number of index metadata descriptors that are protected by one spinlock.

Summary	Description
Default value	100
Valid values	1–214748364
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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 Multiprocessor Servers*.

1.308 open object spinlock ratio

(Multiprocessing systems only) Specifies the number of object descriptors that are protected by one spinlock.

Summary	Description
Default value	100
Valid values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
Configuration group	Meta-Data Caches

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 Multiprocessor Servers*.

1.309 optimization goal

Determines which optimization goal SAP ASE uses.

Summary	Description
Default value	<code>allows_mix</code>
Valid values	<code>allows_oltp</code> , <code>allows_dss</code>
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:

- `allows_oltp` – the most useful goal for purely OLTP queries.
- `allows_dss` – the most useful goal for operational DSS queries of medium-to-high complexity.

1.310 optimize temp table resolution

Allows stored procedures that reference temporary tables created outside the procedure to not require recompiling for each execution.

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

Trace flag 299 previously supported this functionality.

1.311 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	Description
Default value	10
Valid 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.

1.312 optimize dump for fast load

Optimizes `dump database` commands for a faster `load database` time.

Summary	Description
Default value	0
Valid values	0 – 50
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Backup/Recovery

`optimize dump for faster load` reduces the amount of recovery time required for a `load database` command by including the data that changed while the `dump database` command ran (that is, the dump has more up-to-date contents). The `optimize dump for faster load` default of 0 (no changed data is copied), provides no reduction in the `load database` recovery time.

SAP ASE supports multiple iterations of `dump database` copying pages to the archive. That is, at the end of the first iteration, changes that have been made to the database while the first iteration was active are copied by Backup Server to the archive in a second iteration. The set of changes in each iteration is less than those in the previous iteration, giving `load database` less changes to recover. A maximum of four iterations can be performed.

When `optimize dump for fast load` is set to 0, a single iteration is done.

When `optimize dump for fast load` is set to a non-zero value, it represents the minimum percentage of data page changes (compared to total reserved pages) that have occurred during the `dump database` operation that are added to the dump. SAP ASE dumps these additional changed pages up to four times if the percentage of changed pages exceeds the configuration value. The maximum value of 50 indicates that more than 50% of the database content has changed, and results in much larger dump files and a much longer time to complete the dump operation.

1.313 optimizer level

Determines the level of optimization the query processor uses.

Summary	Description
Default value	<code>ase_default</code>

Summary	Description
Valid values	<ul style="list-style-type: none"> • <code>ase_current</code> – enables all optimizer changes through the current release. • <code>ase_default</code> – disables all optimizer changes since version 1503 ESD #1. • <code>ase1503esd2</code> – enables all optimizer changes through version 15.0.3 ESD #2. • <code>ase1503esd3</code> – 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.

For example, the following command uses the placeholder 0 as the second parameter in the procedure `sp_configure`. This example enables all optimizer changes through the current release by setting the optimizer level to the value of `ase_current`:

```
sp_configure "optimizer_level", 0, "ase_current"
```

See *SAP ASE Performance and Tuning Series: Query Processing and Abstract Plans > Controlling Optimization* for information about optimization levels.

1.314 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	Description
Default value	200
Valid values	2 – 2147483647
Status	Dynamic
Display level	Comprehensive

Summary	Description
Required role	System administrator
Configuration group	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*.

Related Information

[number of locks \[page 195\]](#)

[page lock promotion lwm \[page 220\]](#)

[page lock promotion pct \[page 221\]](#)

1.315 page lock promotion lwm

`page lock promotion lwm` (low-water mark), with `page lock promotion hwm` (high-water mark) and `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	Description
Default value	200
Valid values	2 – value of <code>page lock promotion hwm</code>

Summary	Description
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	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*.

Related Information

[number of locks \[page 195\]](#)

[page lock promotion hwm \[page 219\]](#)

[page lock promotion pct \[page 221\]](#)

1.316 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..

Summary	Description
Default value	100
Valid values	1 – 100
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
Configuration group	Lock Manager, SQL Server Administration
<p>If the number of locks held on an object is between <code>page lock promotion lwm</code> (low-water mark) and <code>page lock promotion hwm</code> (high-water mark)</p> <p>See <i>Performance and Tuning Series: Locking and Concurrency Control > Locking Configuration and Tuning</i>.</p> <p>The default value for <code>page lock promotion pct</code> is appropriate for most applications.</p> <p>You can also configure lock promotion at the object level for page locked objects. See <code>sp_setpglockpromote</code> in the <i>Reference Manual: Procedures</i>.</p>	

Related Information

[page lock promotion hwm \[page 219\]](#)

[page lock promotion lwm \[page 220\]](#)

1.317 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	Description
Default value	95
Valid 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:

$$100 * \text{used pages} / (\text{used pages} + \text{unused pages})$$

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.

