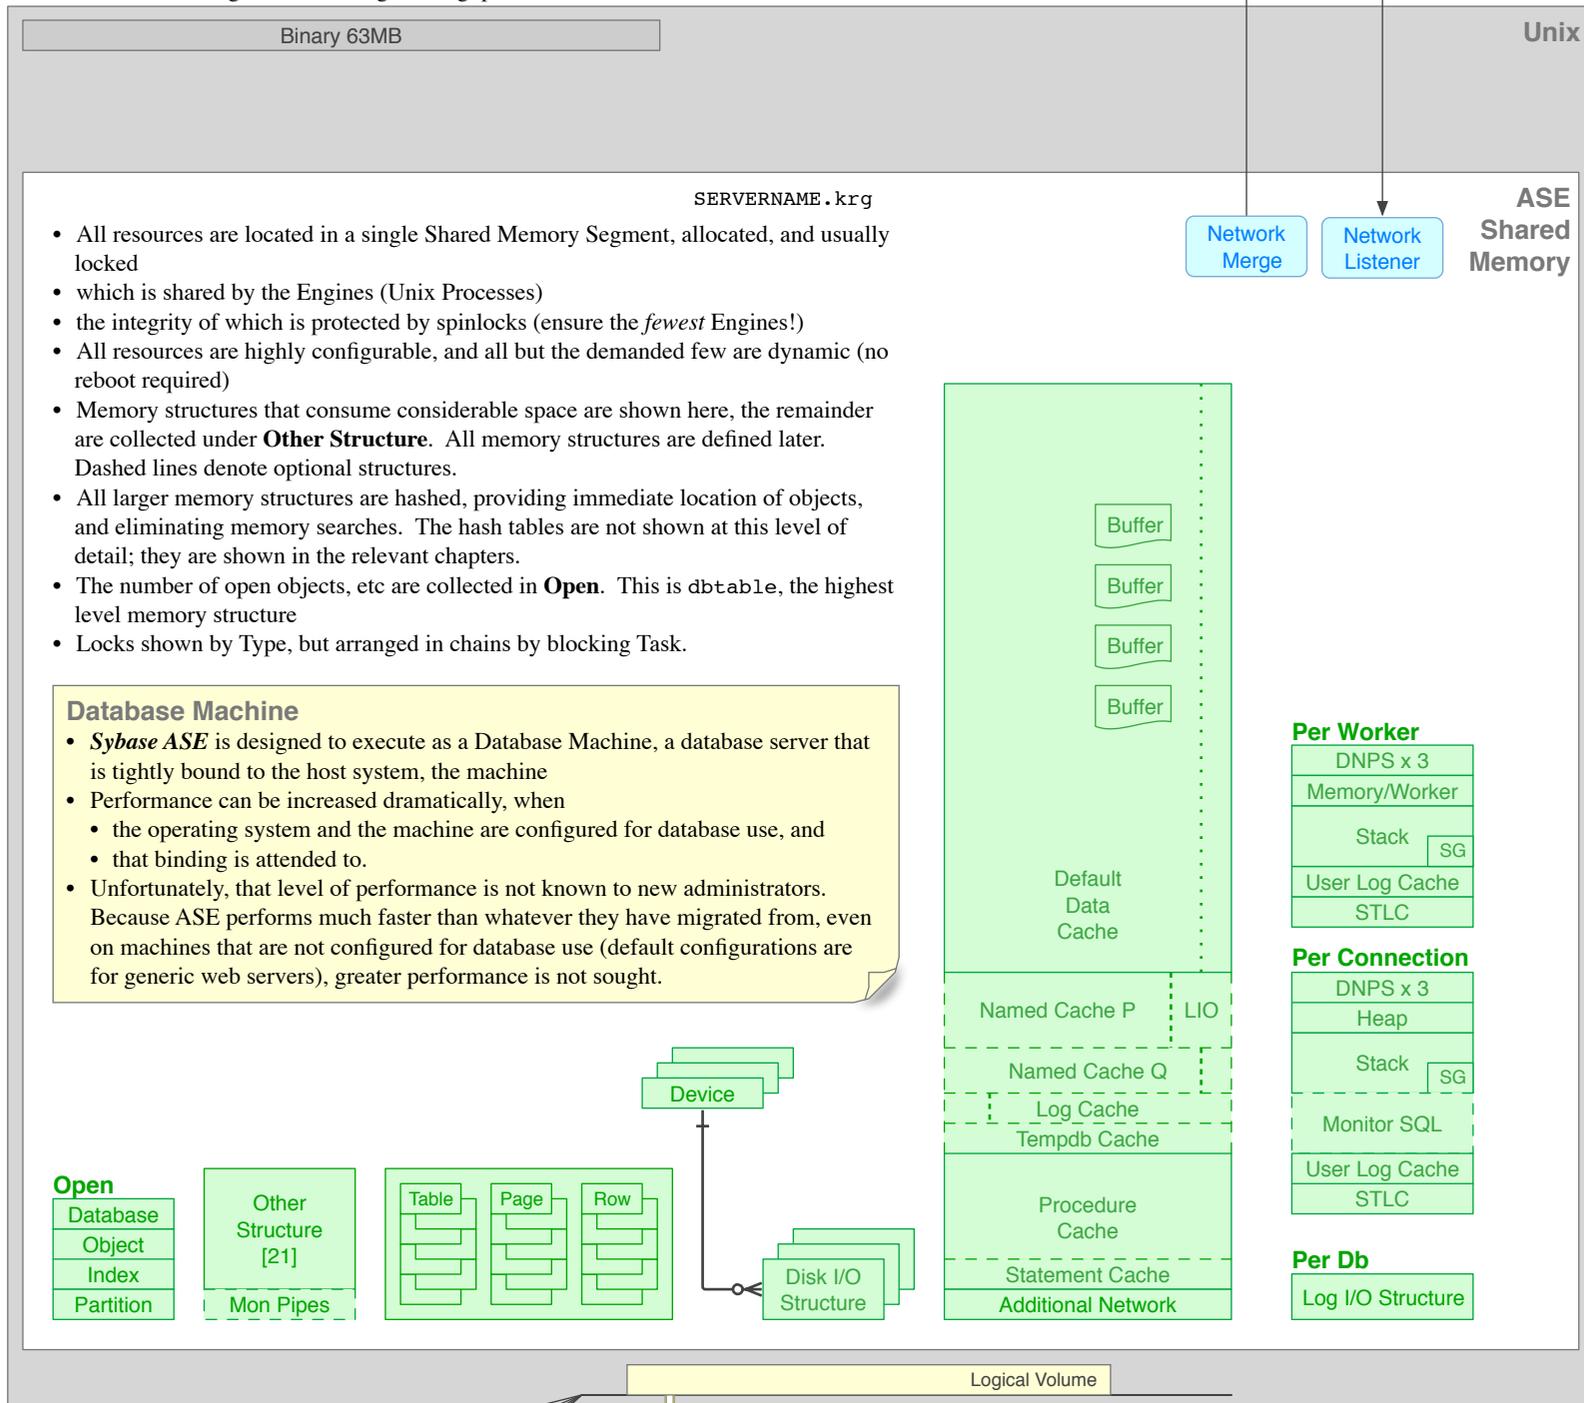
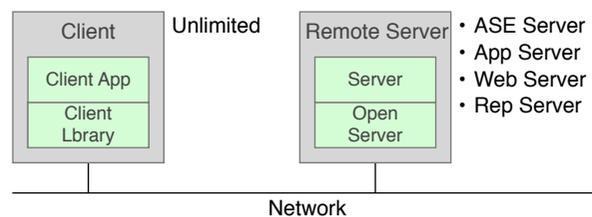


Sybase ASE Architecture Foundation



Besides being SQL-compliant, *Sybase Adaptive Server Enterprise* is a genuine database server. Unfortunately these days, with cooks and bottle-washers writing postulant "servers", the word *genuine* is demanded. Like Unix, and vastly different to alleged "servers", it:

- consists of a single binary, that runs as pure Multi-Threaded¹ code
- has a Kernel and a Scheduler, which are tightly bound to the o/s and the chipset
- Engines that execute as the *fewest possible* Unix Processes
- designed architected as a database machine, and operates the best on a machine that is thusly suited and configured for the high throughput of a database server

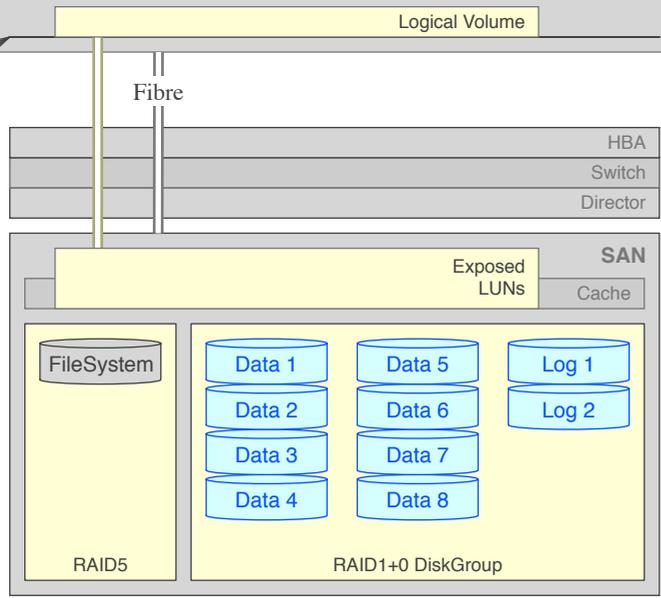


- SERVERNAME.krg
- All resources are located in a single Shared Memory Segment, allocated, and usually locked
 - which is shared by the Engines (Unix Processes)
 - the integrity of which is protected by spinlocks (ensure the *fewest* Engines!)
 - All resources are highly configurable, and all but the demanded few are dynamic (no reboot required)
 - Memory structures that consume considerable space are shown here, the remainder are collected under **Other Structure**. All memory structures are defined later. Dashed lines denote optional structures.
 - All larger memory structures are hashed, providing immediate location of objects, and eliminating memory searches. The hash tables are not shown at this level of detail; they are shown in the relevant chapters.
 - The number of open objects, etc are collected in **Open**. This is dbtable, the highest level memory structure
 - Locks shown by Type, but arranged in chains by blocking Task.

Database Machine

- *Sybase ASE* is designed to execute as a Database Machine, a database server that is tightly bound to the host system, the machine
- Performance can be increased dramatically, when
 - the operating system and the machine are configured for database use, and
 - that binding is attended to.
- Unfortunately, that level of performance is not known to new administrators. Because ASE performs much faster than whatever they have migrated from, even on machines that are not configured for database use (default configurations are for generic web servers), greater performance is not sought.

- Backplane/IOController/Hub
- PCI Card & Extender
- Fibre Channel
- Local/Remote Device
- Ensure ASE has a clear path to the data
- Ensure the fewest servers compete, at **each layer** in the hardware stack *and* the software stack
- RAID1+0 is recommended over any other
- A Private DiskGroup for each is essential
- Raw Partitions are far superior to /fs files
- *Especially* for tempdb (noting the absurdity of recommendations otherwise)
- Mirroring, to either a local or remote SAN, is provided, as an alternative to SAN-level mirroring



- Physical I/O is the slowest link in the performance chain.
- SAN Configuration is very important for any database server, it is the subject of a 50-page document for customers.

