# Scalable Storage with Ethernet Disk Drives – Kinetic on the Move

Erik Riedel
EMC
erik.riedel@emc.com


from flickr/Blude, floppy disks for breakfast

from flickr/erinhillaw, Math/Physics Bike Rack

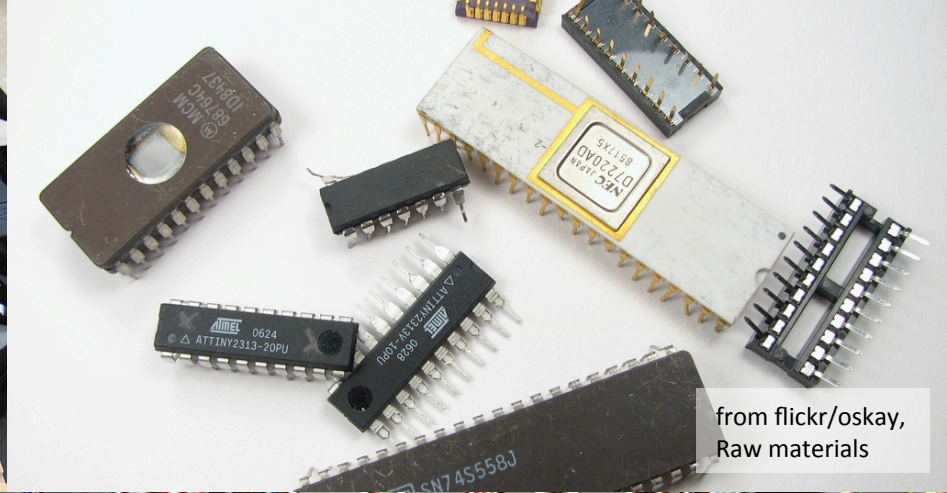


from flickr/oskay, Raw materials

**We Build Big Clouds From Little Pieces**

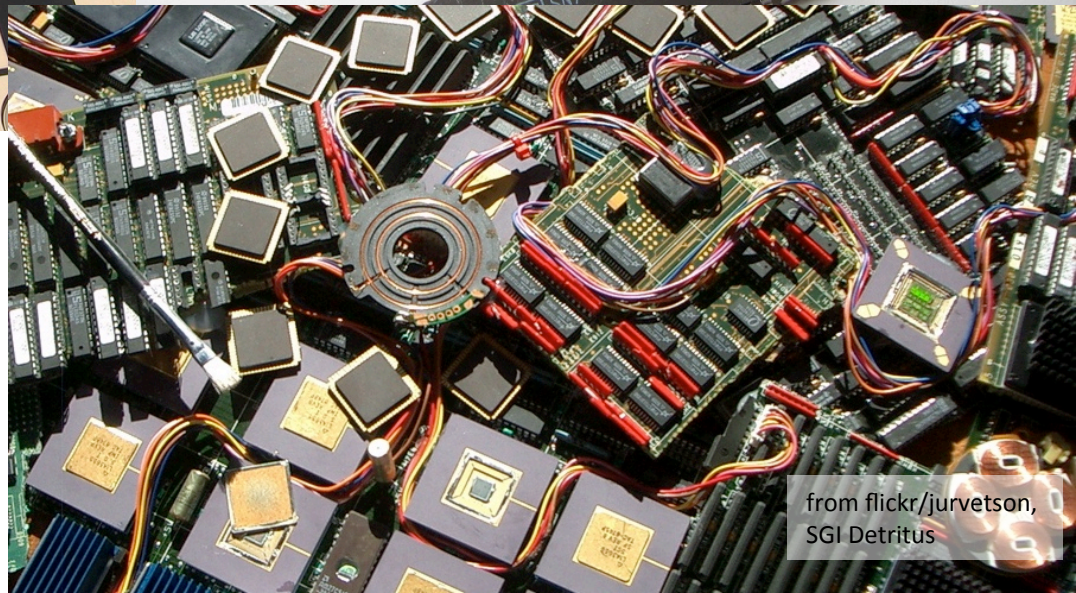

from flickr/jurvetson, SGI Detritus

# Background

- EMC Object Storage
  - Centera – since April 2002
  - Atmos – since November 2008
  - ECS – Elastic Cloud Storage – since June 2014
- over $3b object storage sold across three products
- over 3 EB exabytes of deployed capacity
- all three are Linux-based ethernet object stores


EMC Elastic Cloud Storage (ECS)
OVERVIEW
SOLUTIONS
SINGLE. SHARED. GLOBAL.
The third generation object platform from EMC, Elastic Cloud Storage (ECS) is designed for next-gen applications and traditional workloads with unmatched storage efficiency, resiliency and simplicity. It can be deployed as a turnkey storage appliance or software-only solution designed to run on industry-standard hardware.
EMC²
ECS

since 2002


ATMOS
BIG. SMART. ELASTIC.
WATCH VIDEO
EMC Strengthens Object Strategy with new Atmos Compliance Security and
OVERVIEW
DETAILS
EMC Atmos: Object-base
Atmos provides the esse
and public cloud storage.
EMC²
ATMOS

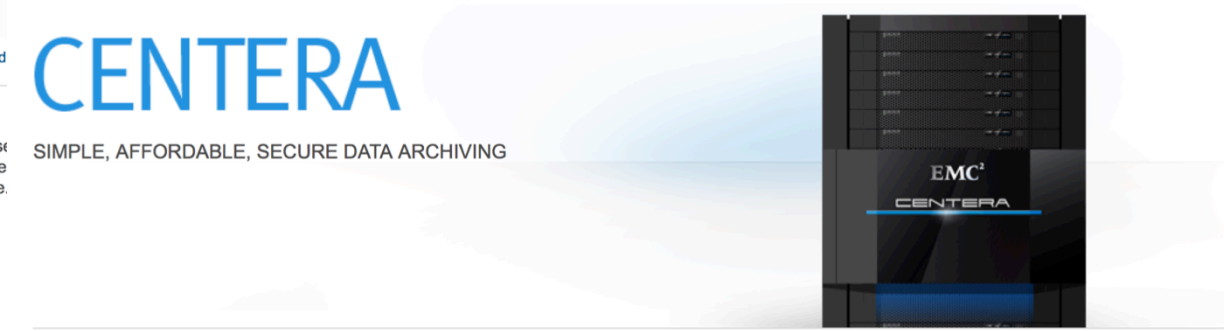
CENTERA
SIMPLE, AFFORDABLE, SECURE DATA ARCHIVING
EMC²
CENTERA
OVERVIEW
OPTIONS
EMC Centera: Content-addressable storage (CAS) platform for data archiving. Centera provides content authenticity, governance and compliance, long-term retention, and high availability with maximum efficiency and low total cost of ownership (TCO).