1.318 partition groups

Specifies the maximum number of partition groups that can be allocated by SAP ASE.

Summary	Description
Default value	1024
Valid values	1 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.319 partition spinlock ratio

For SAP ASE servers running with multiple engines, sets the number of rows in the partition descriptors that are protected by one spinlock.

Summary	Description
Default value	10
Valid values	1-2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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 Multiprocessor Servers*.

Related Information

[number of open partitions \[page 203\]](#)

1.320 pci memory size

Sets the size of the pluggable component interface (PCI) memory pool.

Summary	Description
Default value	64 MB
Valid values	0 – 2147483647
Status	Dynamic

Summary	Description
Display level	Comprehensive
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.

1.321 per object statistics active

Controls whether SAP ASE collects statistics for each object.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

1.322 percent database for history

Specifies the percentage of the total space available in `sybmgmt.db` that is reserved for the `js_history` table.

Summary	Description
Default value	20
Valid values	0 – 100

Summary	Description
Status	Dynamic
Display level	Comprehensive
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.

1.323 percent database for output

Specifies the percentage of the total space available in `sybmgmt.db` that is reserved for job output.

Summary	Description
Default value	30
Valid values	0 – 100
Status	Dynamic
Display level	Comprehensive
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.

1.324 percent history free

Specifies the percentage of reserved space in `sybmgmt.db` to be kept free.

Summary	Description
Default value	30
Valid values	0 – 100

Summary	Description
Status	Dynamic
Display level	Comprehensive
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 `sybmgmt.db` is filled.

1.325 percent output free

Specifies the percentage of reserved space kept free in `sybmgmt.db` that is reserved for Job Scheduler output.

Summary	Description
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 `sybmgmt.db` is filled.

1.326 performance monitoring option

Enables the license for the BMC DBXray graphical performance monitoring and diagnostic tool .

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

Summary	Description
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

1.327 permission cache entries

Determines the number of cache protectors per task, increasing the amount of memory for each user connection and worker process.

Summary	Description
Default value	64
Valid values	1–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.328 physical lock cushion

`physical lock cushion` allows you to allocate a sufficient number of locks to prevent out-of-lock conditions, which can occur during instance startup, shutdown, failover, or on badly partitioned environments.

Summary Information

Default value	100
Range of values	0 - 100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	Shared Disk Cluster

The value for `physical lock cushion` is a percentage, and can be any value between 0 (the default) and 100.

1.329 plan text pipe active

Determines whether SAP ASE collects query plan text.

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

Related Information

[plan text pipe max messages \[page 230\]](#)

1.330 plan text pipe max messages

Determines the number of query plan text messages SAP ASE stores per engine.

Summary	Description
Default value	0
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

Related Information

[sql text pipe max messages \[page 265\]](#)

1.331 point query rate threshold

Sets the percentage threshold for HCB auto tuning point query rate.

Summary	Description
Default value	30
Valid values	0 – 100
Status	Dynamic

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

A point query rate below the value of `point query rate threshold` indicates index hash caching may offer little improvement because few queries will benefit from using a hash table, and may cause the HCB auto tuning task to disable index hash caching on those indexes.

1.332 prevent automatic upgrade

Allows you to prevent an upgrade that is triggered when starting an older SAP ASE installation (without the `--upgrade-ok` parameter) with a newer version of the SAP ASE `dataserver` binary.

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

A value of:

- 0 – (the default) allows SAP ASE to trigger an upgrade when starting an old installation with a newer version of the SAP ASE `dataserver` binary.
- 1 – (if you do not include the `--upgrade-ok` parameter) any attempt to start SAP ASE with a newer version of the SAP ASE `dataserver` binary fails and prints message 110325:

```
The current database version '%d' is lower than the server's version '%d',
and the server is configured to require upgrade permission. Boot the server
with flag '--upgrade-ok' to permit this boot."
```