Green Memory Component
 Blue Function
 DNPS Default Network Packet Size
 STLC Session Tempdb Log Cache
 SG Stack Guard

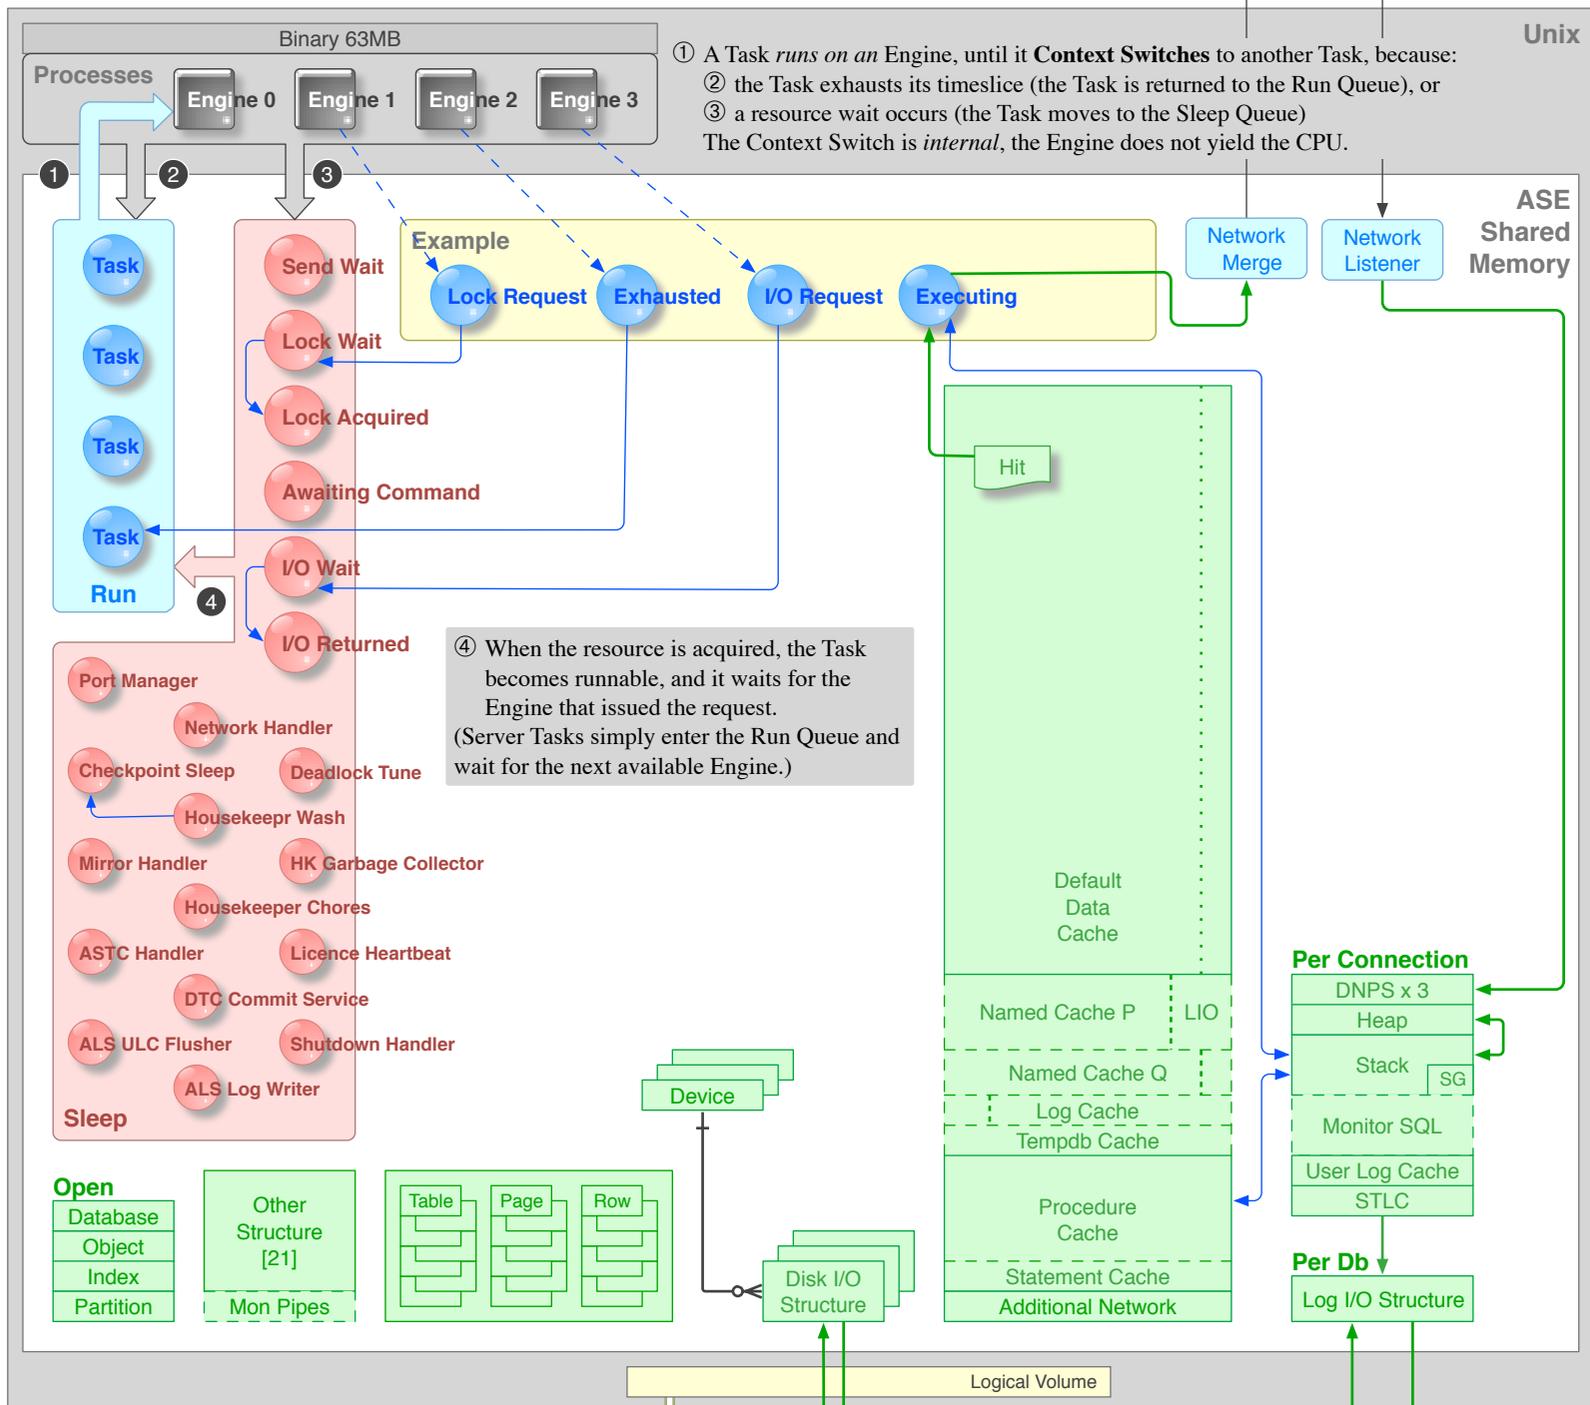
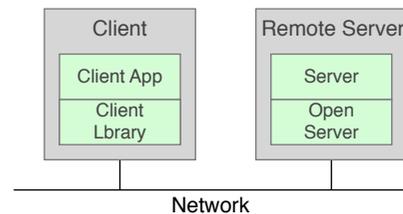
¹ The technical term established since 1969, not the private definition used in Oracle documentation. Contrast with Chip Multi-Threading.

Sybase ASE Architecture

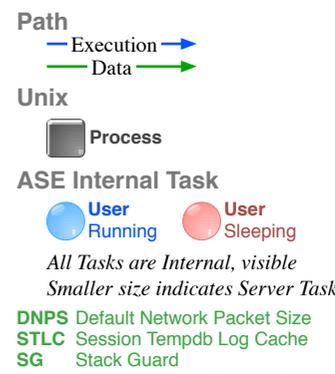
Task Context



- Engines execute as the Unix Processes, which should be the *fewest possible* for the load. A Server with four Engines configured is shown
- User Tasks (for each connection and worker) and a few server Tasks (for independent functions), operate as internal processes
- The Example shows states of Tasks, that have been achieved while executing
- 3D articles are live objects. 2D articles are the more fixed context within which they operate