# SCALE

| | srvs | encs | disks | TB drv | TB rack |
|---|---|---|---|---|---|
| Atmos gen1 | 16 | 16 | 240 | 1TB | 240TB |
| Atmos gen2 | 6 | 24 | 360 | 2TB | 720TB |
| Atmos gen3 | 8 | 8 | 480 | 4TB | 1.9PB |
| Atmos gen4 | 8 | 8 | 480 | 6TB | 2.9PB |
| ECS gen5 | 8 | 8 | 480 | 8TB | 3.8PB |
| ECS gen5+ | 8 | 8 | 784 | 8, 10TB | 6.3PB |
| ECS gen6 | 4 | 7 | 588 | 8, 12TB | 7.0PB |

Gen1 (2008)
1TB

Gen2 (2010)
2,3TB

Gen3 (2012)
3,4,6TB

Gen4 (2014)
6TB

Gen5 (2015)
8TB

Gen5+ (2016)
8TB

Gen6 (2017)
12TB

# Mechanicals

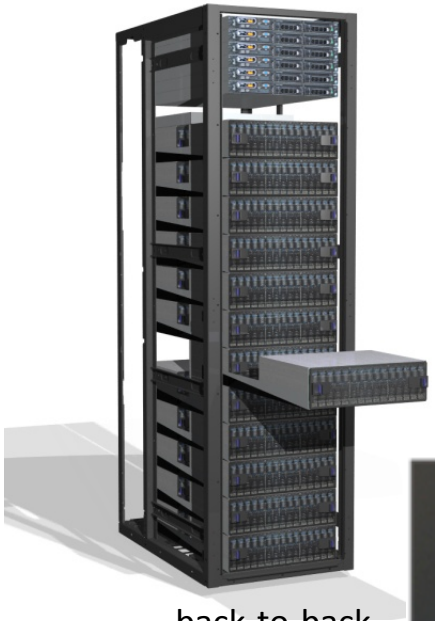

back-to-back

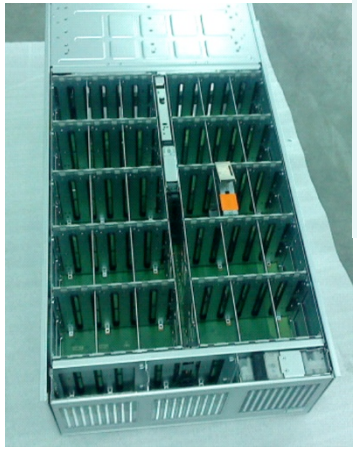

drawers

drawers

sleds

modular

trays

modular

# Density

Updated from “Long-Term Storage”, presented at Library of Congress Workshop in September 2012

| 2012 | Disks (raw) @ 3TB | Disks (protected) | Racks @ 480 disks |
|---|---|---|---|
| 5 PB | 1,700 disks | 2,700 disks | 6 racks |
| 20 PB | 6,700 disks | 11,000 disks | 23 racks |
| 50 PB | 17,000 disks | 27,000 disks | 56 racks |

| 2014 | Disks (raw) @ 6TB | Disks (protected) | Racks @ 480 disks |
|---|---|---|---|
| 5 PB | 830 disks | 1,300 disks | 3 racks |
| 20 PB | 3,300 disks | 5,300 disks | 12 racks |
| 50 PB | 8,300 disks | 13,000 disks | 28 racks |

| 2016 | Disks (raw) @ 10TB | Disks (protected) | Racks @ 780 disks |
|---|---|---|---|
| 5 PB | 500 disks | 700 disks | 1 rack |
| 20 PB | 2,000 disks | 2,800 disks | 4 racks |
| 50 PB | 5,000 disks | 7,000 disks | 9 racks |

# Scale Out

| RU | NILE DENSE |
|---|---|
| 40 | GbE |
| 39 | 10 GbE |
| 38 | 10 GbE |
| 37 | Empty |
| 36 | |
| 35 | Rinjin 4 Blade |
| 34 | |
| 33 | Blank |
| 32 | |
| 31 | |
| 30 | |
| 29 | Blank |
| 28 | |
| 27 | |
| 26 | |
| 25 | Blank |
| 24 | |
| 23 | |
| 22 | |
| 21 | Blank |
| 20 | |
| 19 | |
| 18 | |
| 17 | Voyager 10 Disk |
| 16 | |
| 15 | |
| 14 | |
| 13 | Voyager 10 Disk |
| 12 | |
| 11 | |
| 10 | |
| 9 | Voyager 10 Disk |
| 8 | |
| 7 | |
| 6 | |
| 5 | Voyager 10 Disk |
| 4 | |
| 3 | |
| 2 | |
| 1 | Not Used |

480 TB/4n60d
U400

| RU | NILE DENSE |
|---|---|
| 40 | GbE |
| 39 | 10 GbE |
| 38 | 10 GbE |
| 37 | Rinjin 4 Blade |
| 36 | |
| 35 | Rinjin 4 Blade |
| 34 | |
| 33 | Voyager 10 Disk |
| 32 | |
| 31 | |
| 30 | |
| 29 | Voyager 10 Disk |
| 28 | |
| 27 | |
| 26 | |
| 25 | Voyager 10 Disk |
| 24 | |
| 23 | |
| 22 | |
| 21 | Voyager 10 Disk |
| 20 | |
| 19 | |
| 18 | |
| 17 | Voyager 10 Disk |
| 16 | |
| 15 | |
| 14 | |
| 13 | Voyager 10 Disk |
| 12 | |
| 11 | |
| 10 | |
| 9 | Voyager 10 Disk |
| 8 | |
| 7 | |
| 6 | |
| 5 | Voyager 10 Disk |
| 4 | |
| 3 | |
| 2 | |
| 1 | Not Used |