1.333 print deadlock information

Prints deadlock information to the error log.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on), 2 (on, print summary)
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	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.
```

1.334 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)

Summary	Description
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.

1.335 procedure cache size

Specifies the size of the procedure cache, in 2 KB pages.

Summary	Description
Default value	7000
Valid values	7000 – 2147483647
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	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.

⚠ Caution

SAP ASE performance degrades if `procedure cache size` is too small.

If you are upgrading, `procedure cache size` is set to the size of the original procedure cache at the time of upgrade.

i Note

SAP recommends that you set the `procedure cache size` parameter to a value less than 1800G.

1.336 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	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

1.337 process wait events

Controls whether SAP ASE collect statistics for each wait event for every task.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Memory Use, Monitoring

You can get wait information for a specific task using `monProcessWaits`.

See *Transact-SQL Users Guide > Using Stored Procedures*.

1.338 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	Description
Default value	20
Valid values	
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

1.339 quorum heartbeat interval

Specifies the number of seconds between quorum heartbeats.

Summary	Description
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.

1.340 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	Description
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.

1.341 quoted identifier enhancements

Enables and disables quoted identifiers use in SAP ASE.

Summary	Description
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

i 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 the *Reference Manual: Blocks*.

Related Information

[enable functionality group \[page 76\]](#)

1.342 rapidlog buffer size

Specifies the buffer size for the output of diagnostic for measuring 'Proc Cache Header' performance.

Summary	Description
Default value	1024
Valid values	1024–2000000
Status	Dynamic
Display level	
Required role	System administrator
Configuration group	Error Log

See *Performance and Turning Series: Basics > Memory Use Performance > Procedure Cache > Diagnostic for 'Proc Cache Header' Memory Pool*.

1.343 rapidlog max files

Specifies the maximum number of files for the output of diagnostic for measuring 'Proc Cache Header' performance.

Summary	Description
Default value	99999999
Valid values	1–99999999
Status	Dynamic
Display level	
Required role	System administrator
Configuration group	Error Log

See *Performance and Tuning Series: Basics > Memory Use Performance > Procedure Cache > Diagnostic for 'Proc Cache Header' Memory Pool*.

1.344 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	Description
Default value	0 (off)
Valid values	0 (off), 1(on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
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*.

1.345 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	Description
Default value	5
Valid values	1-32767
Status	Dynamic
Display level	Basic
Required role	System administrator
Configuration group	Backup/Recovery

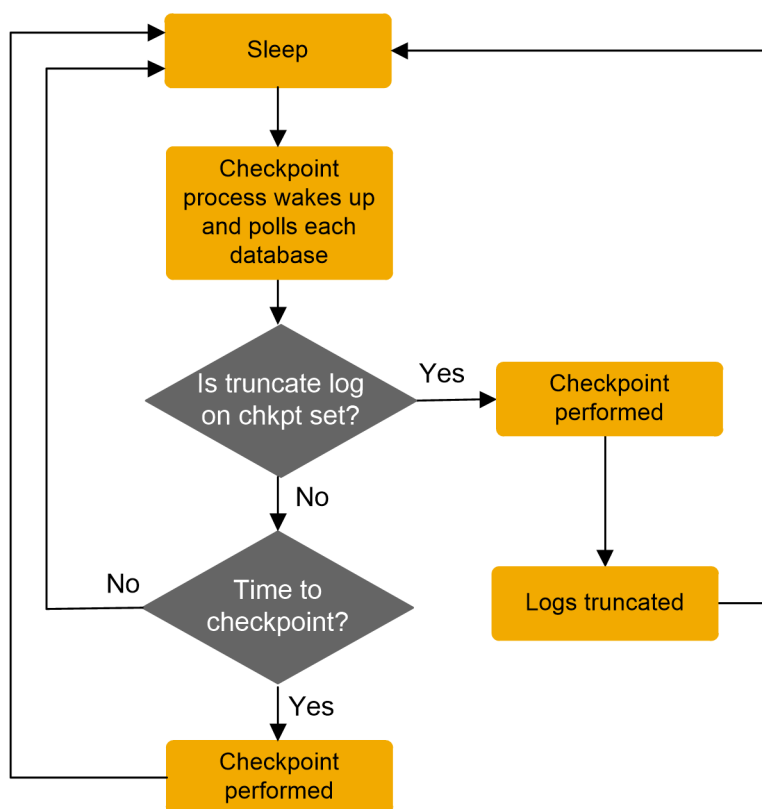
The recovery procedure rolls transactions backward or forward, starting from the transaction that the checkpoint process indicates as the oldest active transaction. The recovery process has more or less work to do, depending on the value of `recovery interval in minutes`.

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.

i 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.



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*.

1.346 recovery prefetch size

Sets the look-ahead size (in numbers of log records) to be used by the recovery prefetch scan.

Summary	Description
Default value	0 (use dynamic prefetch)
Valid values	0 – 20,000
Status	Dynamic

Summary	Description
Display level	Comprehensive
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.

1.347 remote server pre-read packets

Determines the number of packets that are pre-read by a site handler during connections with remote servers.

Summary	Description
Default value	3
Valid values	3–255
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	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.

1.348 replication agent memory size

Determines the amount of memory that SAP ASE allocates to the RepAgent thread pool for a multithreaded RepAgent.

Summary	Description
Default value	4096

Summary	Description
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System Administrator
Configuration group	User Environment

1.349 restore database options

Restores the database options that are set by `create database`, `alter database`, and `sp_dboption` when you load a database or transaction dump.

Summary	Description
Default value	0
Valid values	1, 0
Status	Dynamic
Display level	Comprehensive
Required role	System administrator or <code>manager server configuration</code>
Configuration group	Backup/Recovery

When you enable `restore database options`, all database options in the dump that can be externally set are restored in the target database, except:

- `durability`
- `no_recovery`
- `template`
- `read only`
- `dbo use only`

The encryption status of the database is copied from the dump, regardless of the `restore database options` value. The `allow wide dol rows` option will only be restored if it is enabled in the dump. If `allow wide dol rows` is disabled in the dump, the value set in the target database is used, regardless of the `restore database options` value.

1.350 restricted decrypt permission

Enables or disables restricted decrypt permission in all databases. You must have the `sso_role` to set this parameter.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
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, you should protect data privacy by not granting decrypt permission to the DBO. Access to keys through user passwords prevents the DBO and others from accessing the data unless they have a key's password. It may be convenient, however, to allow 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.

For consistency purposes, you should 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.

1.351 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	Description
Default value	200
Valid values	2-2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	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*.

Related Information

[number of locks \[page 195\]](#)

[row lock promotion lwm \[page 245\]](#)

[row lock promotion pct \[page 246\]](#)