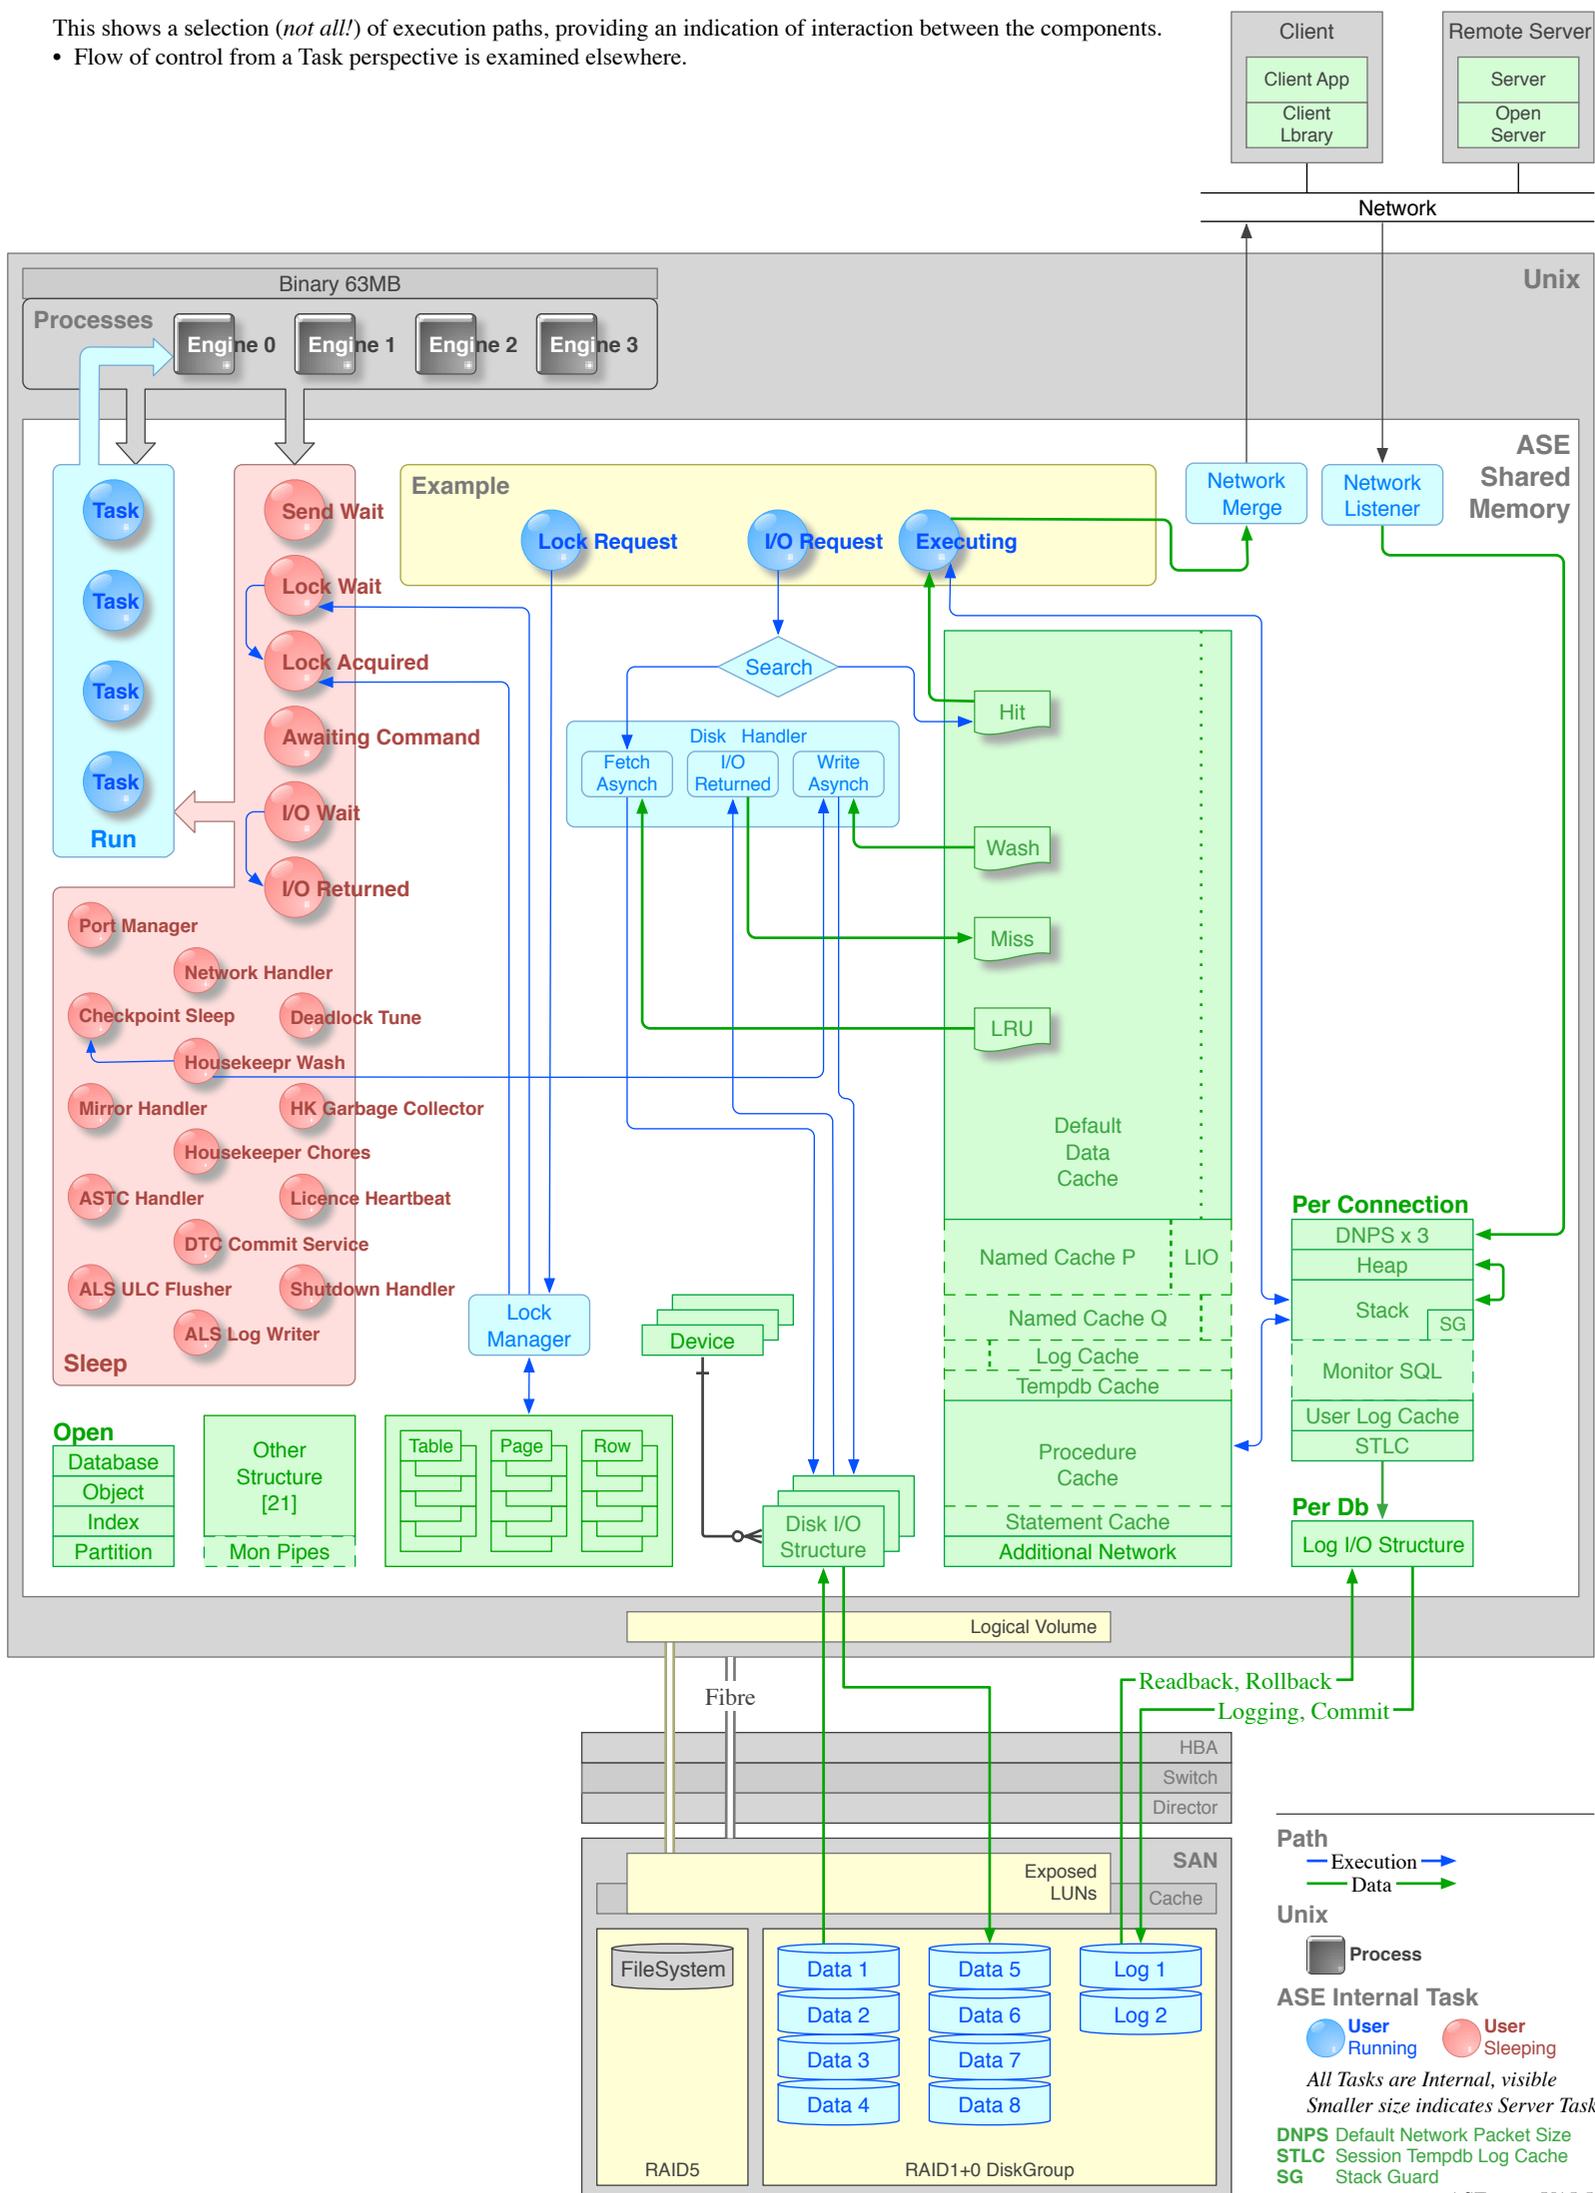
- Engine Available**
Initiating (for resource wait) or any
- Exhausted**
- Context Switch**
Cache Miss
Disk Device
Disk System Write
Exceed BatchSize
Lock Normal (Logical Lock)
Lock APL Index (Address Lock)
Lock Latch
Lock LW Protection (Modify Conflicts)
Log Group Commit
Log IO Structure (Log Semaphore)
Log Last Page Write
Network Receive
Network Send
Network Sleep (Network Services)
Other
User Log Cache (PLC Lock)
Voluntary Yield
- Scheduled**
Or when resource wait is complete