960 TB/4n120d
U900

| RU | NILE DENSE |
|---|---|
| 40 | GbE |
| 39 | 10 GbE |
| 38 | 10 GbE |
| 37 | Rinjin 4 Blade |
| 36 | |
| 35 | Rinjin 4 Blade |
| 34 | |
| 33 | Voyager 30 Disk |
| 32 | |
| 31 | |
| 30 | |
| 29 | Voyager 30 Disk |
| 28 | |
| 27 | |
| 26 | |
| 25 | Voyager 30 Disk |
| 24 | |
| 23 | |
| 22 | |
| 21 | Voyager 30 Disk |
| 20 | |
| 19 | |
| 18 | |
| 17 | Voyager 30 Disk |
| 16 | |
| 15 | |
| 14 | |
| 13 | Voyager 30 Disk |
| 12 | |
| 11 | |
| 10 | |
| 9 | Voyager 30 Disk |
| 8 | |
| 7 | |
| 6 | |
| 5 | Voyager 30 Disk |
| 4 | |
| 3 | |
| 2 | |
| 1 | Not Used |

1.9 PB/4n240d
U1900

| RU | NILE DENSE |
|---|---|
| 40 | GbE |
| 39 | 10 GbE |
| 38 | 10 GbE |
| 37 | Rinjin 4 Blade |
| 36 | |
| 35 | Rinjin 4 Blade |
| 34 | |
| 33 | Voyager 60 Disk |
| 32 | |
| 31 | |
| 30 | |
| 29 | Voyager 60 Disk |
| 28 | |
| 27 | |
| 26 | |
| 25 | Voyager 60 Disk |
| 24 | |
| 23 | |
| 22 | |
| 21 | Voyager 60 Disk |
| 20 | |
| 19 | |
| 18 | |
| 17 | Voyager 60 Disk |
| 16 | |
| 15 | |
| 14 | |
| 13 | Voyager 60 Disk |
| 12 | |
| 11 | |
| 10 | |
| 9 | Voyager 60 Disk |
| 8 | |
| 7 | |
| 6 | |
| 5 | Voyager 60 Disk |
| 4 | |
| 3 | |
| 2 | |
| 1 | Not Used |