1.352 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	Description
Default value	200
Valid values	2 – value of <code>row lock promotion hwm</code>
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	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*.

Related Information

[number of locks \[page 195\]](#)

[row lock promotion hwm \[page 244\]](#)

[row lock promotion pct \[page 246\]](#)

1.353 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	Description
Default value	100
Valid values	1–100
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	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_setrowlockpromote` in the *Reference Manual: Procedures*.

Related Information

[row lock promotion hwm \[page 244\]](#)

[row lock promotion lwm \[page 245\]](#)

1.354 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	Description
Default value	600
Valid values	600 – 4026531839
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	SQL Server Administration

1.355 rules file

Specifies the location of the `rules` file, which is used by the autotuning feature. SAP ASE internal. Do not change.

1.356 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	Description
Default value	2000 (default value of 3 for the Cluster Edition)
Valid values	0 – 2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

i 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.

i 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*.

1.356.1 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.

1.357 sampling percent

Is the numeric value of the sampling percentage, such as 5 for 5 percent, 10 for 10 percent, and so on.

Summary	Description
Default value	0
Valid 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 `sysabstats` and `sysstatistics` system tables.

1.358 scavenge temp objects

Enables or disables the server from scavenging temporary tables from LRU and MRU page chains.

Summary	Description
Default value	1 (off)
Valid values	1 (off), 0 (on)
Status	Dynamic

Summary	Description
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

Setting `scavenge temp objects` to 1 means that temporary tables are not scavenged from the LRU and MRU page chains. Setting `scavenge temp objects` to 0 means that temporary tables are removed from LRU and MRU page chains.

1.359 secure default login

Specifies a default login for all users who are preauthenticated but who do not have a login in `master..syslogins`.

Summary	Description
Default value	0
Valid 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`.

1.360 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	Description
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*.

i 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.

Related Information

[enable functionality group \[page 76\]](#)

1.361 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	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
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 1 to restrict `select` permission to the object owner and the system administrator.

1.362 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	Description
Default value	1 (on)
Valid 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.

1.363 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 allotted, SAP ASE fails the connection.

Summary	Description
Default value	600
Valid values	0 – 32767
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Shared Disk Cluster

1.364 session tempdb log cache size

Specifies the size for each session `tempdb` log cache.

Summary	Description
Default value	The logical page size, in bytes
Valid 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.

Related Information

[user log cache size \[page 287\]](#)

1.365 shared memory starting address

Determines the virtual address where SAP ASE starts its shared memory region. It is unlikely that you will ever reconfigure this option; do so only after consulting with SAP Technical Support.

Summary	Description
Default value	0
Valid values	Platform-specific
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	Physical Memory

1.366 show deferred compilation text

Enables and disables SAP ASE to display the text of deferred compilation statements as they run.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

1.367 sigstack csmd min size

Configures the minimum amount of stack that is required to handle an address violation signal without stack overflow and captures a configured shared memory dump (CSMD).

Summary	Description
Default value	Platform (all values in bytes) <ul style="list-style-type: none">• Linux: 80000• Sun SPARC: 70000• Solaris AMD: 80000• IBM AIX: 80000• HP-UX: 100000• IBM Linux on Power Systems: 80000
Range of values	Platform (all values in bytes) <ul style="list-style-type: none">• Linux: 30000 – 700000• Sun SPARC: 30000 – 700000• Solaris AMD: 50000 – 800000• IBM AIX: 50000 – 800000• HP-UX: 30000 – 700000• IBM Linux on Power Systems: 50000 – 800000
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

The `sigstack csmd min size` parameter is used to capture a CSMD during signal handling. If the amount of available stack upon entering signal handler is less than the configured value, the task becomes inoperative. If required, you can increase or decrease the stack size.

By default, the `sigstack csmd min size` is enabled. Disable it using trace flag 2913.

i Note

Contact SAP technical support before making changes to these values.

1.368 sigstack min size

Configures the minimum amount of stack that is required to handle an address violation signal without overflowing the stack.

Summary	Description
Default value	Platform (all values in bytes) <ul style="list-style-type: none">• Linux: 30000• Sun SPARC: 30000• Solaris AMD: 35000• IBM AIX: 37000• HP-UX: 80000• IBM Linux on Power Systems: 30000
Valid values	Platform (all values in bytes) <ul style="list-style-type: none">• Linux: 20000 - 400000• Sun SPARC: 10000 - 300000• Solaris AMD: 30000 - 400000• IBM AIX: 30000 - 400000• HP-UX: 30000 - 300000• IBM Linux on Power Systems: 20000 - 400000
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

Use `sigstack min size` to increase the default stack size that SAP ASE uses for a particular platform.

For example, on Linux you can increase the minimum amount of stack required for handling an address violation signal, without overflowing the stack, to a maximum value of 400,000 bytes.

By default, the `sigstack min size` is enabled. Disable it using trace flag 2913.

i Note

SAP recommends that you contact SAP technical support before making changes to these values.

1.369 simplified native access plan

Enables or disables just-in-time compilation of lava execution plans into native code, so that SAP ASE can then invoke these native code plans directly in subsequent executions, allowing for faster execution of extreme

online transaction processing (XOLTP) queries. This compiled queries feature is available as part of MemScale licensed option.

