Help Index

The Help Index procedure is part of our ***Audit Suit***e. This procedure allows experienced DBAs to administer the tables and indices in a database by obtaining a quantitative and qualitative inventory at the index level. The procedure is driven by the entries in `sysindexes` (actual discrete data storage structures within *Sybase* ASE [1]); thus it lists all data storage structures. It is particularly useful before and after making changes to the database objects (e.g. rebuilding or de-fragmentation of the database or a cluster of tables; improving indexing; implementing a tranche of performance improvements), to verify that the intended changes were made. A quantitative summary is provided at the end; the show/hide options determine the summary rows displayed, but not the summary values.

The report lists all data storage structures (indices and heaps) in the specified database. The columns provided are based on switches: reports for specific purposes may be obtained. The columns in the report are logically divided, first into table and index column groups. The index columns are further divided into four column groups: index; performance; fragmentation; and usage (activity) characteristics.

```
[exec] [db_name..]sp_helpindex_sg [-defFkmnprstu]
```

`Table` Table Name
`Lock` Lock scheme: {`APL`|`DPL`|`DRL`}
`Index` Index Name
`Type` Data storage structure type:
- `Heap` the structure is a data heap
- `Clst` the structure is an APL Clustered Index
- `Cl-G` the structure is a Global Clustered Index
- `Cl-L` the structure is a Local Clustered Index
- `Plac` the structure is a DPL/DRL Placement Index [1]
- `NCxx` the structure is a Nonclustered Index
- `Text` the structure is the single text/image chain for the table

`Un` Index is Unique [3]
`Cn` Constraint:
- `PK` Index is a Primary Key constraint [2]
- `Un` Index is a Unique Key constraint [2]
- `pk` Index is a declared primary key [3]

`FK` Index supports a Foreign Key constraint [4]
`Columns` Key columns

Additional columns (in the relevant position) are provided, based on switches:

`d` Show DML usage statistics: (default is Hide) [5] [6]
`Inserts`
`Updates`
`Deletes`

`e` Show error columns [7] (and summary rows):
Table errors:
- `xI` No Index
- `xUI` No Unique Index
- `xCI` APL/No Clustered Index

---

1 An APL Clustered index (`Clst|Cl-G|Cl-L`) is a single entry in `sysindexes`, reflecting the single data storage structure (the index is 'clustered' with the data, the leaf level of the index is the data row); a DPL/DRL Placement Index (`Plac`) and its data `Heap` are two separate entries in `sysindexes` reflecting two separate data strorage structures, each of which is administered separately. The DPL/DRL Placement Index is an ordinary Nonclustered Index (the leaf level is a pointer to the data row in the Heap) with the additional criteria of (a) identifying the preferred (not actual) placement of data rows in the Heap, and (b) being placed on the same segment as the `Heap`; it is not 'clustered' with the data rows. In order to avoid the confusion introduced by evangelists, and consistent with *Sybase* Engineering, we use the term 'Placement', not 'Clustered'.

2 This separation shows clear identification of indices created via `create index` vs `constraint unique` or `constraint primary key` syntax; and allows continued administration under such separation.

3 Regardless of how the index is created, the key can be identified in standard-compliant manner via `sp_primarykey`.

4 The current entry is the unique index in the parent table; the Foreign Key supported is in the child tables (not the current entry).

5 For an APL Clustered index, which is a single entry, the figure shown is the aggregate for the index and the data.

6 If Software Gems' `monTable` is in place , it is used (showing true aggregate values); else if monitoring (MDA) is enabled, it is used (showing current snapshot values); else the request is abandoned.

7 The single column allows a single error to be identified; where there is more than one error, the worst error is shown (when it is corrected them worst remaining error is shown, etc).

Index errors:

- `DUP` Duplicate Index (the column set is a duplicate of another Index/column set)
- `dup` Duplicate Index (the column set is a subset of another Index/column set)
- `xUn` APL/Non-Unique Clustered Index
- `xUs` Index Unused by Optimiser (requires {`d|m|u`} Switch)

`f` Show fragmentation and statistics columns [8]:

- `Fwd` Forwarded rows (`DPL`/`DRL`)
- `Del` Deleted rows (`DPL`/`DRL`)
- `DPCR` Data Page Cluster Ratio (`Heap` & `Clst`)
- `IPCR` Index Page Cluster Ratio (`Plac` & `NCI`)
- `DRCR` Data Row Cluster Ratio (`Plac` & `NCI`)
- `LGIO` Large I/O Efficiency
- `SPUT` Space Utilisation (`Heap & Clst`)
- `Delta` Index: `MAX( DATACHANGE() )` of *all* columns in the *Index*
  Heap: `DATACHANGE()` of all partitions in the Heap
- `StatUpdt` Date statistics updated (oldest across partitions)

`F` Provide Flat File (Spreadsheet) format (default is nice)

`k` Show space usage in Kilobytes (default is Megabytes)

`m` Show Monitor stats: (default is Hide) [5] [6]

- `LogcReads` Logical Reads
- `PhysReads` Physical Reads
- `APFReads` Asynchronous Pre-Fetch Reads
- `PhysWrites` Physical Writes