3.8PB/8n480d
U4000

17 PB petabytes
48 nodes
3,840 disks

# Scale Out Further

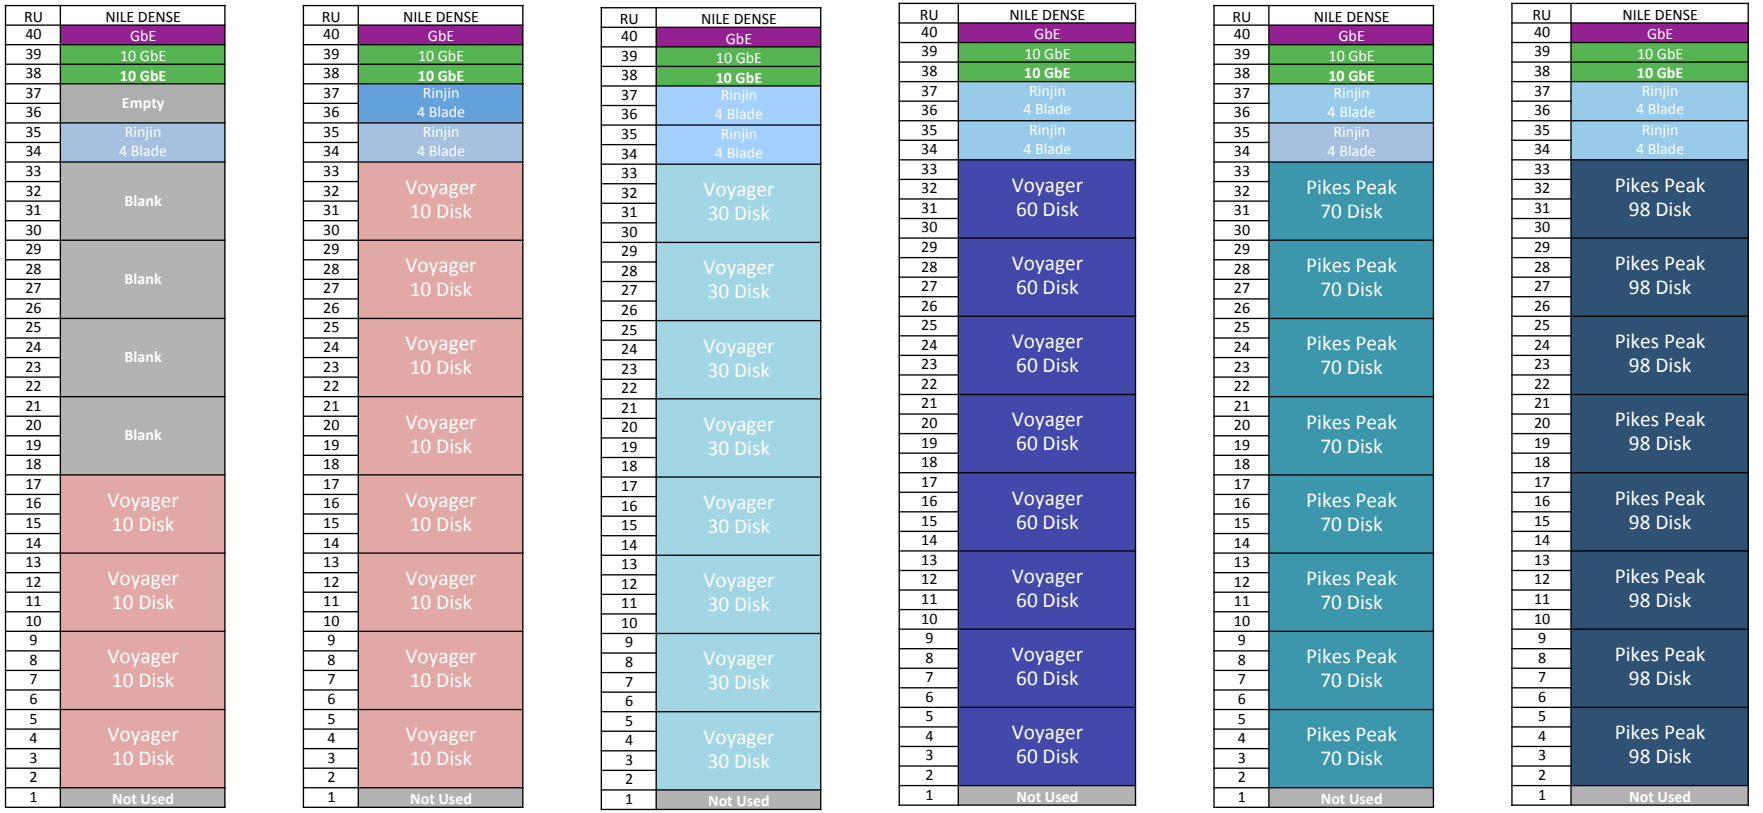


320 TB
U400

640 TB
U400T

1.9 PB
U2000

3.8PB
U4000

4.5PB
D4500

6.2PB
D6200

# FLEXIBILITY

# Experiments in Flexibility


320 TB
U400




400 TB
U400E




480 TB
U480E




640 TB
U400T




1.9 PB
U2000




3.8PB
U4000


2.0 PB/24n
F2000-24

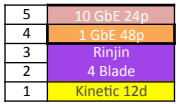


96 TB/4n
F100

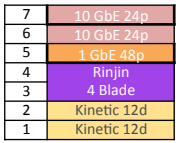


192 TB/4n
F200

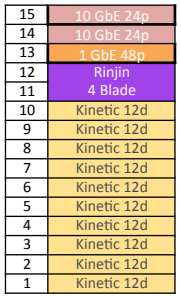


960 TB/4n
F1000

| | | |
|---|---|---|
| 40 | turtle | |
| 39 | rabbit | |
| 38 | left | right |
| 37 | lehi | murray |
| 36 | layton | logan |
| 35 | orem | ogden |
| 34 | provo | sandy |
| 33 | eight | |
| 31 | | |
| 29 | seven | |
| 27 | | |
| 25 | six | |
| 23 | | |
| 21 | five | |
| 19 | | |
| 17 | four | |
| 15 | | |
| 13 | three | |
| 11 | | |
| 09 | two | |
| 07 | | |
| 05 | one | |
| 03 | | |
| 01 | EMPTY | |

# Experiment – SAS Switching

| name | hardware | #drives | drive type | port A | port B |
|---|---|---|---|---|---|
| one | Voyager DAE | 60 | Muskie SATA 2TB | left 9 | right 9 |
| two | Voyager DAE | 60 | Muskie SATA 2TB | left 10 | right 10 |
| three | Voyager DAE | 60 | Muskie SATA 2TB | left 11 | right 11 |
| four | Voyager DAE | 60 | Muskie SATA 2TB | left 12 | right 12 |
| five | Voyager DAE | 60 | Muskie SATA 2TB | left 13 | right 13 |
| six | Voyager DAE | 60 | Muskie SATA 2TB | left 14 | right 14 |
| seven | Voyager DAE | 60 | Muskie SATA 2TB | left 15 | right 15 |
| eight | Voyager DAE | 60 | Muskie SATA 2TB | left 16 | right 16 |


left


right

2012

| | |
|---|---|
| 40 | |
| 39 | 10 GbE |
| 38 | C6100-10GbE |
| 37 | 2/4 Node |
| 36 | C6100-4 Node |
| 35 | 10 GbE |
| 34 | SAS 32 MAX |
| 33 | Voyager |
| 32 | |
| 31 | |
| 30 | |
| 29 | Voyager |
| 28 | |
| 27 | |
| 26 | |
| 25 | Voyager |
| 24 | |
| 23 | |
| 22 | |
| 21 | Voyager |
| 20 | |
| 19 | |
| 18 | |
| 17 | Voyager |
| 16 | |
| 15 | |
| 14 | |
| 13 | Voyager |
| 12 | |
| 11 | |
| 10 | |
| 9 | Voyager |
| 8 | |
| 7 | |
| 6 | |
| 5 | Voyager |
| 4 | |
| 3 | |
| 2 | |
| 1 | GbE |
| | |
| | 480 |

