

Document Map

Select any link item to go directly to the webpage.

Sybase ASE Architecture v2.0



Genuine architectural definition of *Sybase ASE*. in 10 pages: all important components & how they relate to other components; definition of both the Process Kernel & Threaded Kernel. Explains common performance problems; componentry; etc.

Purpose

These documents are intended to complement the *Sybase* manuals, and to correct them, as follows:

- they contain information that is not in the manuals (ie. they overcome the lack of information)
- where the manuals contain contradictory information, the correct version *only*. is provided, *the goal is to eliminate confusion and half-truths!*
- where misleading or false technical terms are used, *correct technical terms are used instead*
- they bring all the relevant information about a subject together, in one place.

Sybase ASE Product Manuals

These documents do not replace the product manuals, or repeat the info therein. Of course, the latter remains the product reference.

Oracle vs Sybase ASE v1.0



For those who think freeware/vapourware/nowhere "sql" "platforms" and Oracle (it has the same non-architecture, except that it is commercial) are servers, and thus are prevented from learning genuine server architecture, this is a pre-requisite.

Data Storage & Fragmentation v2.5



- three documents combined
- references & links integrated
- higher quality vector graphics

Data Storage V2.5

- A **Sybase Data Storage** Defines how data is stored, and the DataStructures used
 - 1 **Data Storage Unit** Defines all Data Storage Units; the hierarchy; their relations; Control Pages
 - 2 **DataStructure Introduction** Introduces the five DataStructures; Partitions; and their relations
 - 2.1 **Segment** Introduces Segment in context
 - 2.2 **Device** Introduces Device in context
 - 3 **DataStructure Definition** Defines each DataStructure comparing APL & DPL/DRL forms, and how it relates to other DataStructures
 - 3.1 **Heap** Defines the APL vs DPL/DRL form of the Heap, fresh and fragmented
 - 3.2 **Clustered Index** Defines the Clustered Index (APL only), and its advantages
 - 3.3 **Nonclustered Index** Defines the Non-Clustered Index, and it relation to Clustered Index (APL) vs Heap/Placement Index (DPL/DRL)
 - 3.4 **Placement Index** Defines the Placement Index(DPL/DRL only), and it relation to the Heap
 - 3.5 **Text/Image Chain** ()
 - 4 **Catalogue** Identifies the catalogue pertaining to Data Storage
 - 5 **Data Model** Explains DataStructures and their relations, in terms of a Data Model