`n` Hide unpopulated tables (default is Include)

`p` Show performance columns (and summary rows):

Table performance columns:

- `Fix` Fixed length columns
- `Var` Variable length columns [9]
- `Trg` Existence of `[I][U][D]` trigger
- `Id` Table has an `IDENTITY` Column

Index performance columns:

- `Var` Variable length columns in Index
- `Id` Index is an `IDENTITY` Column
- `SO` Sorted Object
- `SD` Index was created with Sorted Data option
- `SI` Sorted Index
- `Lx` Large I/O (APF) disabled
- `Mx` MRU Strategy disabled
- `AI` Ascending Inserts
- `Ptn` Number of Partitions
- `Segment` Segment Name or Number of Segments
- `Siz` Actual Data (`Clst/Heap`) or Leaf (`NCI`) row size [10]

`r` Show space usage columns (and summary rows):

- `Rows` Table row count
- `Kb/Mb` Space usage in Kilo/Megabytes for the Index/Heap [5] [11] [12]

`s` Show system tables (default is Hide)

`t` Show text/image chains (default is Hide)

`u` Show Optimiser selection and usage statistics [6] [13] for Heap or Index (default is Hide):

- `OptSelect` `monOpenObjectActivity.OptSelectCount`
- `OptUsed` `monOpenObjectActivity.UsedCount`

(0 in both columns indicate object is neither Selected nor Used)

• This is an example of one component of our ***Audit Suite***. The intention is to demonstrate the quality and maturity of our scripts.

• As consultants performing either Technical Audit or Performance Enhancement exercise, it is the first report that would be obtained; it comprises a baseline; and readily identifies areas for address at the index level.

---

8 *Sybase* function return values are provided: an update count as a delete plus an insert, thus the Percentage figure may be absurd. Percentages are constrained to 999 for sanity.

9 Indicates columns that could be Fixed to improve performance; excludes Text/Unitext/Image columns, which reside in a single text chain.

10 Absolute size for fixed rows; average size for variable rows.

11 The actual space usage returned by `reserved_pages()` for the object is shown (not the possible space usage if the object was de-fragmented).

12 If Mb is chosen, small values are shown in Kb prefixed with a slash.

13 In this instance, since the index (not data heap) selection figure is sought, the APL Clustered index figure is shown (ie, it is not aggregated with the heap figure).

Table	Lock	Fix	Var	Trg	Id	Rows	Fwd	Del	Err	Index	Type	Un	Cn	FK	Var	Id	Siz	Mb	Err	Columns
z_sessions_web	APL	4	6	IUD		0				z_sessions_web_pkey	Clst	Un	PK		1		61	0		session_id
z_settings	APL	5	6	IUD		437				z_settings_pkey	Clst	Un	PK		1		95	0		id
										tz_settings	Text							0		
z_settings_category	APL	1	3	IUD		25				z_settings_category_pkey	Clst	Un	PK		1		74	0		id
										tz_settings_category	Text							0		
z_settings_tools	APL	11	9	IUD		433				z_settings_tools_pkey	Clst	Un	PK		1		152	0		id
										tz_settings_tools	Text							1		
z_spandex	APL	1	1			100				z_spandex_pkey	Clst	Un	PK				12	0		numcount
z_sql_error_codes	APL	1	3			38			xCI		Heap						25	0		
										z_sql_error_code_pkey	NC1	Un	PK		1		19	0		dbms,sqldbcode
z_sql_error_stack	APL	3	2			0				z_sql_error_stack_pkey	Clst	Un	PK		1		151	0		session_id,app_number,stack_inst
z_strings	APL	5	6	IUD		5817				z_strings_pkey	Clst	Un	PK		1		171	1		site_prefix,language,key_string
z_system_checkin	APL	1				1				z_system_checkin_p	Clst	Un	PK				10	0		last_checkin
z_system_messages	DRL	2	1			1	0	0			Heap						29	0		
										z_system_messages_p	Plac	Un	PK		1		9	0		person_id
										tz_system_messages	Text							0		
z_system_performance_data	APL		11	IUD		0			xI		Heap						243	0		
z_system_performance_run	APL	2	4	IUD		32				z_system_perform_run_pkey	Clst	Un	PK	FK	1		88	0		workstation_id,when_run
z_system_performance_score	APL	2	3	IUD		522				z_system_perform_score_pkey	Clst	Un	PK		2		7			
z_system_performance_test	APL	2	5	IUD		32				z_system_perform_test_pkey	Clst	Un	PK	FK	1		8			
z_table_logging	APL	2	1	IUD		342				z_table_logging_pkey	Clst	Un	PK		1		2			
z_table_upgrade	APL	1	9	IUD		21165			xCI		Heap						5			
										z_table_upgrade_ADF_p	NC1	Un	PK		2		3			
z_table_upgrade_new	APL	4		IUD		10			xCI		Heap						82	0		
										z_table_upgrade_new_pkey	NC1	Un	PK				67	0		dbtable,config_id
z_temporary	DRL		2			6	0	0			Heap						38	0		
										z_temporary_pkey	NC1	Un	PK		1		32	0		variable_name
z_transtemp	APL		3			0				z_transtemp_cluster	Clst	Un			2		171	0		session_id,name
z_txn_detail	APL	1	16			0				z_txn_detail_clustindx	Clst				2		626	0	xUn	actual_table,key1
										z_txn_detail_sessionindx	NC1				1		14	0		session_id
z_txn_header	APL	2	10			0				z_txn_header_pkey	Clst	Un	PK		4		120	0		session_id,txn,top_table,key1
										z_txn_header_when_indx	NC1				1		13	0		when_filed
z_validation_rule	APL	2	6	IUD		1				z_validation_rule_pkey	Clst	Un	PK		1		171	0		rule_id
z_validations	APL		4	IUD		22				z_validations_p	Clst	Un	PK		1		157	0		id
z_voice_profiles	APL	2	9	IUD		1			xCI		Heap						58	0		
										z_voice_profiles_pkey	NC1	Un	PK		1		16	0		profile_id
										z_voice_profiles_person_indx	NC2	Un			3		22	0		person_id,workstation_id,profile_id
										z_voice_profiles_mspk_indx	NC3	Un			3		39	0		medspeak_username,person_id,profile_id
										z_voice_profiles_name_indx	NC4	Un			2		30	0		name,profile_id