Sybase ASE Architecture Execution



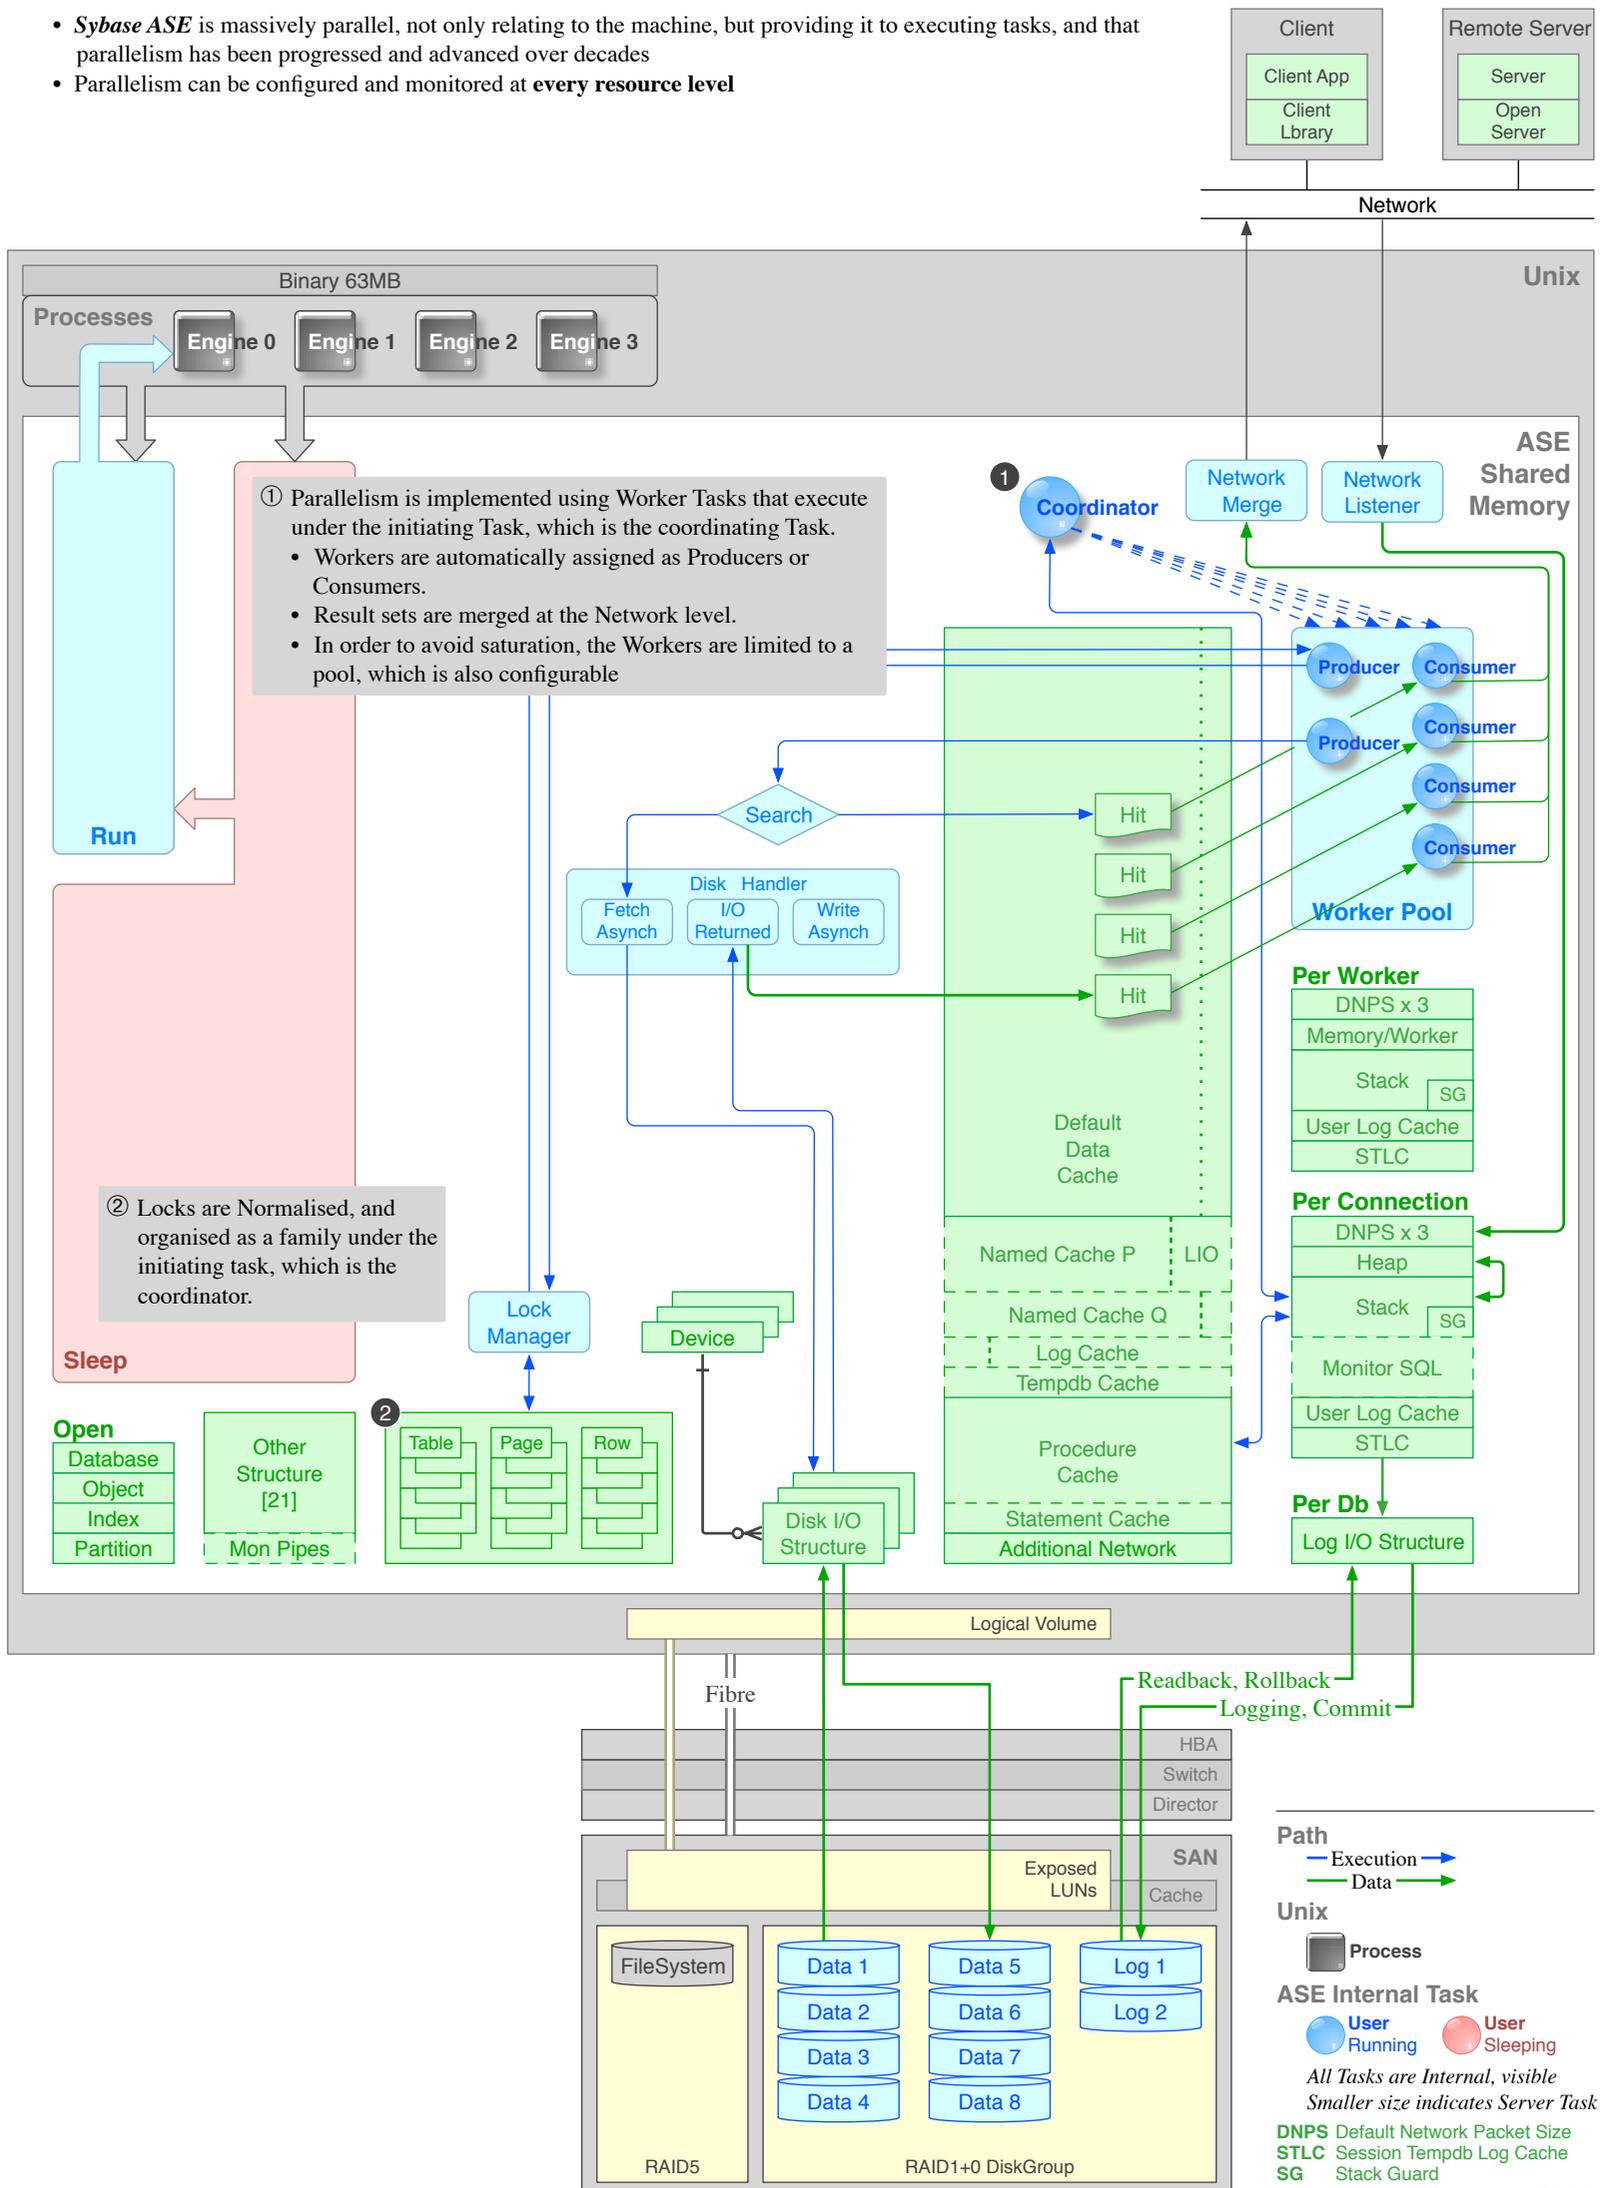
This shows a selection (*not all!*) of execution paths, providing an indication of interaction between the components.
 • Flow of control from a Task perspective is examined elsewhere.



Sybase ASE Architecture Parallelism

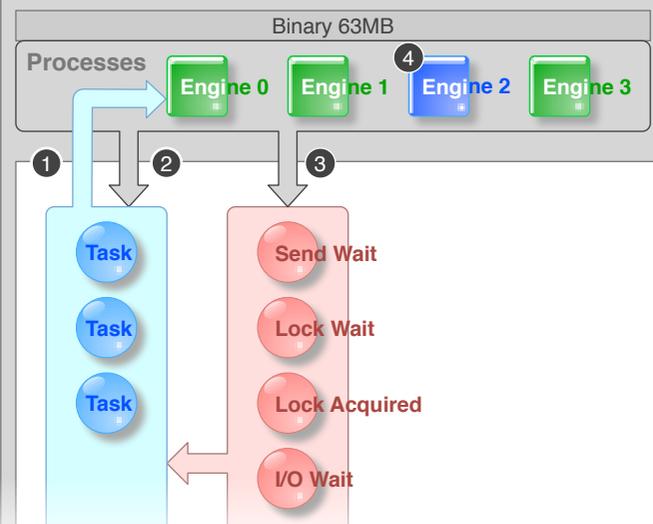


- **Sybase ASE** is massively parallel, not only relating to the machine, but providing it to executing tasks, and that parallelism has been progressed and advanced over decades
- Parallelism can be configured and monitored at **every resource level**



This section explains a commonly misunderstood issue, and the compounding negative effects of incorrect machine or ASE configuration.

- Note that ASE sees only logical CPUs (whatever Cores or Threads that have been configured on the host system), and it will take full advantage of max online engines.



① A Task runs on an Engine, until it Context Switches, because:

- ② the Task exhausts its timeslice, or
- ③ a resource wait occurs

The Context Switch is *internal*, the Engine does not yield the CPU. If the Engine finds a task that it can run, the cycle continues. If not, it holds the CPU and waits for one. The method used is:

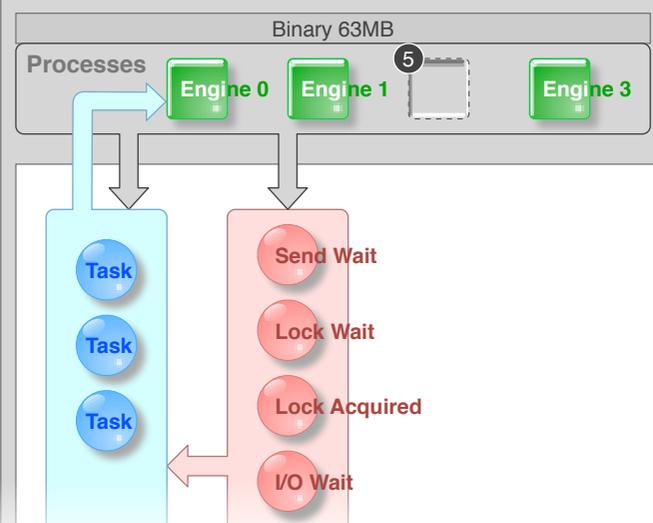
- ④ it loops `runnable process search count` times.

Eventually

- ⑤ the Engine yields the CPU.