Fragmentation V2.5

- B **Sybase Fragmentation** Defines the types of fragmentation, in three Levels; their determination; causes; correction; and prevention
 - Part A Data Storage is a pre-requisite*
 - 6 **Definition** Introduces the types of fragmentation, and the three Levels at which it occurs
 - 6.1 **Impact** Defines the impact of fragmentation, at each Level
 - 6.2 **Fragmentation Type** Defines types of fragmentation; section of DataStructure affected; best & worst conditions; effect; correction
 - 7 **Determination** Defines accurate determination; all relevant metrics; metrics that should be ignored
 - 7.1 **Level I** All metrics for Level I fragmentation. **Customers only**
 - 7.2 **Level II Space** Basic space metrics for Level II fragmentation
 - 7.3 **Level II DerivedStat** Derived Stat metrics for Level II fragmentation
 - 7.4 **Level III** Metrics for Level III fragmentation
 - 7.5 **Evaluation** DataStructure fragmentation is three levels taken together
 - 7.6 **Partition** Determination of fragmentation metrics for partitioned DataStructures
 - I **Level I - Allocation Unit** Defines Level I fragmentation under three headings
 - 8 **Allocation Unit** Defines fragmentation at AllocationUnit & Extent levels, identifies causes
 - 8.1 **Fresh** Illustrates a fresh AllocationUnit
 - 8.2 **Fragmented** Illustrates a fragmented AllocationUnit
 - 8.3 **DataStruct Perspective** Illustrates an AllocationUnit from the perspective of a DataStructure
 - 9 **Drop-Create Clustered Index** Illustrates why Drop-Create CI does not return Asynch Pre-Fetch & Large I/O
 - 9.1 **Common Issue** Defines the common misunderstanding, why Drop-Create CI does not produce expected results
 - 9.2 **BCP-Out, Drop** Illustrates the result of bcp-out, followed by Drop Clustered Index
 - 9.3 **BCP-In, Create CI** Illustrates the result of bcp-in, followed by Create Clustered Index
 - 9.4 **Drop, Create CI** Defines correct method for de-fragmenting a Clustered Index
 - 9.5 **Drop, Create PI** Illustrates the result of bcp-in, followed by Create Placement Index
 - 9.6 **Next Create CI** Illustrates the effect of the *next* Create Clustered Index
 - 10 **Segment** Defines method for correction & prevention of Level I fragmentation: Segments
 - 10.1 **Normal Growth** Illustrates normal growth of DataStructures within Segments
 - 10.2 **Fragmented** Illustrates fragmentation of DataStructures within Segments
 - 10.3 **Fresh** Illustrates the result of de-fragmentation of DataStructures within Segments
 - 11 **Level I Summary** Summarises Level I fragmentation, and de-fragmentation operations
 - II **Level II - PageChain & Unused Space** Defines Level II fragmentation under four headings
 - 12 **Page Chain** Defines PageChain fragmentation
 - 12.1 **Fresh** Illustrates the fresh Clustered Index PageChain (APL) vs Heap/Placement Index (DPL/DRL)
 - 12.2 **Fragmented** Illustrates the fragmented Clustered Index PageChain (APL) vs Heap/Placement Index (DPL/DRL)
 - 12.3 **Effect/Range Query & Table Scan** Illustrates the effect of Range Queries & Table Scans on the fragmented Clustered Index
 - 12.4 **Effect/Covered Query** Illustrates the effect of a Covered Query on the NCI (APL) and NCI or PI (DPL/DRL)
 - 13 **Overflow Page** Illustrates the non-unique Clustered Index vs Row Forwards in the Heap/Placement Index (DPL/DRL)
 - 14 **Unused Space/Extent** Illustrates Unused Space in the Extent, for CI & NCI (APL) vs Heap/PI & NCI (DPL/DRL)
 - 14.1 **Effect** Defines the effect of Unused Space in the Extent, APL vs DPL/DRL
 - 15 **Unused Space/Page** Illustrates Unused Space in the Page, for CI (APL) vs Heap (DPL/DRL)
 - 15.1 **Effect** Defines the effect of Unused Space in the Page, for CI (APL) vs Heap (DPL/DRL)
 - 16 **Level II Summary** Summarises Level II fragmentation, and de-fragmentation operations
 - III **Level III - Page** Defines this level of fragmentation (DPL/DRL only) in different scenarios, compared with APL
 - 17.1 **Clustered Index vs Heap & Placement Index** Illustrates a fresh Clustered Index (APL) vs Heap/Placement Index (DPL/DRL)
 - 17.2 **Next Sequential Insert** Illustrates effect of next-sequential INSERT into Clustered Index (APL) vs Heap/PI (DPL/DRL)
 - 17.3 **Interspersed Insert/Space** Illustrates effect of interspersed INSERT into CI (APL) vs Heap/PI (DPL/DRL), when space available
 - 17.4 **Interspersed Insert/No Space** Illustrates effect of interspersed INSERT into CI (APL) vs Heap/PI (DPL/DRL), when no space available
 - 17.5 **Interspersed Delete** Illustrates effect of interspersed DELETE into CI (APL) vs Heap/PI (DPL/DRL)
 - 17.6 **Interspersed Update** Illustrates effect of interspersed UPDATE (expand row) into CI (APL) vs Heap/PI (DPL/DRL)
 - 17.7 **Page Fragmentation** Illustrates no fragmentation in CI (APL) Page vs fragmentation in Heap/PI (DPL/DRL) Page
 - 18 **Level III Summary** Summarises Level III fragmentation, and de-fragmentation operations for Heap/PI (DPL/DRL)

Index Type v2.5

- C **Index Type** (Older doc) Comparison of Index Types per Lockscheme on one page