# Experiment – SAS Switching

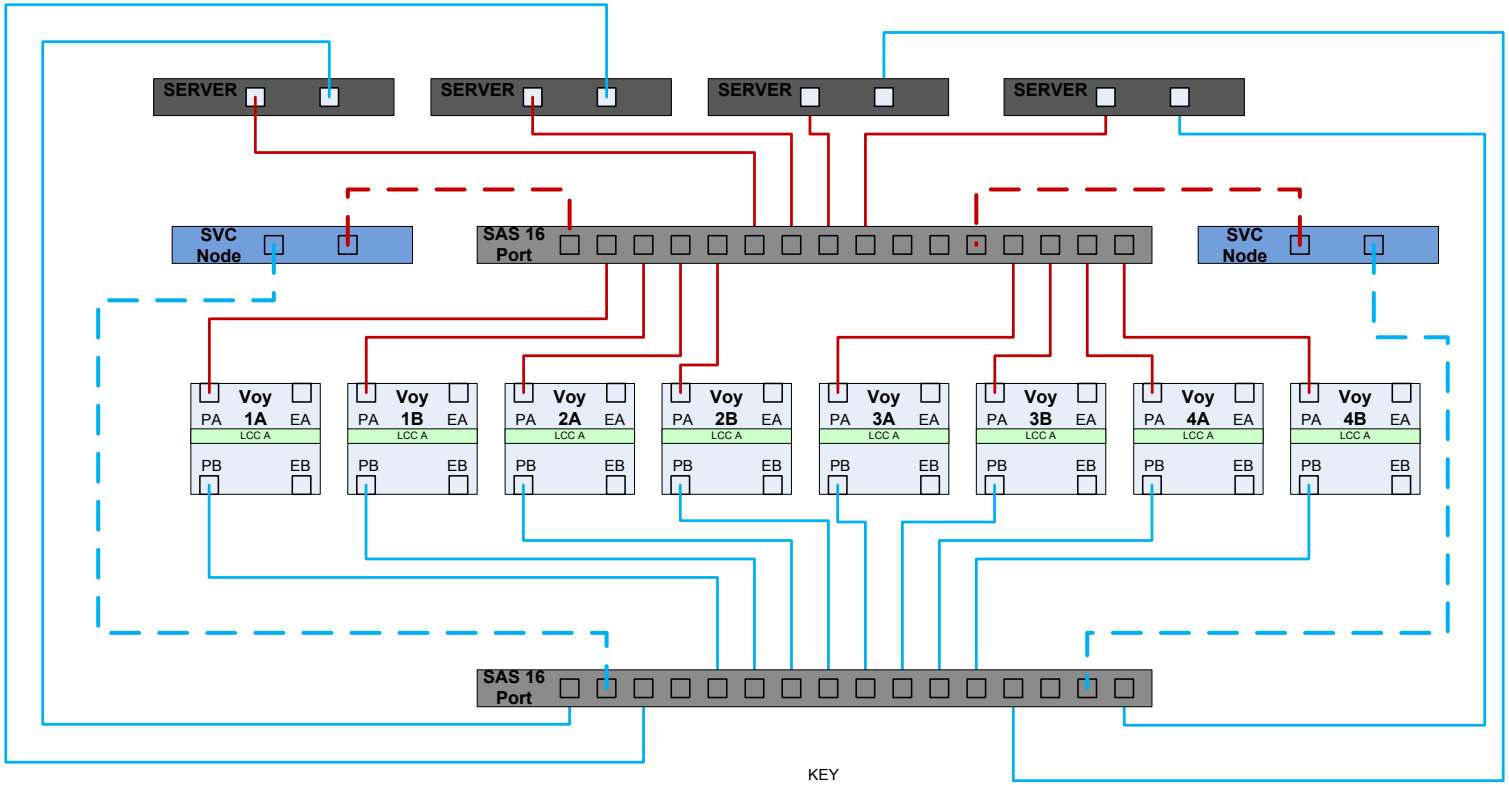


## Rack Config & SAS Wiring Diagram

(A3-480/LightsOut Config)

2012

# Experiment – SAS Zoning

Entry Point = 4 x 15 disk nodes = 60 disks

2013

2 x 15 disk nodes illustrated below

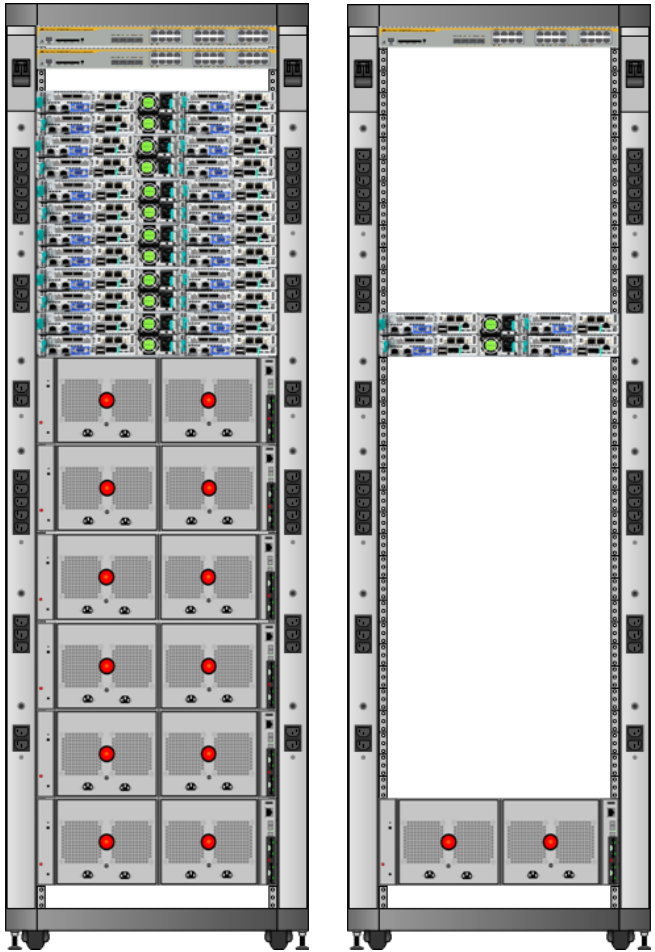

Max = 24 x 15 disk nodes = 360 disks

4 Server “PHOENIX” Chassis

60 Disk JBOD, “VOYAGER”

# Experiment – Kinetic

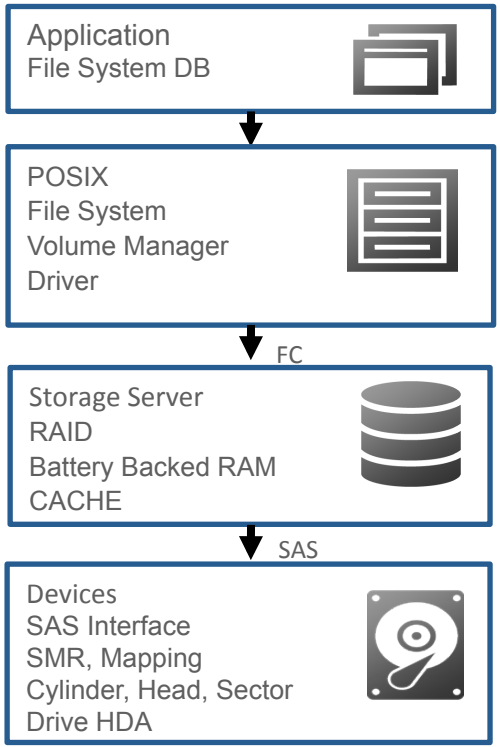


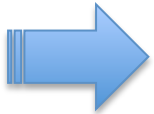

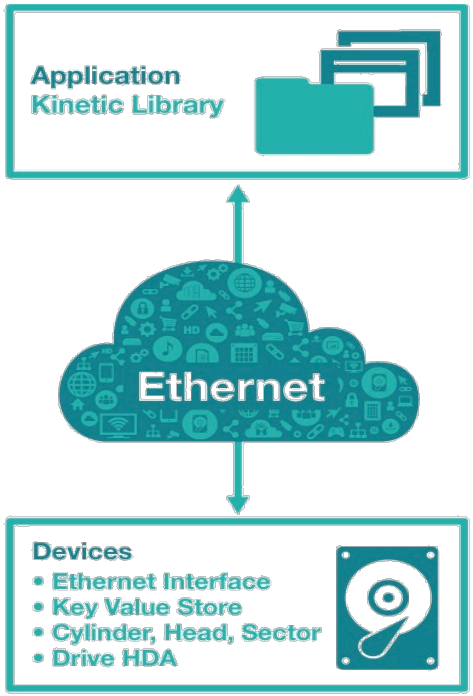


From Seagate material “Kinetic Open Storage - Enabling Break-through Economics in Scale-out Object Storage”