Unix

ASE Shared Memory



CPU Yield

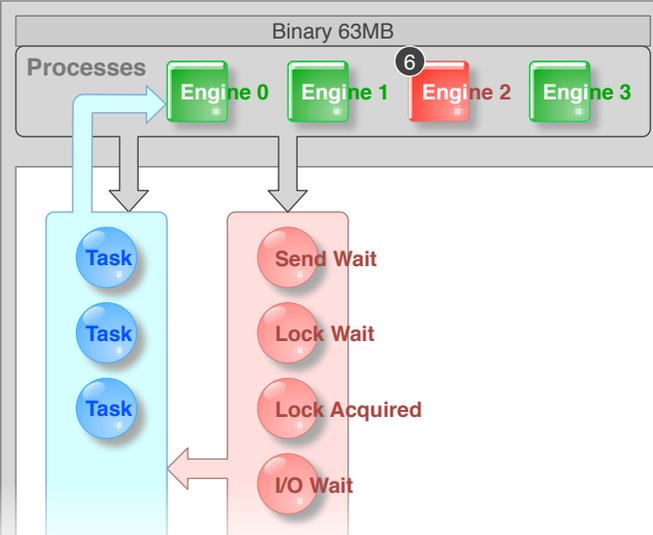
Sybase ASE is designed to execute as a Database Server, the only server on the host system, tightly bound to the machine.

Not Yielding the CPU is a design principle:

- It avoids o/s level Context Switches, which are the most expensive operation at the Unix level. Once the CPU has been yielded, it has to wait until it is scheduled to execute again.
- (which is one reason why a product without a server architecture, is a circus of clowns, bumping into each other, and why a circus needs a machine ten times more powerful than Sybase for the same load).
- Yielding the CPU is undesirable for a high performance database server
- Note that voluntarily yielding the CPU is quite different to being forced off the CPU by Unix, due to other processes (planned or unplanned) running on the system.

Unix

ASE Shared Memory



CPU Monopolisation

Failure to understand this, leads to **two common configuration errors**, especially on large machines, each of which, although negative, may be invisible to the untrained eye, but the combination is highly visible to all: high CPU usage *without* a proportionate increase in throughput:

- `runnable process search count` set too high:
 - ⑥ Engines are prevented from yielding the CPU when there is no work.
- `max online engines` set too high for the load:
 - over-subscription of CPUs, which means many Engines are very busy doing nothing at all.
 - Many Engines are prevented from yielding the CPU.

The two errors in combination results in *many* Engines being busy doing nothing, *and* they are prevented from yielding the CPU.

CPU monopolisation is desirable when the architecture is understood and it is configured appropriately, and it is a disaster when not. *The results of placing people with no tertiary technical qualifications in technical positions are catastrophic.*

Unix

ASE Shared Memory

Goal: Fewest Engines, Not Yielding

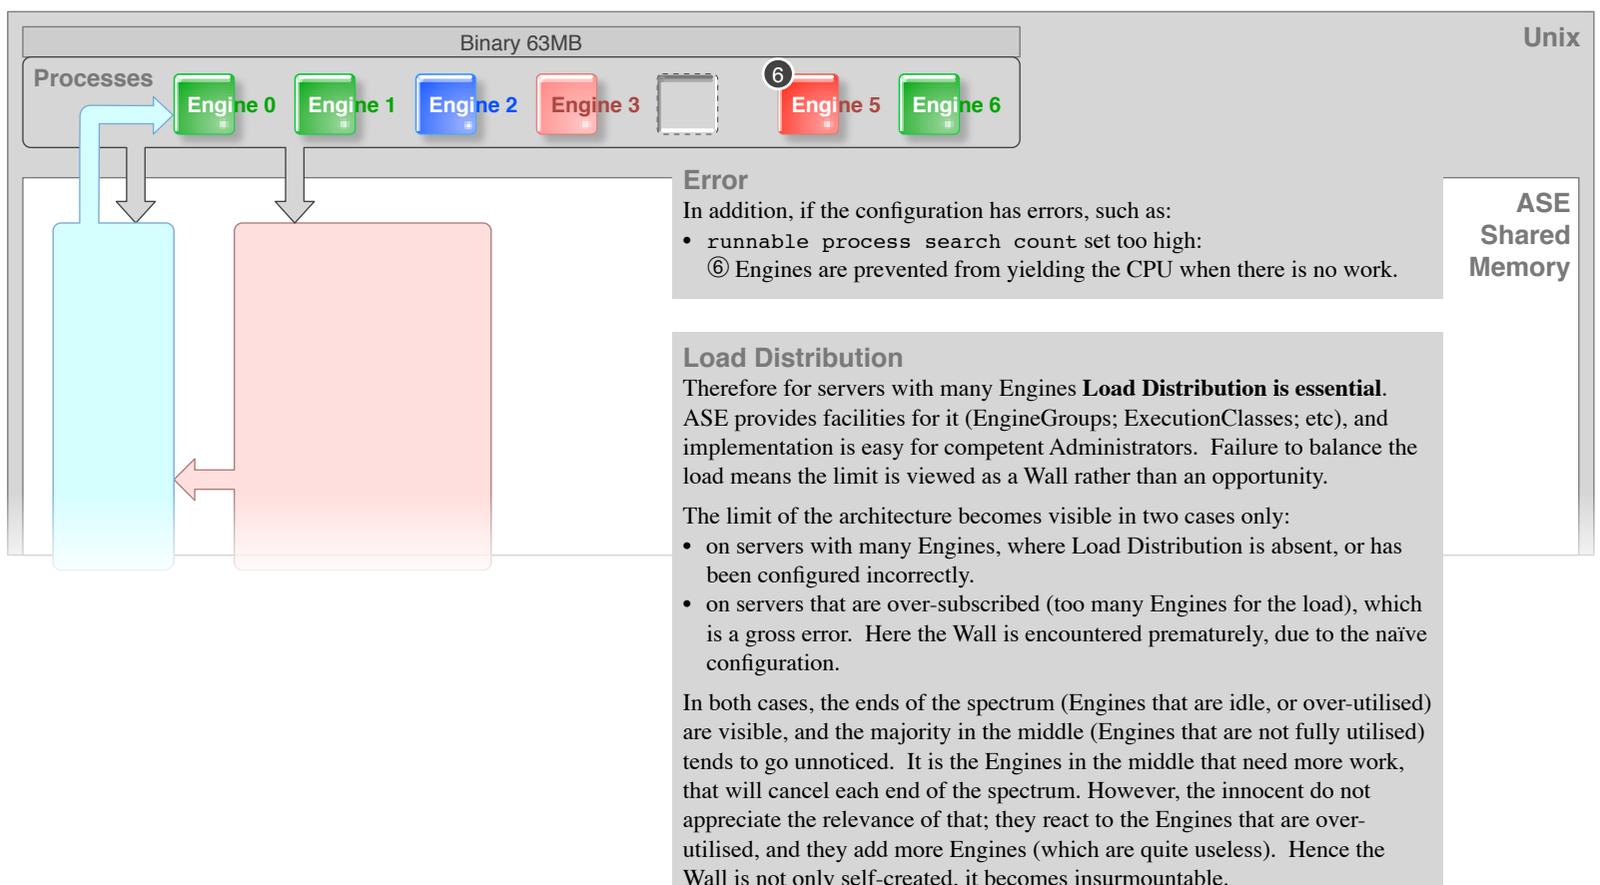
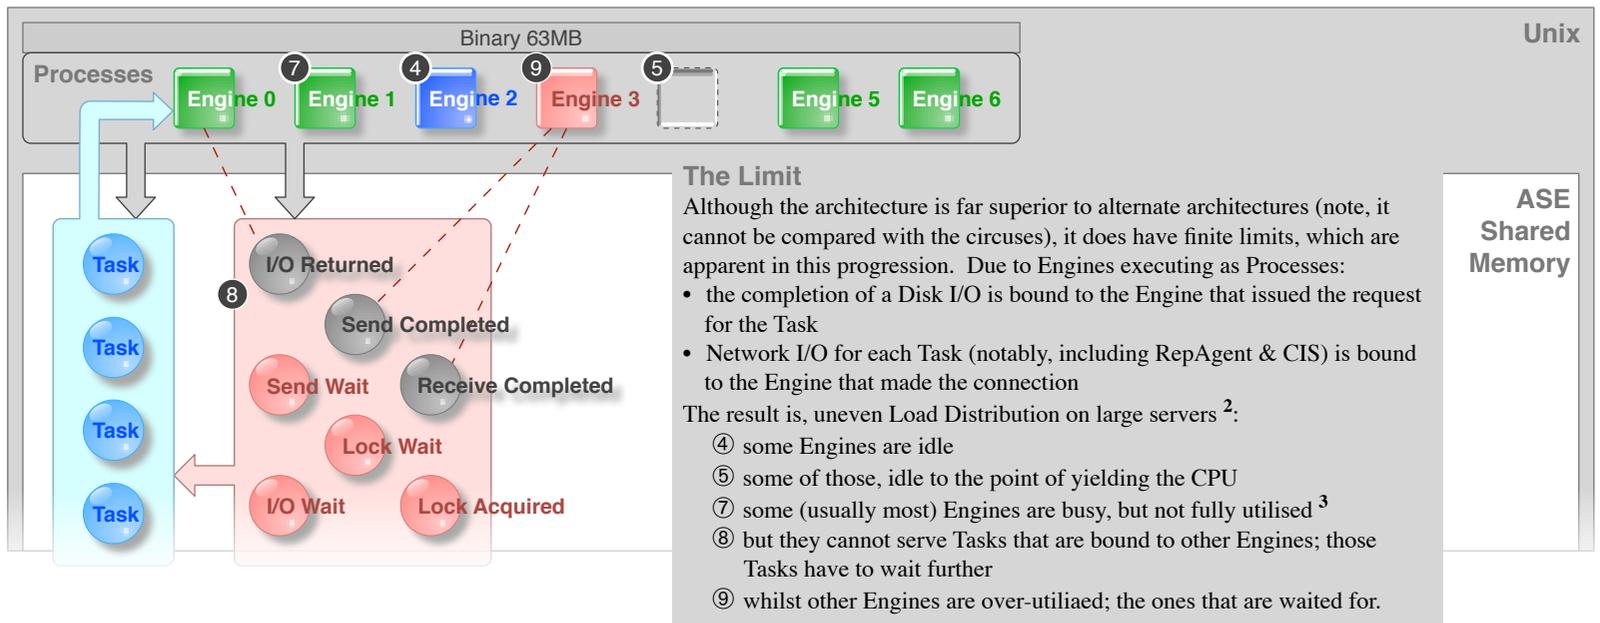
The architecture is brilliant, it allows tuning of the constituent issues discussed, and it operates uneventfully under the watch of competent Administrators.

- For the machine, configure the fewest Threads per Core; the fewest Cores per CPU, for the load.
- In ASE, configure the fewest Engines for the load.
- Aim for CPU Utilisation to be in the 80-95% range; it is an unyielding Database Server.
- Configure a small `runnable process search count` that is appropriate for your specific machine, o/s and load. Start at 5 or 10, and work upwards.

- ④ Idle; searching for runnable Task, with limit
- ⑤ Idle; yielded
- ⑥ Idle; searching for runnable Task, without limit

This section explains the limits of the Process Kernel architecture, that motivated the architecture in the next release.

- This should not be confused with the results of configuration errors, described in the previous section.



- ④ Idle; searching for runnable Task, with limit
- ⑤ Idle; yielded
- ⑥ Idle; searching for runnable Task, without limit
- ⑦ Busy; moderate CPU Usage
- ⑨ Busy; high CPU Usage

² The architectural limits are experienced only at the high end of throughput, ie. on servers that have a large number of Engines. While smaller systems can operate quite satisfactorily without implementing Load Distribution, such work is demanded for larger systems (anything over 8 Engines). Stated another way, the limit is only a limit due to absence of such work, it is easily overcome..

³ If most Engines are not utilising the CPU to 80 to 95%, the server is over-subscribed, and this will lead to an array of problems (refer to the previous section).

Sysmon Metric

In order to understand **Sybase ASE** and its components, the best avenue by far, is the examination and comprehension of the Metrics reported in **sysmon**. Note that there is, of course, a deep and meaningful, performance-related reason why each Metric is captured, and reported.