Cache Manager v2.4



- Sybase Cache Behaviour** Illustrated exposé of all aspects of cache operation; referencing sysmon metrics
 - 1 **Introduction** The essentials of cache operation, and an overview
 - 2 **Wash Area** Dual use, delayed writes, disk writes, and considerations
 - 3 **Normal Strategy** LRU strategy & considerations
 - 4 **Fetch & Discard Strategy** MRU strategy & considerations
 - 5 **Overload** Overload conditions & determination
 - 6 **Large I/O** Importance and use; focus on Reading
 - 6.1 **Large I/O Write** Focus on writing using Large I/O, under each strategy
 - 7 **Cache Partition** Configuration for high concurrency
 - 8 **Asynch Pre-Fetch** Configuration for high throughput; Causes of under-utilisation; final notes

Statistics Demystified v2.5



- Sybase Statistics Demystified** A complete, illustrated, guide to *Sybase ASE* statistics. Magic, replaced by science.
 - 1 **Introduction** Purpose; Document Layout; Status
 - 2 **Concept** Which every DBA must be thoroughly familiar with
 - 2.1 **Statistic Type** All possible 'stats' types identified, actual Stats defined
 - 2.2 **Column Stats** Definition of Stats for a single Column
 - 2.2.1 **Density**
 - 2.2.2 **Distribution**
 - 2.3 **Histogram** Definition of the Histogram
 - 2.3.1 **Histogram Type**
 - 2.3.2 **Cell Content**
 - 2.3.3 **Range Cell**
 - 2.3.4 **Frequency Cell**
 - 2.3.5 **Partitioned Table**
 - 2.3.6 **Automatic Tuning**
 - 2.4 **ColumnGroup** Definition of Stats for a group of Columns
 - 2.4.1 **ColumnGroup/Indexed**
 - 2.4.2 **ColumnGroup/Indexed Stats**
 - 2.4.3 **ColumnGroup/Not Indexed**
 - 2.4.4 **ColumnGroup/Not Indexed Stats**
 - 2.5 **optdiag Stats** Explanation os optdiag reports
 - 2.6 **Summary** Summary and requirement for maintenance of Stats
 - 2.6.1 **Summary**
 - 2.6.1 **Cell Requirement**
 - 3 **Catalogue** A deeper understanding, essential for writing your won Stats reports
 - 3.1 **The Logical** The logical tables in *sysstatistics*
 - 3.2 **The Physical** The physical *sysstatistics* table
 - 3.3 **Row** Definition of all row types
 - 3.4 **Statistics** Definition of Stats row
 - 3.5 **Histogram** Definition of Histogram row
 - 4 **Update Statistics** Detailed explanation, and effect, of each flavour of the command
 - 4.1 **Example** Definition of a moderately complex example, that is used throughout
 - 4.2 **Update Stats - Column** UPDATE STATISTICS for a single column
 - 4.3 **Update Stats - ColumnGroup** UPDATE STATISTICS for a group of columns (unindexed)
 - 4.4 **Update Stats - Index** UPDATE INDEX STATISTICS
 - 4.5 **Update Stats - Index (Deprecated)** UPDATE STATISTICS for an index
 - 4.6 **Update Stats - Table** UPDATE TABLE STATISTICS
 - 4.7 **Update Stats - All** UPDATE ALL STATISTICS
 - 5 **Goal** Directions to achieve maximum effect in minimum window (customers only)
 - 6 **SG Report** *Software Gems'* Statistic reports, showing how to structure info with logical grouping
 - 6.1 **Concept** Concept of the reports
 - 6.2 **HelpStatistic** *Software Gems'* HelpStatistic report
 - 6.2.2 **HelpStatistic - Partition** Partition Stats displayed using the HelpStatistic report
 - 6.3 **HelpHistogram** *Software Gems'* HelpHistogram report
 - 6.3.2 **Advantage over optdiag** Advantage of this form, over optdiag
 - 6.3.3 **Skew Value** Handling skew values in the Histogram, displayed using the HelpHistogram report
 - 6.3.4 **Repeated Value** Handling repeated values, displayed using the HelpHistogram report

Lock Manager v2.9



- Sybase Lock Manager** Fully illustrated guide to *Sybase ASE* Locking
 - 1 **Lock Escalation** Determination of all lock types & how they relate; Isolation; keywords
 - 2 **Lock Promotion** Illustrates lock promotion; the modulation of isolation in context; the effect of qualifiers
 - 3 **Configuration** Defines all contention resolution mechanisms from the perspective of resource usage; contention & resolution

MDA Collector v1.4



MDA Collector



MDA Collector Source

RAID



RAID5 vs RAID1+0