2014

# Experiment – Kinetic


- Newisys EDA-4605 Enclosure
  - 60-disk
  - dual 10 GbE controllers
  - 4x 10 GbE uplinks

NEWISYS

**Newisys® EDA-4605**
4U-60 3.5" Ultra Dense
Ethernet Drive Storage Array


Ultra-Efficient Ethernet-Direct HDD Storage System delivering Object storage to Cloud and Big Data Deployments

The Newisys EDA-4605 paves the way for the latest developments in Storage technology targeting Ethernet-direct Hard Disk Drive technology. Optimized for Object storage, the EDA-4605 implements an ultra-efficient storage platform for Cloud and Big Data deployments. It is ideal for scale-out and distributed storage solutions.

Unlike traditional storage boxes, the EDA-4605 provides redundant Ethernet fabrics that connect directly to the disk drives, eliminating many layers of overhead and enabling new levels of storage scalability. Whether deployed in Cloud installations, for Big Data or the traditional data center, the EDA-4605 delivers object storage at unprecedented efficiency.

The first disk product to leverage the EDA-4605 is the new Seagate Kinetic Open Storage drive.

With up to 60 x 3.5" Seagate Kinetic Open Storage drives per 4U enclosure, the industry-unique Newisys EDA-4605 is an ultra-dense, space and power saving, complexity-reducing storage solution. The Newisys EDA-4605 fits nicely into a standard 19" wide, 1m deep, rack that easily creates a 15 drives/U object storage building block. With 4TB drives, this can deliver 2.4 Petabyte per standard 42U rack, and can easily scale out beyond that.

**Product Highlights**

- Optimized for Ethernet-Direct HDDs, such as the novel Seagate Kinetic Open Storage Drives
- Ideal building block for Object Storage deployments
- Reduces complexity and improves efficiency for Cloud and Big Data Object Storage installations
- Full-featured, highly available, high performance Object Storage

**Product Features**

- Four 10GbE connections to the datacenter
- Redundant, 1GbE connections to each of the 60 HDDs
- Dual, redundant, hot-pluggable Ethernet Switch and Management (ESM) modules
- Dual, redundant, hot-pluggable, high efficiency power supply and fan units (PSU)
- Redundant hot-pluggable, system blowers implemented in the PSUs
- Modular design increases product configuration flexibility
- Standard chassis customization and branding available

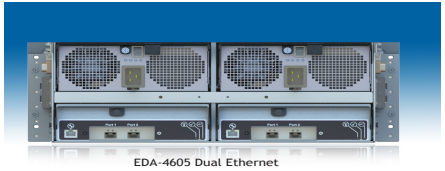

EDA-4605 Dual Ethernet Switch and Management

- Seagate Kinetic Ethernet drive
  - 4TB in October 2014 (2x 1 GbE net)
  - 8TB in September 2015 (2x 2.5 GbE net)

2014

# PRESENT DAY

# Storage w/ Kinetic - Prototype

2016

12n504d
240 cores
4,032 TB raw
160 Gbps

**4.0PB/12n**
**F4000-12**

| RU | ECS FLEX |
|---|---|
| 40 | 10 GbE |
| 39 | 10 GbE |
| 38 | GbE |
| 37 | Rinjin 4 Blade |
| 36 | |
| 35 | Rinjin 4 Blade |
| 34 | |
| 33 | Rinjin 4 Blade |
| 32 | |
| 31 | |
| 30 | Titan/Kinetic 84 Disk |
| 29 | |
| 28 | |
| 27 | |
| 26 | |
| 25 | Titan/Kinetic 84 Disk |
| 24 | |
| 23 | |
| 22 | |
| 21 | |
| 20 | Titan/Kinetic 84 Disk |
| 19 | |
| 18 | |
| 17 | |
| 16 | |
| 15 | Titan/Kinetic 84 Disk |
| 14 | |
| 13 | |
| 12 | |
| 11 | |
| 10 | Titan/Kinetic 84 Disk |
| 9 | |
| 8 | |
| 7 | |
| 6 | |
| 5 | Titan/Kinetic 84 Disk |
| 4 | |
| 3 | |
| 2 | |
| 1 | |

| RU | ECS FLEX |
|---|---|
| 40 | 10 GbE |
| 39 | 10 GbE |
| 38 | GbE |
| 37 | Rinjin 4 Blade |
| 36 | |
| 35 | Rinjin 4 Blade |
| 34 | |
| 33 | Rinjin 4 Blade |
| 32 | |
| 31 | |
| 30 | Titan/Kinetic 14 Disk |
| 29 | |
| 28 | |
| 27 | |
| 26 | |
| 25 | Titan/Kinetic 14 Disk |
| 24 | |
| 23 | |
| 22 | |
| 21 | |
| 20 | Titan/Kinetic 14 Disk |
| 19 | |
| 18 | |
| 17 | |
| 16 | |
| 15 | Titan/Kinetic 14 Disk |
| 14 | |
| 13 | |
| 12 | |
| 11 | |
| 10 | Titan/Kinetic 14 Disk |
| 9 | |
| 8 | |
| 7 | |
| 6 | |
| 5 | Titan/Kinetic 14 Disk |
| 4 | |
| 3 | |
| 2 | |
| 1 | |

| RU | ECS FLEX |
|---|---|
| 40 | 10 GbE |
| 39 | 10 GbE |
| 38 | GbE |
| 37 | Rinjin 4 Blade |
| 36 | |
| 35 | Rinjin 4 Blade |
| 34 | |
| 33 | Rinjin 4 Blade |
| 32 | |
| 31 | |
| 30 | Titan/Kinetic 14 Disk |
| 29 | |
| 28 | |
| 27 | |
| 26 | |
| 25 | Titan/Kinetic 14 Disk |
| 24 | |
| 23 | |
| 22 | |
| 21 | |
| 20 | Titan/Kinetic 14 Disk |
| 19 | |
| 18 | |
| 17 | |
| 16 | |
| 15 | Titan/Kinetic 14 Disk |
| 14 | |
| 13 | |
| 12 | |
| 11 | |
| 10 | Titan/Kinetic 14 Disk |
| 9 | |
| 8 | |
| 7 | |
| 6 | |
| 5 | Titan/Kinetic 14 Disk |
| 4 | |
| 3 | |
| 2 | |
| 1 | |

