

SFD6

NEC

*Gideon Senderov

NEC \$1.4B/yr in R & D

Over 55 years in servers and storage (1958)

SDN, Servers, Storage, Software
M-Series and HYDRAsstor

*Chauncey Schwartz

MX10-Series

New models are M110, M310, M510, M710

M710 - Max 960 Drive, 8Gb/16Gb FC, Max 48 ports, Max 192GB Cache Memory

SAS - from 6 - 12 Gbps on the back-end - all 4 models

Big perf improvements over previous models

Support for Icehouse release of OpenStack

- Cinder drivers for FC and iSCSI (supported for a fee by NEC)

- Base OS support (RHEL 7, Ubuntu 14)

Software bundles

- Base Software (migration software included)
- Advanced Perf pack
- Local replication pack
- Remote replication
- Data security
- Virtual qos
- Virtual management pack

Forklift upgrade from older model

Base Software

- Storage Manager Express
- Storage Manager
- Storage Manager Command
- ThinProvisioning
- StoragePowerConserver
- PathManager
- AccessControl
- VMware Cooperation
- PerformanceMonitor
- PerformanceNavigator
- Data Migration (New to the bundle)

- DynamicDataReplication Express

Advanced Performance

- PerforOptimizer
- PerforCache

Local Replication

- DynamicDataReplication / SnapVolume
- ReplicationControl SQL
- ReplicationControl File System

Remote Replication

- RemoteDataReplication
- DisasterRecovery

Data Security Pack

- SecureEraser
- VolumeProtect

Virtual QoS

- IO Load Manager
- Virtual Cache Partitioning

Virtual Machine Management

- Analyzer for VMware vCenter Operations

PerforOptimizer (Auto Data Tiering) - Move data to the most effective storage automatically made on access characteristics

PerforCache (W/R SSD Cache) - Use SSD as secondary cache for performance improvement

256MB chunk size for tiering

ThinProvisioning, VAAI Support and VMware Operability

NEC Storage Analyzer for single pane of glass performance view

I/O flow control - controlling I/O flow on a logical disk basis to stabilise performance - based on IOPS only, not throughput

QoS kicks in immediately, graph displays every 10 seconds

I/O limiting is strictly a CLI feature at the moment

No plan for RESTful API or PowerShell at the moment

*Gideon - HYDRAsstor

Looking to address

- inadequate scalability of capacity and perf

- — cannot scale perf to keep up with growth
- — multi products with diff architectures
- — more siloed capacity to manage
- limited deduplication scope

Local vs Global data deduplication

dedupe is done in-line

Resiliency - traditional RAID is not sufficient for dedupe data

Scalable Grid Storage Architecture

community of smart nodes

nodes

- industry std servers
- multiple types allowed
- heterogeneous and open

software

- fully distributed
- self-aware and self-organisation
- data management services
- virtualises hardware

On-line upgrade / expansion with multi-generation nodes

- in-place tech refresh with no migration
- greener, faster storage
- continuous data availability
- reduce capex and opex
- non-disruptive scaling from zero provisioning storage

No virtual edition, want to control the whole thing

Hands-Free Management

- Simple, fast deployment
- self-discovering capacity
- self-tuning and resource management
- self-healing
- web-browser gui

Biggest system in the US is 50 nodes

Accelerator nodes (no user data on them) or Storage Nodes now combined as hybrid nodes, no expansion nodes

Advanced Erasure-Coded Data Resiliency - “Distributed Resilient Data (TM)”

- “User dialable” disk / node protection - default protection against 3 concurrent failures
- greater protection with lower overhead - default 25% capacity overhead
- faster self-healing with less performance degradation - only data is reconstructed,

not entire drives

Right now, supports NFS v2 and SMB v1

Scalable Inline Global Data Deduplication

Distributed 2-Tier Architecture

- Independent linear scalability of perf and cap

Global Deduplication

- data deduplication across all data in all nodes

Distributed hash table

- data routed to responsible storage node
- deduplication and hash table processing scales linearly with storage nodes

Uses gigabit on the back-end (4 x 1 gigabit connections between nodes)

Uses SHA-1 (160 bits) for hashing and weekly background CRC scrub

Deduped then compressed (using LZ0)

HYDRAsstor does not use self-encrypting drives, but there is a data at rest encryption option (occurs after dedupe and compression)

Per controller - 40TB per hour inline backup speed (Using HS8-4000 and OST), PB per hour when using 165 nodes (sounds theoretical)

WAN-optimized replication (per files system)

- asynchronous grid-to-grid WAN-optimized replication for DR
- Dedupe across all replicated HYDRAsstor grids
 - minimises network bandwidth requirements
 - minimises DR site capacity requirements
- Policy-based data selection - file system granularity
- in-flight encryption

OST Suite - components separately or packaged at a lower price

- Deduped Transfer
- Express I/O
- Dynamic I/O (no charge)
- Optimized Synthetics
- Optimized Copy
- A.I.R.

Hybrid Node and Storage Node

Also a mini-HYDRA (HS3-410) - single controller, 8 - 24TB RAW, 104 - 312TB Effective

No direct path to tape or NDMP, media server would dump to tape

Working on an object storage interface - S3 and SWIFT

Common code base

Software features supported for entire product line

New HS6-4000A - Scale-out archive

- inherited scale out grid architecture from HS8
- focus on large archive data vs backup data

*Demo of HYDRAsstor and OST by Nobu