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

i Note

Simplified native access plan is supported only on Linux.

For more information, see *Performance and Tuning Series: Query Processing and Abstract Plans > Understanding Query Processing > Lava Query Execution Engine > Simplified Native Access Plans*.

Related Information

[enable mem scale \[page 90\]](#)

1.370 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	Description
Default value	10
Valid values	1–38
Status	Dynamic
Display level	Intermediate
Required role	System administrator

Summary	Description
Configuration group	SQL Server Administration

The maximum value that can be inserted into an IDENTITY column is $10^{\text{precision}} - 1$. 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.

1.371 size of global fixed heap

Specifies the memory space for internal data structures and other needs.

Summary	Description
Default value	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
Configuration group	Java Services, Memory Use

If you change size of the `global fixed heap`, change the `total logical memory` by the same amount.

Related Information

[total logical memory \[page 280\]](#)

1.372 size of process object heap

Specifies the total memory space for all processes using the Java VM.

Summary	Description
Default value	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 group	Java Services, Memory Use

If you change size of process object heap, change the total logical memory by the same amount.

Related Information

[total logical memory \[page 280\]](#)

1.373 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	Description
Default value	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

Summary	Description
Required role	System administrator
Configuration group	Java Services, Memory Use

If you change the `size of shared class heap`, change the `total logical memory` by the same amount.

Related Information

[total logical memory \[page 280\]](#)

1.374 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	Description
Default value	0
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Memory Use, Unicode

Those using Asian languages 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.

1.375 solaris async i/o mode (Solaris Only)

Allows you to select various asynchronous IO modes.

Summary	Description
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 I/O 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 I/O 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

i 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.

1.376 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	Description
Default value	40

Summary	Description
Valid values	0 – 4000
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Query Tuning

1.377 sql batch capture

Controls whether SAP ASE collects SQL text.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

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.

Related Information

[max sql text monitored \[page 166\]](#)

1.378 sql perfmon integration (Windows Only)

Enables and disables the ability to monitor SAP ASE statistics from the Windows Performance Monitor.

Summary	Description
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.

1.379 sql server clock tick length

Specifies the duration of the server's clock tick, in microseconds.

Summary	Description
Default value	Platform-specific
Valid 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`.

i 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.

Related Information

[cpu grace time \[page 47\]](#)

1.380 sql text pipe active

Controls whether SAP ASE collects SQL text.

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

1.381 sql text pipe max messages

Specifies the number of SQL text messages SAP ASE stores per engine.

Summary	Description
Default value	0
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

1.382 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	Description
Default value	4096
Valid values	0–2147483647
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

i Note

(UNIX only) You should 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.



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.

1.383 stack size

Specifies the size, in bytes, of the execution stacks used by each user process on SAP ASE.

Summary	Description
Default value	Platform-specific
Valid 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, you should increase the default by 50 percent. If you are not using JDBC or CIS, the standard default value is usually sufficient.

1.384 start xp server during boot

Determines whether XP Server starts when SAP ASE starts.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Static
Display level	
Required role	System administrator
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`.

1.385 startup delay

Controls when RepAgent is started during the server start.

Summary	Description
Default value	0 (off)
Valid 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.

1.386 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	Description
Default value	0
Valid values	Size of cache in 2K pages
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Memory Use, SQL Server Administration

i 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.

SAP recommends that you set the `statement cache size` parameter to a value less than 1800G.

1.387 statement pipe active

Controls whether SAP ASE collects statement-level statistics.

Summary	Description
Default value	0 (off)
Valid 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.

Related Information

[statement pipe max messages \[page 270\]](#)

1.388 statement pipe max messages

Determines the number of statement statistics messages SAP ASE stores per engine.

Summary	Description
Default value	0
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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.

Related Information

[sql text pipe max messages \[page 265\]](#)

[statement pipe active \[page 269\]](#)

1.389 statement statistics active

Controls whether SAP ASE collects monitoring table statement-level statistics.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic

Summary	Description
Display level	Comprehensive
Required role	System administrator
Configuration group	Monitoring

Use `monProcessStatement` to get statement statistics for a specific task.

1.390 streamlined dynamic sql

Enables the statement cache to store dynamic SQL statements.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Application Functionality

i 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 *Performance and Tuning Series: Basics > Memory Use and Performance*.

Related Information

[enable functionality group \[page 76\]](#)

1.391 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	Description
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.

1.392 suppress js max task message