## Parts List

- 9x Rinjin servers (36 nodes)
- 2x 10 GbE per node SFP+
- 6x 10 GbE data switches (Arista 64-port SFP+)
- 3x 1 GbE mgmt switches (Arista 48-port Cat6)
- 18x Titan enclosures (dual controller, 4x 10 GbE uplinks)
- 3x loose Titan enclosures
- 84 * 6 + 14 * 15 = 714
- 714x Kinetic/8TB drives
- SFP+ twinax cables (data)
- Cat6 cables (mgmt)

## Adding

- 6x 40 GbE data switches (Arista 32-port QSFP+)
- updating nodes to 2x 40GbE

# PERFORMANCE

# Enclosure + Server Bandwidth

| | enclosure | server | #drives | drives total bandwidth | enclosure physical | #nodes | enclosure bandwidth |
|---|---|---|---|---|---|---|---|
| gen4 | voyager | phoenix | 60 | 8,400 MB/s | 4x 3G SATA | 1 | 980 MB/s |
| gen5 | voyager | rinjin | 60 | 8,400 MB/s | 8x 6G SAS | 1 | 4,100 MB/s |
| gen5+ | pikes peak | rinjin | 98 | 13,720 MB/s | 8x 12G SAS | 1 | 5,800 MB/s |
| gen6 | titan | rinjin | 84 | 10,080 MB/s | 4x 10G ethernet | 4 | 4,000 MB/s |

*all bandwidths are micro-benchmark “OS-level” measurements, end-to-end

- ethernet-based enclosure aggregate performance is comparable to SAS-based enclosure
- using **2.5 GbE** drive links and **10 GbE** enclosure uplinks

# MANAGEABILITY

from flickr/purplemattfish, Broken hard drive?

# Topology Map – getrackinfo

```
provo-vanilla:~ # getrackinfo -a
Node private      Node              Public                                RMM
Ip Address        Id       Status   Mac                 Ip Address        Mac                 Ip Address          Node Name
===============   ======   ======   =================   ================  =================   =================   =========
192.168.219.1     1        MA       00:1e:67:9f:01:96   10.249.248.111    00:1e:67:69:29:8f   10.249.248.101      provo-vanilla
192.168.219.2     2        SA       00:1e:67:9f:01:a2   10.249.248.112    00:1e:67:69:28:72   10.249.248.102      sandy-vanilla
192.168.219.3     3        SA       00:1e:67:9e:ff:9e   10.249.248.113    00:1e:67:69:29:99   10.249.248.103      orem-vanilla
192.168.219.4     4        SA       00:1e:67:9e:ff:7f   10.249.248.114    00:1e:67:69:29:5f   10.249.248.104      ogden-vanilla
192.168.219.5     N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.6     N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.7     N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.8     N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.9     9        O        N/A                 N/A               00:1e:67:69:28:4f   10.249.248.109      boston-vanilla
192.168.219.10    10       O        N/A                 N/A               00:1e:67:69:29:6d   10.249.248.110      chicago-vanilla
192.168.219.11    11       O        N/A                 N/A               00:1e:67:69:29:3f   10.249.248.111      houston-vanilla
192.168.219.12    12       O        N/A                 N/A               00:1e:67:69:27:2c   10.249.248.112      phoenix-vanilla
192.168.219.13    13       O        N/A                 N/A               00:1e:67:69:28:8c   10.249.248.113      dallas-vanilla
192.168.219.14    14       O        N/A                 N/A               00:1e:67:69:28:3f   10.249.248.114      detroit-vanilla
192.168.219.15    15       O        N/A                 N/A               00:1e:67:69:29:8a   10.249.248.115      columbus-vanilla
192.168.219.16    16       O        N/A                 N/A               00:1e:67:69:29:8c   10.249.248.116      austin-vanilla
192.168.219.17    N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.18    N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.19    N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.20    N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.21    N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.22    N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.23    N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
192.168.219.24    N/A      noLink   N/A                 N/A               N/A                 N/A                 N/A
Status:
M - Master, S - Slave
E – Epoxy
I - Initializing, U - Updating, A - Active
P - On, O - Off
```