- This is an example of the report produced by the ***Help Index*** script, from a recent assignment (used with permission).
- This shows a selection of possible columns.

Copyright © 2009 Software Gems Pty Ltd

Copyright © 2009 Software Gems Pty Ltd

Table	Lock	Fix	Var	Trg	Id	Rows	Fwd	Del	Err	Index	Type	Un	Cn	FK	DPCR	IPCR	DRCR	LGIO	SPUT	Siz	Mb	Err
z_webconfig	APL	1	6			0			xCI		Heap				100.00			100.00	0.00	439	0	
										z_webconfig_pkey	NC1	Un	PK		100.00	0.00	100.00		40	0	
z_workstation_group	APL	1	5	IUD		19				z_workstation_group_pkey	Clst	Un	PK		0.00			12.50	58.18	62	0	
z_workstation_group_printer	APL	1	3	IUD		71				z_workstation_groupprint_pkey	Clst	Un	PK		0.00			12.50	83.23	24	0	
z_workstation_group_setting	APL	2	6	IUD		0				z_wsgroup_setting_pkey	Clst	Un	PK		100.00			100.00	0.00	308	0	
z_workstations	APL	10	16	IUD		332				z_workstations_p	Clst	Un	PK		50.00			22.22	57.96	116	0	
z_workstations_addons	APL	1	4			278				z_workstations_addons_pkey	Clst	Un	PK		71.43			33.33	53.97	31	0	
zi_dbma_connect	APL	3	4			0			xCI		Heap				100.00			100.00	0.00	80	0	
										zi_dbma_connect_pkey	NC1	Un	PK		100.00	0.00	100.00		17	0	
										tzi_dbma_connect	Text									0	
zi_dbma_customer_db	APL	2	10	IUD		0				zi_dbma_customer_db_pkey	Clst	Un	PK		100.00			100.00	0.00	223	0	
										tzi_dbma_customer_db	Text									0	
zi_dbma_log	APL	4	1			0			xCI		Heap				100.00			100.00	0.00	34	0	
										zi_dbma_log_pkey	NC1	Un	PK		100.00	0.00	100.00		21	0	
										tzi_dbma_log	Text									0	

Metric	APL	DPL	DRL	Total
Table	707		49	756
Table Column	6185		376	6561
Table Text/Image Column	229		18	247
Table Foreign Key	119			119
Index	963		66	1029
Index Unique	722		48	770
Index Primary Key Constraint	678		37	715
Index Support Foreign Key	54			54
Index Column	1549		125	1674
Index Clst	611			611
Index Heap	96		49	145
Index Plac			42	42
Index NC1	280		12	292
Index NC2	46		8	54
Index NC3	13		3	16
Index NC4	4		1	5
Index NC5	2			2
Index NC6	2			2
Index NC7	2			2
Index NC8	2			2
Index NC9	1			1
Text/Image Chain	182		18	200

Performance	APL	DPL	DRL	Total
Table VAR LENGTH	701		47	748
Table Trigger	574		29	603
Table IDENTITY	14		19	33
Table Column VAR LENGTH	4519		228	4747
Index IDENTITY	14		19	33
Index Column VAR LENGTH	1448		80	1528

Error		APL	DPL	DRL	Total
Table NO INDEX	xI	12		1	13
Table NO UNIQUE INDEX	xUI	1			1
Table APL NO CLUST INDEX	xCI	96			96
Index APL NONUNIQUE CLUST INDEX	xUn	135			135
Index DUPLICATE (full column set)	DUP	2			2
Index DUPLICATE (column subset)	dup	1			1

Size	APL	DPL	DRL	Total
Data Size Mb	24841		203	25044
Index Size Mb	10267		82	10349
Text/Image Size Mb	13562		1112	14674

- This shows a different selection of possible columns.
- The summaries on the last pages are shown.
- Index Errors and issues that affect Performance and shown separately