Although invaluable, **sysmon** poses problems for some people. A program that processes the reports, such as our **Sysmon Processor** overcomes them. Some obstacles posed in **sysmon**, and their manner of address are as follows:

- *Metric names are not consistent across the board, the meaning is sometimes obscured*
- *The organisation of some Metrics in poor. Together, the interpretation of Metrics is hindered*

Resolution: Metrics names have been completely Normalised; they are grouped logically and in the relevant hierarchy (notice the indentation), such that it parallels the structure of the server.

- *The reports are difficult to navigate*
- *The Metrics across many reports, which are 40 to 70 pages each, especially when 24 or 48 are being collected per day, are difficult to correlate*

Resolution: The **Sysmon Processor** produces all reports for the day in a grid, and allows various groupings, such as by period, etc. [Example Processed Report](#)

The structure of the server, and therefore the structure of the processed reports, fall into three categories:

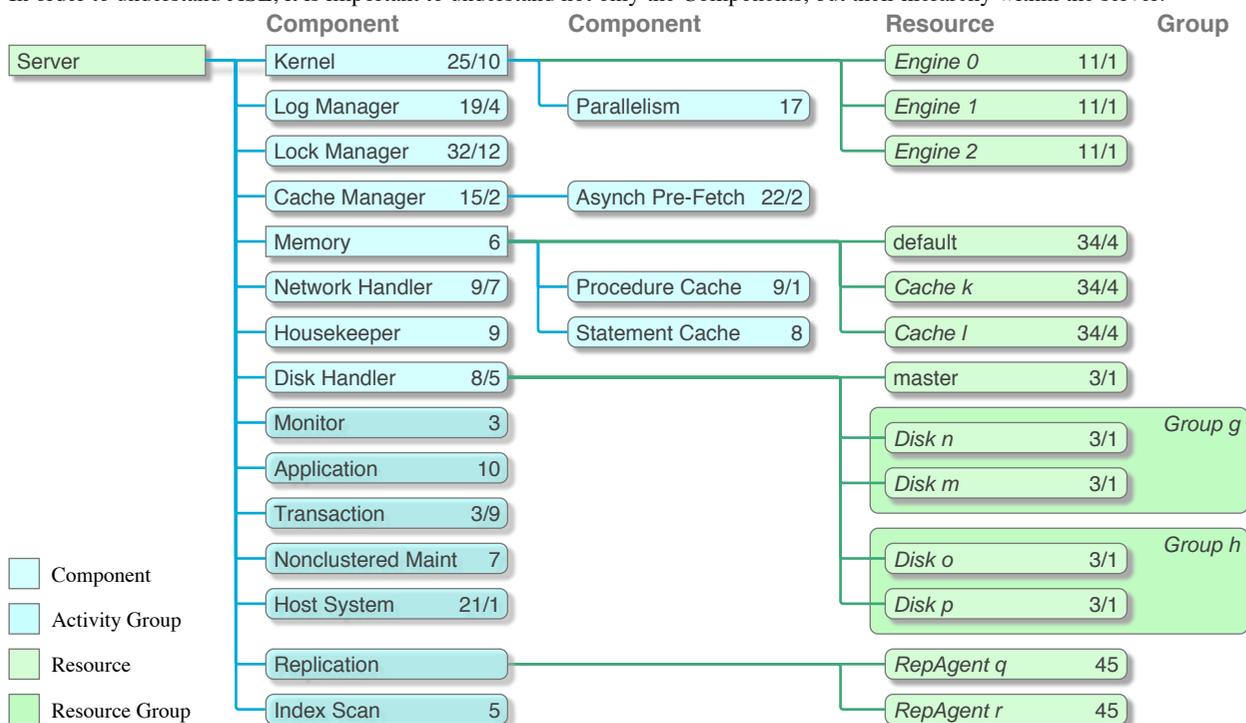
- **Component** Actual components of ASE: these exist in every server, and are available as soon as it is installed. A set of Metrics is collected for each.
- **ResourceType** There are five ResourceTypes: Disk; Cache; Engine, and optionally, ReplicationAgent and Application.
 - **Resource** Of course, everything in the server is a resource, in the normal English sense. Named Resources are specifically those resources that are added by the administrator, after installation; they are specific to each server (the first of each set is added during the installation). A set of Metrics is collected for each Resource. Since there are multiple Resources within each ResourceType, this forms a repeating group of Metrics.
 - **ResourceGroup** Allows Resources to be grouped by usage, type, etc. Essential for large numbers of Resources; and for Load Distribution.
- **Activity** Several Metrics are collected, which are neither Components nor Resources; they are grouped logically, and presented the same as Components.

Component	ComponentMetric
18 Server Components	289 Server level Metrics, grouped by Component
Asynch Pre-Fetch	
Application	Disk Check
Cache Manager	Disk Check Returned IO
Disk Handler	Disk IO Outstanding
Host System	Disk IO Request
Housekeeper	Disk IO Completed
Index Scan	...
Kernel	Log ULC Record
Log Manager	Log ULC Grant
Lock Manager	Log ULC Wait
Memory	Log Log IO Structure Grant
Monitor	Log Log IO Structure Wait
Network Handler	Log Log Allocation
Nonclustered Maint	...
Procedure Cache	Lock Request Table
Parallelism	Lock Request Table-Excl
Statement Cache	Lock Request Table-Excl-Int
Transaction	

ResourceType	Resource	ResourceMetric
Application 18 Metrics		Appl CPU Busy Appl IO Busy Appl Idle ...
Cache 52 Metrics	Server-specific, as configured	Cache Wait Cache Search Cache Hit ...
Disk 4 Metrics		Disk Read APF Disk Read Disk Write ...
Engine 12 Metrics		Engine Busy Engine CPU Busy Engine IO Busy ...
RepAgent 45 Metrics		RepAgent Log Scan RepAgent Record Scanned RepAgent Log Truncation Wait ...

Structure of Sybase ASE

In order to understand ASE, it is important to understand not only the Components, but their hierarchy within the server:



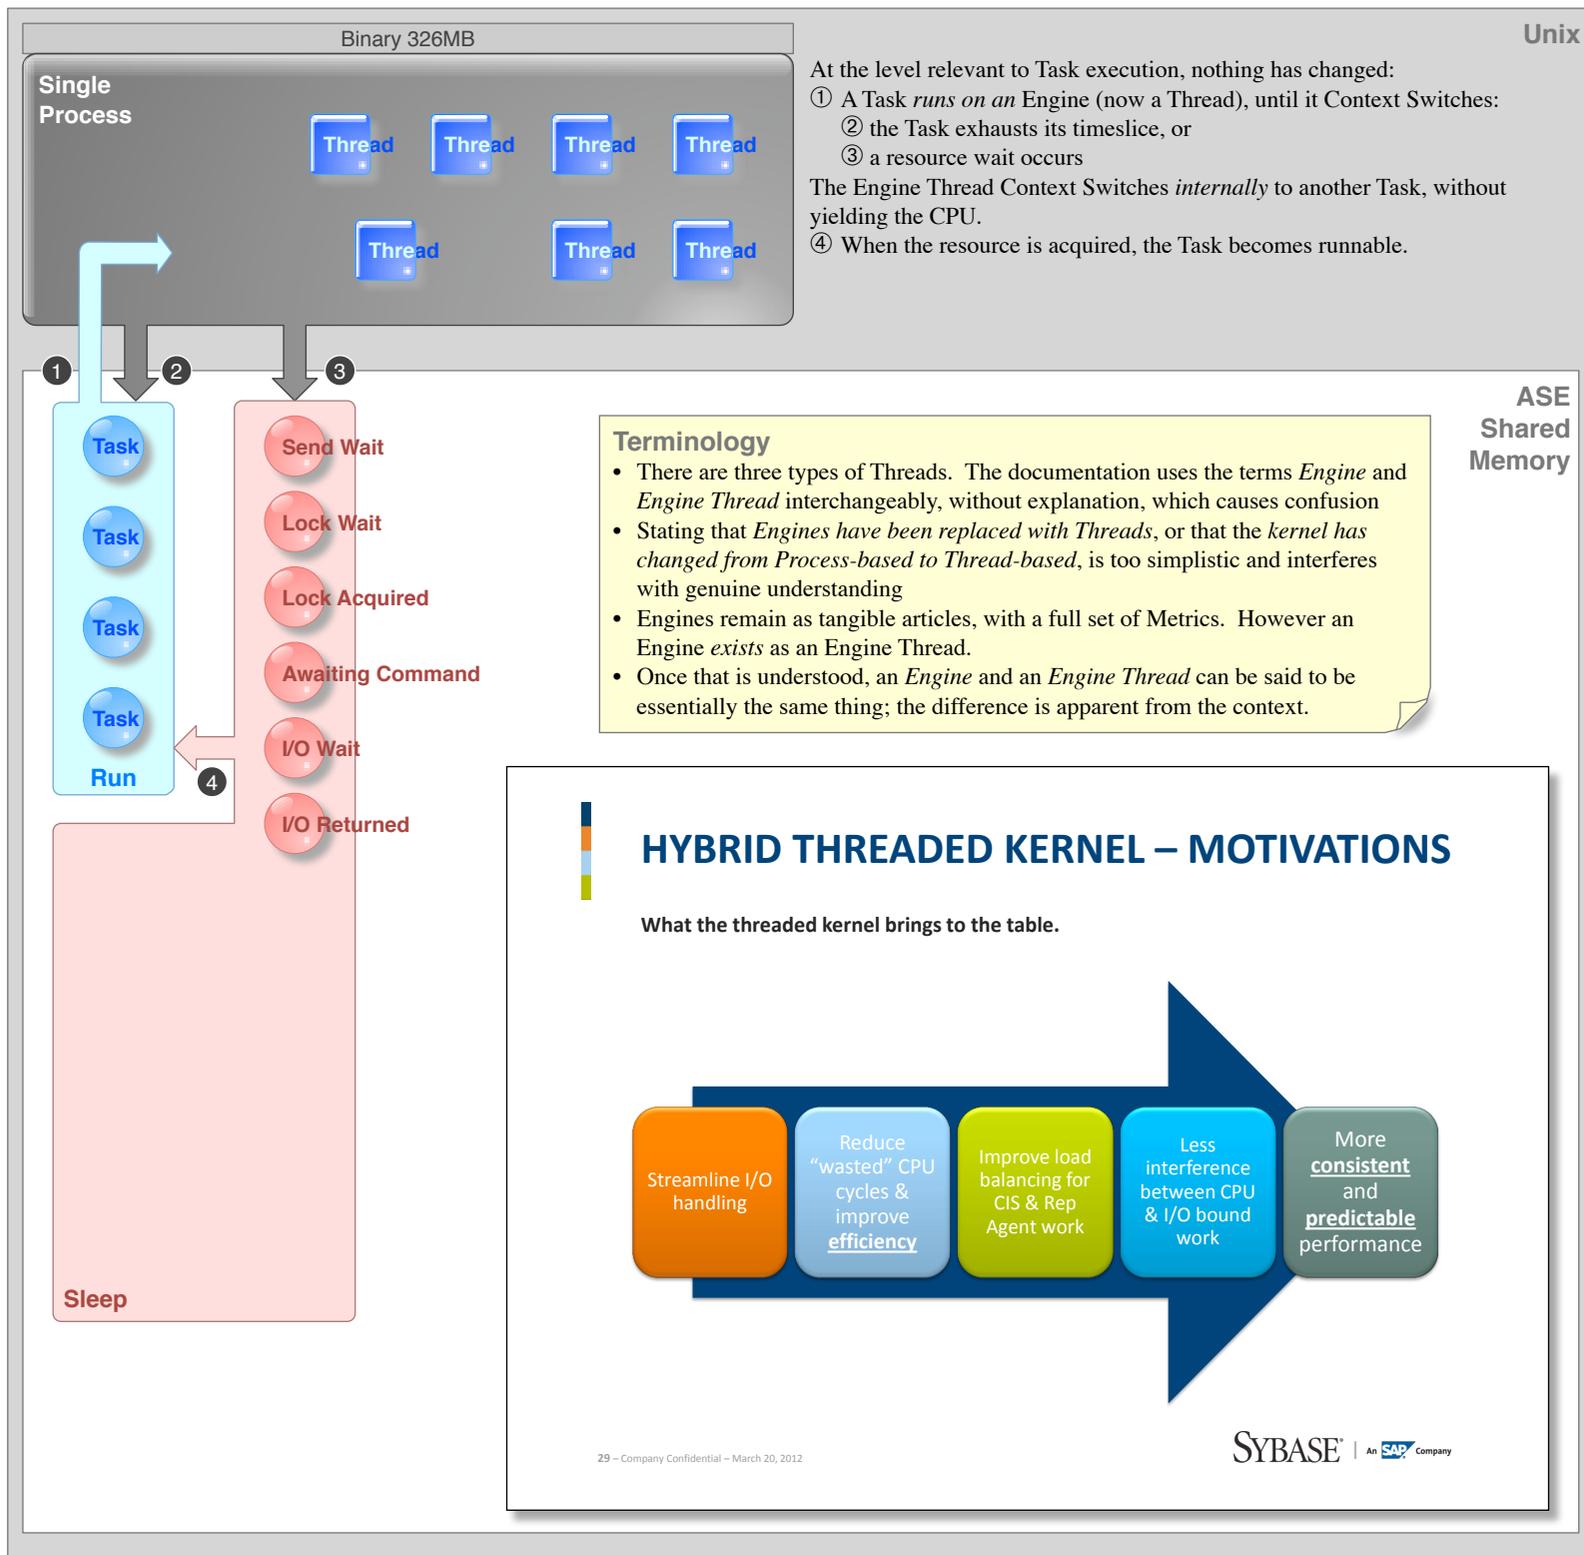
- *The Host System is of course outside the Sybase server. In order to allow Host System Metrics (vmstat, iostat) and ASE Metrics to be examined together, and to be charted or graphed together, it is treated as a Component.*
- *If it is implemented, Application level metrics are captured, (each Application is treated as a Resource). Not shown.*
- *While this diagram serves to identify the Structure of ASE, to some degree, and hopefully increases your ability to monitor it and improve its performance, it does not constitute an Architecture or Componentry diagram.*

