
Open Stack on DCS

Greg Althaus and Brent Douglas
DCS Solution Engineer



Problem Definition

- What hardware would you use for a first cut recommendation for Swift and Nova?
- Couple of discussions with Rack Space about what they do today
- Can we do better?
- CAVAETS (as always)
 - No testing has been done (only experience with platforms)
 - Limited initial choice set to meet deadlines 😊



Chasing Bottlenecks

- Storage
 - Spindles are physical constraint
 - Too little network limits total I/O
 - RAID controller can be bottleneck
 - CPU & RAM are not gates but add cost

- Compute
 - RAM or Disk are the physical constraints
 - Too much CPU can overwhelm network
 - Too little Disk I/O throughput can crush CPU



Baseline

- Storage Node
 - 2 socket quad core
 - 16 GB memory
 - 48 3.5” disks maximized for storage
 - 4U

- Compute Node
 - 2 socket quad core
 - 32 GB memory
 - 8 spindles – Raid 10
 - 2U



Buidling Blocks: Hyperscale-Optimized PowerEdge C Series

1U 2S RACK
POWEREDGE C1100



2U 2S RACK
POWEREDGE C2100



2U SHARED
INFRASTRUCTURE
POWEREDGE C6100



Purpose-built for scale-out rack deployments, large homogenous cloud/cluster application environments where the software stack provides platform availability and resiliency

2U SHARED
INFRASTRUCTURE
POWEREDGE C6105



3U PCIe EXPANSION CHASSIS
POWEREDGE C410X



Hardware Choices - Nova

Component	2100	6100 ½ Pop	6100 Full	6100 + SAN
Base	2U, 12+2 drives 1 system: 12 core, 18 ram	2U, 12 drives 2 sleds: 24 core, 24 ram	2U, 12 drives 4 sleds: 48 core, 48 ram	6U, 60 drives 4 sleds: 48 core, 48 ram
Normalized 6U	2100	6100 ½ Pop	6100 Full	6100 + SAN
Cores	36	72	144	48
Spindles (3½")	36	36	36	60
Spindles/Host	12+2	6	3	15
DIMMs	54	72	144	48
RAM	432	576	1,152	384
Cost	Lowest	Moderate- Low	Moderate- High	Highest
Target	Storage Node	Mixed Use <i>sweet spot!</i>	High Compute	Enterprise Money Tree

Hardware Choices - Swift

- C1100 with direct attached drive trays
 - Two trays per
- MD1200 with 12 3.5” Drives in 2U
 - Big as you can afford – currently 2TB max
- C1100 for front-end
- Questions?
 - Is an C1100 too much in dual CPU form?



Nova POC

- Homogenous Hardware
- 6100 half-populated
 - Maximize CPU and Memory
 - Maximize Storage Capacity
- 10 Nodes for Compute
 - Shrink this as needed
- 1 Extra Node for Controllers
- Networking is flat
 - Really consider 10GB, but could start with 1GB
- Note: Each Node has 2 servers of 4 populated



Swift POC

- Homogenous Hardware
- 7 Nodes
 - 2 C1100 providing front-end services
 - 5 C1100 providing running compute
 - › 2 MD1200 per system
- ~280 TB raw storage



DCS 7100 Series (Nucleon)

Hyper-scale, High Density, Shared Infrastructure Platform

CLOUD AND CLUSTER OPTIMIZED, SHARED INFRASTRUCTURE SERVER

Great for: Hyperscale-inspired building block for high-performance cluster computing (HPCC), Web 2.0 environments and cloud builders where density/efficiency is key

Sled Features (Current Offering)

- Up to (12) Full or (24) Half-Width Sleds
- 4 x 3.5" HDD in Half-Width Sleds
- 12 x 3.5" HDD in Full-Width Sleds
- 16 x 3.5" HDD FW Storage Sled Option
- 1 x 16 Gen II PCI-E (LP)
- 1 x Mezzanine Slot for 10GbE HBA
- iKVM, DCMI, PXE, IPMI 2.0 support

Chassis Features (Current Offering)

- 12U Shared Infrastructure Module
- Compatible with Std 19" Rack
- Up to 22/24 Nodes per 12U module
- Up to 96 Nodes per rack (48U)
- 6000W 5+1 Redundant Power Supply

- N+1 Variable Fan Array