Prevents SAP ASE from printing the Job Scheduler `js_maxtask` error messages to the error log.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	10
Required role	System administrator
Configuration group	Application Functionality

i 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`.

Related Information

[enable functionality group \[page 76\]](#)

1.393 suspend audit when device full

Determines what SAP ASE does when an audit device becomes completely full.

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

i 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.

1.394 syb_sendmsg port number

Specifies the port number that SAP ASE uses to send messages to a User Datagram Protocol (UDP) port with `sp_sendmsg` or `syb_sendmsg`.

Summary	Description
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.

i 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*.

Related Information

[allow_sendmsg \[page 24\]](#)

1.395 `sysstatistics flush interval`

Determines the length of the interval, in minutes, between flushes of `sysstatistics`.

Summary	Description
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.

1.396 `systemwide password expiration`

Sets the number of days that passwords remain in effect after they are changed.

Summary	Description
Default value	0
Valid 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, 2017 at 10:30 a.m., with a password expiration interval of 30 days, the password expires on August 31, 2017 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 seven 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 account passwords expire, users can still log in to SAP ASE but cannot execute any commands until they have used `alter login` to change their 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.

1.397 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	Description
Default value	0
Valid 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 seven days, and you attempt to use the tape before seven 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. If you don't provide a value for `tape retention in days` or include the `retaindays` option in the dump command, the `retaindays` option in Backup Server takes effect. See *System Administration Guide: Volume 2 > Backing Up and Restoring User Databases*.

1.398 tcp no delay

Controls Transmission Control Protocol (TCP) packet batching. The default value means that TCP packets are not batched.

Summary	Description
Default value	1 (on)
Valid values	0 (off), 1 (on)
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	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 Tabular Data Stream (TDS) packets may benefit from disabling TCP packet batching.

i Note

Disabling TCP packet batching means that packets are sent, regardless of size; this increases the volume of network traffic.

1.399 text prefetch size

Limits the number of pages of `text`, `unitext`, and `image` data that can be prefetched into an existing buffer pool.

Summary	Description
Default value	16
Valid values	0–65535
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
Configuration group	Network Communications
SAP ASE prefetches only <code>text</code> , <code>unitext</code> , and <code>image</code> data that was created with SAP ASE 12.x or was upgraded using <code>dbcc rebuild_text</code> .	

1.400 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.

Summary	Description
Default value	0
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Memory User, Monitoring

1.401 threshold event monitoring

Enable or disables SAP ASE from recording threshold events.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
Configuration group	Monitoring

1.402 time slice

Sets the number of milliseconds that the SAP ASE scheduler allows a task to run.

Summary	Description
Default value	100
Valid 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*.

1.403 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	Description
Default value	0
Valid values	0 – 2147483647
Status	Calculated

Summary	Description
Display level	Basic
Required role	System administrator
Configuration group	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:
 - number of devices
 - number of open databases
 - number of open indexes
 - number of open objects
 - number of user connections
 - procedure cache percent
 - total logical memory

A number of other parameters also affect the amount of available memory, but to a lesser extent.

Related Information

[number of devices \[page 185\]](#)

[number of open databases \[page 198\]](#)

[number of open indexes \[page 200\]](#)

[number of open objects \[page 201\]](#)

[number of user connections \[page 210\]](#)

[procedure cache size \[page 233\]](#)

[total logical memory \[page 280\]](#)

1.404 total logical memory

Displays the total logical memory for the current configuration of SAP ASE.

Summary	Description
Default value	N/A
Valid values	N/A

Summary	Description
Status	Read-only
Display level	Intermediate
Required role	System administrator
Configuration group	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.

Related Information

[size of process object heap \[page 259\]](#)

[size of shared class heap \[page 259\]](#)

[total physical memory \[page 281\]](#)

1.405 total physical memory

Is a read-only configuration parameter that displays the total physical memory for the current configuration of SAP ASE.

Summary	Description
Default value	N/A
Valid 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`.

Related Information

[max memory \[page 154\]](#)

[total logical memory \[page 280\]](#)

1.406 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	Description
Default value	4096
Valid values	0 – 2147483647
Status	Dynamic
Display level	Intermediate
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.

1.407 txn to pss ratio

Determines the total number of transaction descriptors that are available to the server.

Summary	Description
Default value	16
Valid values	1 – 2147483647
Status	Static

Summary	Description
Display level	1
Required role	System administrator
Configuration group	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.

1.407.1 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`.

1.408 unified login required

Requires that all users who log in to SAP ASE be authenticated by a security mechanism.

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

The `use security services` parameter must be 1 to use the unified login security service.

Related Information

[use security services \[page 286\]](#)

1.409 update statistics hashing

Enables SAP ASE to gather hash-based statistics.

Summary	Description
Default value	off

Summary	Description
Valid values	One of: <ul style="list-style-type: none"> • <code>off</code> – no hashing • <code>on</code> – hashing on all columns • <code>partial</code> – hashing only for low unique count columns • <code>default</code> – 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
```