The numbers in the cells identify the number of raw and Computed Metrics captured for the Component or Resource in the current version of our **Sysmon Processor**. Additional Metrics are computed at execution time:

- Utilisation, which is provided for all Resource Metrics
- Rate Per Sec for selected Metrics
- Schedule Utilisation

The K21 Threaded Kernel is the latest progression, in a venerable series of progressions, in both Symmetric Multi-Processing and Chip Multi-Threading:

- Modern operating systems are moving away from multi-process parallelism to multi-threaded processes, fully utilising hardware Threads
- ASE is progressing in the same manner
- It executes as a single Unix Process, using genuine o/s Threads
- The Kernel and its Threads are (no surprise) highly configurable
- Again this should be the *fewest possible* for the load.



1 EngineThread Available

2 Exhausted

3 Context Switch

- | | |
|--------------------|---------------------|
| Cache Miss | Log IO Structure |
| Disk Device | Log Last Page Write |
| Disk System Write | Network Receive |
| Exceed BatchSize | Network Send |
| Lock Normal | Network Sleep |
| Lock APL Index | Other |
| Lock Latch | User Log Cache |
| Lock LW Protection | Voluntary Yield |
| Log Group Commit | |

4 Scheduled

After resource wait is complete

Unix

Process Thread

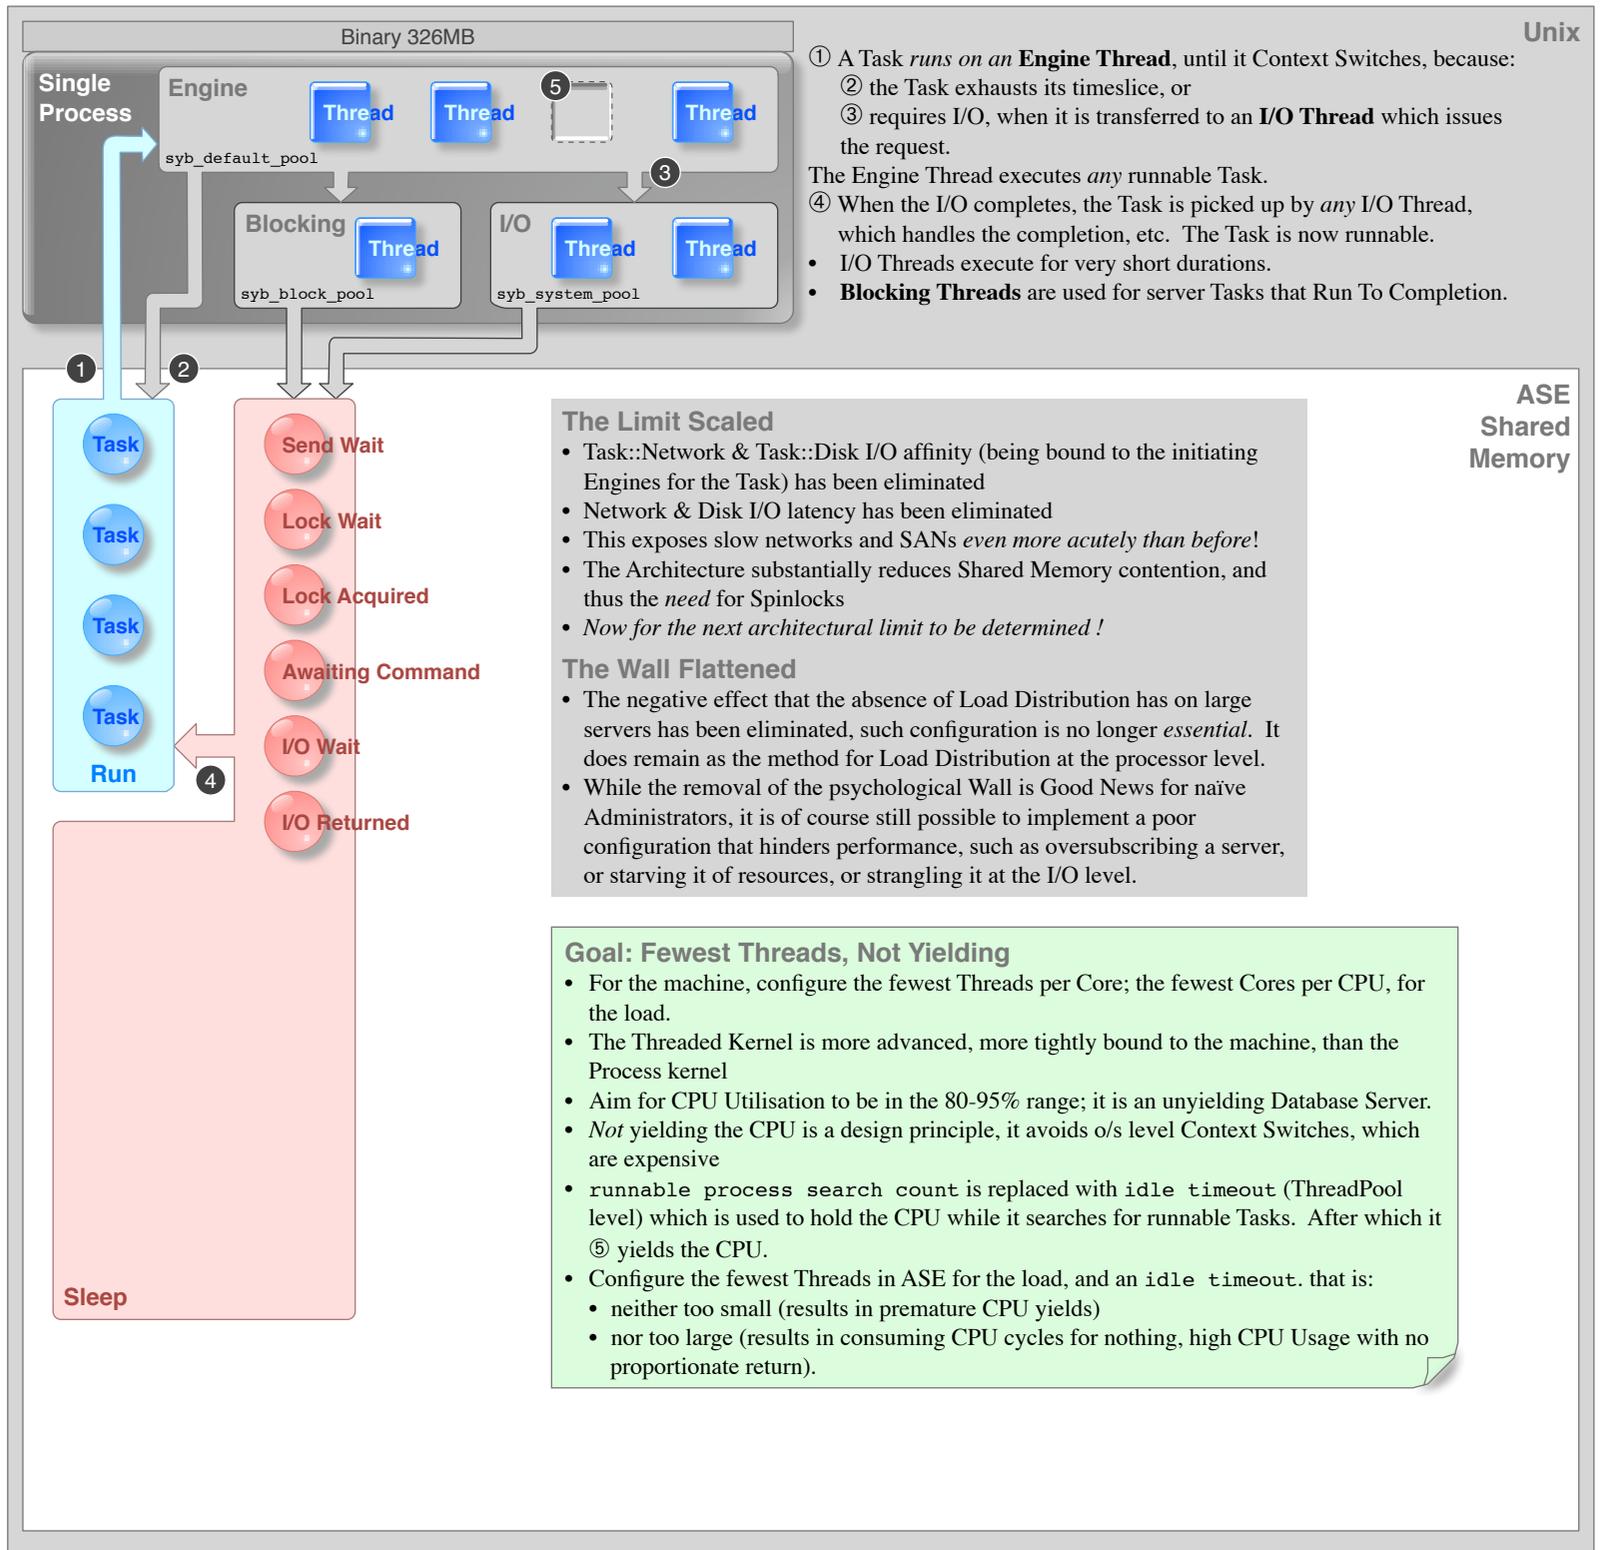
ASE Internal Task

User Running User Sleeping

All Tasks are internal, visible
Smaller size indicates Server Task

ASE V15.7 & subsequent

- Three types of Threads, organised into pools
- Additional Threads can be configured in Engine & I/O Pools
- Additional Engine ThreadPools can be configured
- EngineGroups are replaced by Engine ThreadPools



① EngineThread Available

② Exhausted

④ Scheduled

After resource wait is complete

③ Context Switch

Cache Miss	Log IO Structure
Disk Device	Log Last Page Write
Disk System Write	Network Receive
Exceed BatchSize	Network Send
Lock Normal	Network Sleep
Lock APL Index	Other
Lock Latch	User Log Cache
Lock LW Protection	Voluntary Yield
Log Group Commit	

Unix

Process Thread

ASE Internal Task

User Running User Sleeping

All Tasks are internal, visible
Smaller size indicates Server Task

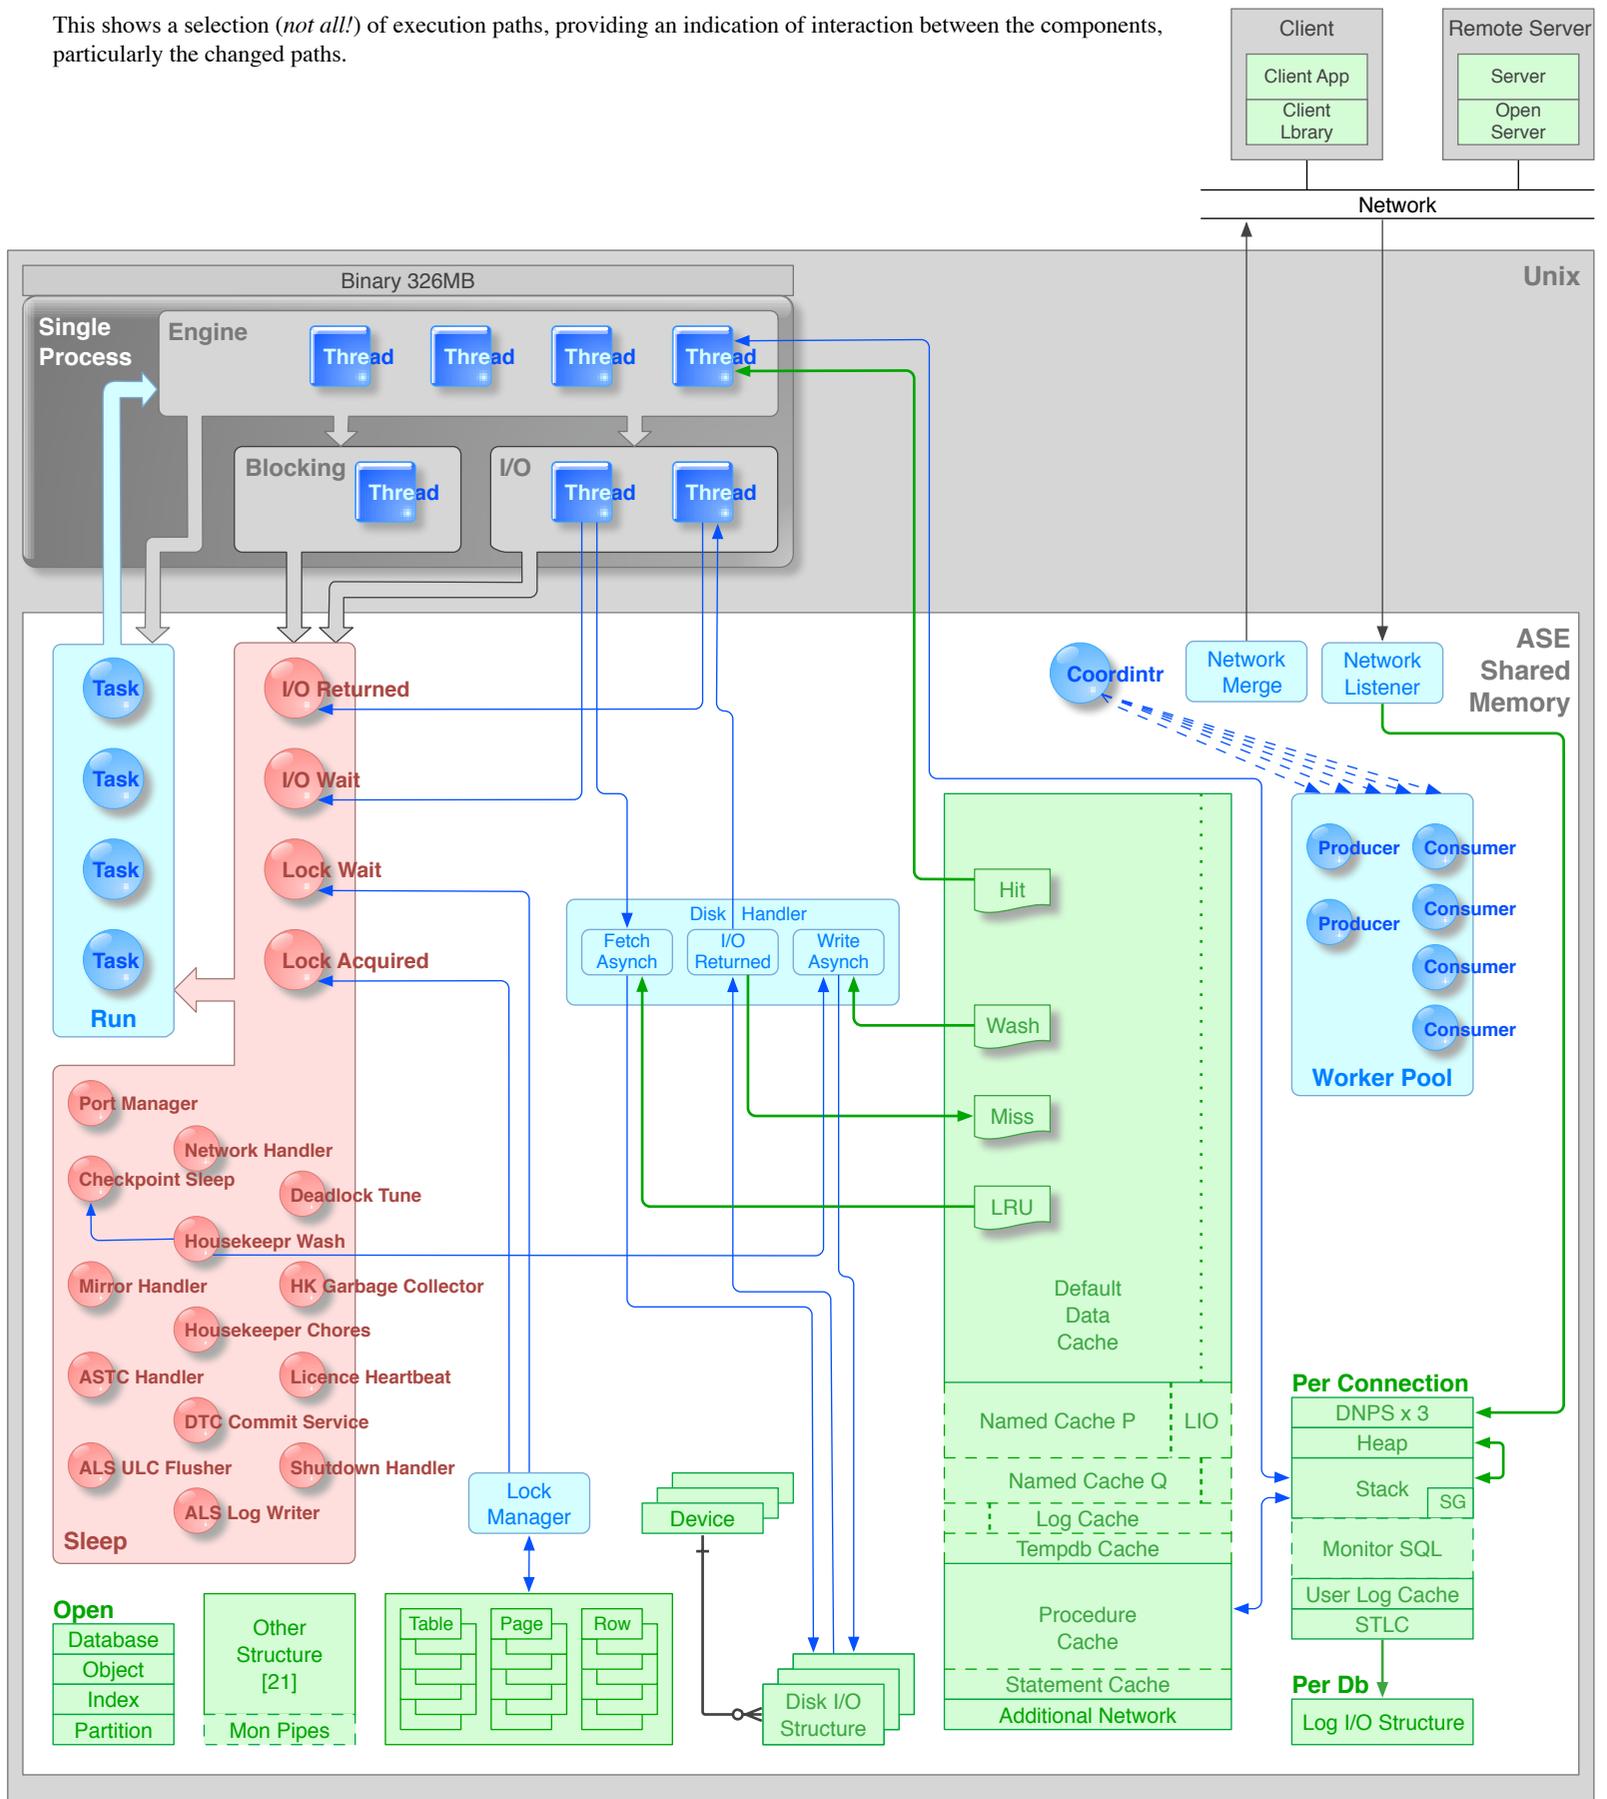
1 This is possible in the Threaded Kernel, because all threads belong to a single Process.

Sybase ASE Architecture

Threaded • Execution



This shows a selection (*not all!*) of execution paths, providing an indication of interaction between the components, particularly the changed paths.



Warning

- If your 15.5 configuration was poor, do not migrate the problem to the Threaded Kernel, **the errors will be magnified**. Fix that first.
- If you do not understand the 15.5 Kernel, and cannot configure it properly for your machine and your load, you will not be able to configure the Threaded Kernel. *If you can't ride a highly-strung horse without drama by yourself, it is not reasonable to attempt show-jumping.*
- The consequence of this Kernel being so advanced and powerful, and ASE being so configurable, is that gross configuration errors will cause it to (a) run into race conditions, such as high CPU Usage with little work being completed, or (b) run slower, and with reduced throughput.
- **Sybase ASE** is not a circus with hundreds or thousands of clowns, it is a performance by a few star performers, and highly configurable. You must decide what kind of race you are running, and configure the server accordingly. *The configuration required for a sprint vs a marathon vs a 1,500m race, are quite different.*

Path
 — Execution —>
 — Data —>

Unix
 [Grey Box] Process [Blue Box] Thread

ASE Internal Task
 [Blue Circle] User Running [Red Circle] User Sleeping
*All Tasks are internal, visible
 Smaller size indicates Server Task*

DNPS Default Network Packet Size
STLC Session Tempdb Log Cache
SG Stack Guard

ASE V15.7 & subsequent