1.410 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	Description
Default value	15000
Valid values	0–2147483647
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	SQL Server Administration

Caution

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"
```

1.411 use security services

Specifies that SAP ASE uses network-based security services.

Summary	Description
Default value	0 (off)
Valid 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.

1.412 user log cache queue size

Determines whether a queuing strategy is used for logging. The user log cache is divided into multiple cachelets or not based on the value you set in this configuration parameter.

Summary	Description
Default value	1
Valid values	0, 1, an integer larger than 1 <ul style="list-style-type: none">0 – disables queuing for user log caches. A single cachelet is used for each user log cache based on the <code>user log cache size</code> in bytes. The user log cache may be flushed prior to a transaction commit due to internal SAP ASE processing such as buffer unpin operations associated with datarows locking tables.1 – enables queuing for user log caches. Multiple cachelets are used for each user log cache based on

Summary	Description
	<p>the <code>user log cache size</code> in bytes and the server page size. A minimum of 4 cachelets is required. If the <code>user log cache size</code> is not big enough, SAP ASE uses a single cachelet, which voids the optimization. When internal SAP ASE processing such as buffer unpin operation takes place, it activates a new log cachelet and the previous log cachelets are linked to the global database queue that is flushed with the next transaction to commit. This reduces the contention on the log semaphore due to buffer unpin operations associated with datarows locking tables.</p> <ul style="list-style-type: none"> an integer larger than 1 – internal use only. Do not set the value higher than 1 except at the advice of SAP Technical Support.
Status	Static
Display level	Comprehensive
Required role	System administrator
Configuration group	User Environment

1.413 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	Description
Default value	Logical page size
Valid values	2048 ^a – 2147483647 a. Minimum determined by server's logical page size
Status	Static
Display level	Intermediate
Required role	System administrator
Configuration group	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.

i 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*.

1.414 user log cache spinlock ratio

For SAP ASE servers 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	Description
Default value	20
Valid values	1–2147483647
Status	Dynamic
Display level	Intermediate
Required role	System administrator
Configuration group	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*.

1.414.1 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, <code>manage server configuration</code> is required. When granular permissions is disabled, <code>sa_role</code> 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`.

1.415 wait event timing

Controls whether SAP ASE collects statistics for individual wait events.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator

Summary	Description
Configuration group	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.

1.416 wait on uncommitted insert

Allows you to control the wait behavior of `select`, `update`, `insert`, and `delete` commands for an uncommitted insert.

Summary	Description
Default value	0 (off)
Valid values	0 (off), 1 (on), 2 (wait on insert)
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	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. The table cannot have a nonclustered unique index.

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 2, the behavior is the same as value 1, with the exception of `insert` also adheres to `wait on uncommitted insert`.

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.

Examples Showing the Behavior For the `insert` and `update` Commands

Example 1:

The following example shows the behavior for the `insert` command:

1. Set the value of configuration `wait on uncommitted insert` to 2.

```
1> sp_configure 'wait on uncommitted insert', 2
2> go
```

2. Create `test_table` and a unique index.

```
1> create table test_table (id int, name varchar(20)) lock datarows
2> go
1> create unique index id_index on test_table(id)
2> go
```

3. Do the following insert, and then rollback the insert into session #1.

```
# session #1
1> begin tran
2> go
1> insert into test_table values (1, 'test1')
2> go
(1 row affected)
```

```
# session #2
1> insert into test_table values (1, 'test1')
2> go
<waiting>
```

```
# session #1
1> rollback
2> go
```

```
# session #2
(1 row affected)
```

The insert in session #2 waited on the uncommitted insert in session #1. The insert in session #2 succeed once the insert in session #1 was rolled back.

4. Do the following insert and commit the insert in session #1.

```
# session #1
1> begin tran
2> go
1> insert into test_table values (2, 'test2')
2> go
(1 row affected)
```

```
# session #2
1> insert into test_table values (2, 'test2')
2> go
<waiting>
```

```
# session #1
1> commit
2> go
```

```

# session #2
Msg 2601, Level 16, State 6:
Line 1:
Attempt to insert duplicate key row in object 'test_table' with
unique index 'id_index'
Command has been aborted.
(0 rows affected)

```

The insert in session #2 waited on the uncommitted insert in session #1. The the insert in session #2 reported error 2601 once the insert in session #1 was committed.

Example 2:

The following example shows the behavior for the update command:

1. Set the value of configuration `wait on uncommitted insert` to 2.

```

1> sp_configure 'wait on uncommitted insert', 2
2> go

```

2. Create `test_table` and a unique index.

```

1> create table test_table (id int, name varchar(20)) lock datarows
2> go
1> create unique index id_index on test_table(id)
2> go

```

3. Do the following insert, and then rollback the insert into session #1.

```

# session #1
1> begin tran
2> go
1> insert into test_table values (1, 'test1')
2> go
(1 row affected)

```

```

# session #2
1> update test_table set name = 'first test' where id = 1
2> go
<waiting>

```

```

# session #1
1> rollback
2> go

```

```

# session #2
(0 row affected)

```

The update in session #2 waited on the uncommitted insert in session #1. The data was not updated in session #2 since the insert in session #1 was rolled back.

4. Do the following insert and commit the insert in session #1.

```

# session #1
1> begin tran
2> go
1> insert into test_table values (2, 'test2')
2> go
(1 row affected)

```

```

# session #2

```

```
1> update test_table set name = 'second test' where id = 2
2> go
<waiting>
```

```
# session #1
1> commit
2> go
```

```
# session #2
(1 row affected)
```

The update in session #2 waited on the uncommitted insert in session #1. The data was updated in session #2 since the insert in session #1 was committed.

1.417 workload manager cache size

Specifies the maximum amount of memory, in 2 KB pages, that the workload manager can use.

Summary	Description
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*.

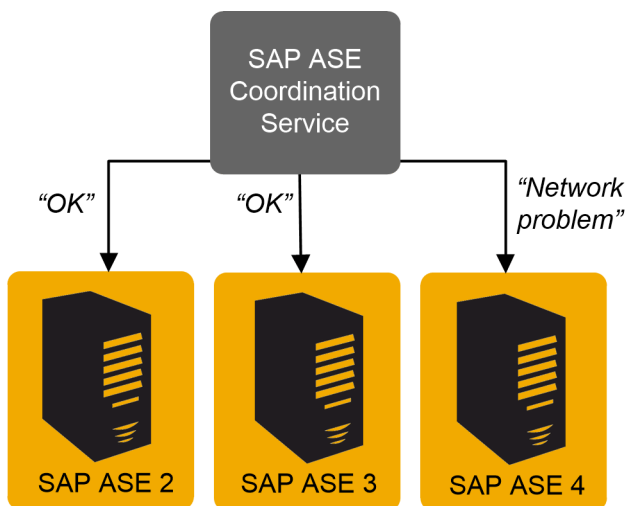
1.418 xact coordination interval

Defines the length of time between attempts to resolve transaction branches that have been propagated to remote servers.

Summary	Description
Default value	60 (seconds)
Valid values	1 – 2147483647 (seconds)
Status	Dynamic

Summary	Description
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 interval`.



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.

1.419 xp_cmdshell context

Sets the security context for the operating system command to be executed using the `xp_cmdshell` system ESP.

Summary	Description
Default value	1
Valid values	0, 1, 2 The values for the context determines under which account the command runs: <ul style="list-style-type: none">• 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.
Status	Dynamic
Display level	Comprehensive
Required role	System administrator
Configuration group	Extended Stored Procedure

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 requester 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.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Gender-Related Language

We try not to use gender-specific word forms and formulations. As appropriate for context and readability, SAP may use masculine word forms to refer to all genders.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

© 2020 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.