# Topology Map – getrackinfo – details

```
provo-vanilla:~ # getrackinfo -a
Node private      Node              Public                                  RMM
Ip Address        Id       Status   Mac                Ip Address           Mac                Ip Address          Node Name
===============   ======   ======   ================   ==================   ================   =================   =========
192.168.219.1     1        MA       00:1e:67:9f:01:96  10.249.248.111       00:1e:67:69:29:8f  10.249.248.101      provo-vanilla
192.168.219.2     2        SA       00:1e:67:9f:01:a2  10.249.248.112       00:1e:67:69:28:72  10.249.248.102      sandy-vanilla
192.168.219.3     3        SA       00:1e:67:9e:ff:9e  10.249.248.113       00:1e:67:69:29:99  10.249.248.103      orem-vanilla
192.168.219.4     4        SA       00:1e:67:9e:ff:7f  10.249.248.114       00:1e:67:69:29:5f  10.249.248.104      ogden-vanilla
192.168.219.5     N/A      noLink   N/A                N/A
192.168.219.6     N/A      noLink   N/A                N/A
192.168.219.7     N/A      noLink   N/A                N/A
192.168.219.8     N/A      noLink   N/A                N/A
192.168.219.9     9        O        N/A                N/A
192.168.219.10    10       O        N/A                N/A
192.168.219.11    11       O        N/A                N/A
192.168.219.12    12       O        N/A                N/A
192.168.219.13    13       O        N/A                N/A
192.168.219.14    14       O        N/A                N/A
192.168.219.15    15       O        N/A                N/A
192.168.219.16    16       O        N/A                N/A
192.168.219.17    N/A      noLink   N/A                N/A
192.168.219.18    N/A      noLink   N/A                N/A
192.168.219.19    N/A      noLink   N/A                N/A
192.168.219.20    N/A      noLink   N/A                N/A
192.168.219.21    N/A      noLink   N/A                N/A
192.168.219.22    N/A      noLink   N/A                N/A
192.168.219.23    N/A      noLink   N/A                N/A
192.168.219.24    N/A      noLink   N/A                N/A
Status:
M - Master, S - Slave
E - Epoxy
I - Initializing, U - Updating, A - Active
P - On, O - Off
```

```
provo-vanilla:~ # getrackinfo -v
=======================================================
NodeName              : provo-vanilla
Node Id               : 1

Interfaces(MAC & IP)
--------------------
public                : 00:1e:67:9f:01:96   10.249.248.111/21
private               : 00:1e:67:69:29:8b   192.168.219.1/24
private ipmi          : 00:1e:67:69:29:8d   192.168.219.101/24
private.4(NAN)        : 00:1e:67:69:29:8b   169.254.186.1/16
remote ipmi           : 00:1e:67:69:29:8f   10.249.248.101/21

Network Services
----------------
NTP Configuration:
       server:    10.254.140.21 10.254.140.22
DNS Configuration:
       domain:
       search:     sea.lab.emc.com corp.emc.com emc.com
       server:     192.168.219.254 10.6.149.11
```

# ECS Hardware Abstraction Layer

```
= provo-silver
Disks(s):
SCSI Device Block Device Enclosure   Partition Name                          Slot Serial Number       SMART
---------- ------------ ----------   ------------------------------------- ---- ------------------- -------
/dev/sg2    /dev/sdc     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A00  NAHJ3TDY            GOOD
/dev/sg3    /dev/sdd     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A01  NAH08A3X            GOOD
/dev/sg4    /dev/sde     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A02  NAHJWW9X            GOOD
/dev/sg5    /dev/sdf     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A03  NAHKNU2Y            GOOD
/dev/sg6    /dev/sdg     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A04  NCG555KS            GOOD
/dev/sg7    /dev/sdh     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A05  NAHKNXSY            GOOD
/dev/sg8    /dev/sdi     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A06  NAHJEJUY            GOOD
/dev/sg9    /dev/sdj     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A07  NCG4VL8S            SUSPECT
/dev/sg10   /dev/sdk     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A08  NAHKDVBY            GOOD
/dev/sg11   /dev/sdl     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A09  NCG335PS            GOOD
/dev/sg12   /dev/sdm     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A10  NAHH72KX            GOOD
/dev/sg13   /dev/sdn     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     A11  NAHKP0AY            GOOD
/dev/sg14   /dev/sdo     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     B00  NCG4ZWGS            GOOD
/dev/sg15   /dev/sdp     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     B01  NAHHU5RY            GOOD
/dev/sg16   /dev/sdq     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     B02  NAHH
/dev/sg17   /dev/sdr     /dev/sg18   ECS:object:cGSk6IN2REyhmawe8r9xXA     B03  NCG4
. . .

total: 60
```

**slot numbers (physical)**

**SMART inquiries for all drives (GOOD, BAD, SUSPECT)**

**physical serial numbers (as shown on the drive label)**

**resilient to /dev/sd* and /dev/sg* changes**

# ECS Basic Health Checks

**resilient to /dev/sd* and /dev/sg* changes**

```
= lehi-silver
Disks(s):
SCSI Device Block Device Enclosure   Partition Name                          Slot Serial Number        SMART
----------- ------------ ----------- ------------------------------------- ---- -------------------- -------
/dev/sg3    /dev/sdd     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A00  NAHWHUMY             GOOD
/dev/sg4    /dev/sde     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A01  NCGBN24S             GOOD
/dev/sg5    /dev/sdf     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A02  NAHWKLSY             GOOD
/dev/sg6    /dev/sdg     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A03  NAHWHK9Y             GOOD
/dev/sg7    /dev/sdh     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A04  NAHW9ELY             GOOD
/dev/sg8    /dev/sdi     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A05  NAHWHN5Y             GOOD
/dev/sg9    /dev/sdj     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A06  NAHWGZPY             GOOD
/dev/sg10   /dev/sdk     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A07  NAHVXAMY             GOOD
/dev/sg11   /dev/sdl     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A08  NAHSURUY             GOOD
/dev/sg12   /dev/sdm     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A09  NAHWW20X             GOOD
/dev/sg13   /dev/sdn     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A10  NAHWUYYX             GOOD
/dev/sg14   /dev/sdo     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     A11  NAHW8RUY             GOOD
/dev/sg15   /dev/sdp     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     B00  NAHWHEXY             GOOD
/dev/sg16   /dev/sdq     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     B01  NAHWH54Y             SUSPECT
/dev/sg17   /dev/sdr     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     B02  NAHV9X9X             GOOD
/dev/sg18   /dev/sds     /dev/sg19   ECS:object:IrHy+LjnQ+mXJK7ekqKwPQ     B03  NAHWLL7X             GOOD
. . .

total: 60
```

**stable slot numbers (physical)**

```
emc@provo-chartreuse:~> sudo -i cs_hal list disks
Disks(s):
SCSI Device Block Device Enclosure   Partition Name                       Slot Serial Number       SMART   DiskSet
----------- ------------ ----------- ------------------------------------ ---- ------------------- ------- ------------
/dev/sg11   /dev/sdl     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   B05  2EG7604R            GOOD
/dev/sg12   /dev/sdm     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   B06  2EG74KUR            GOOD
/dev/sg13   /dev/sdn     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   B07  2EG76ZVR            GOOD
/dev/sg14   /dev/sdo     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   C02  2EG75Z7R            GOOD
/dev/sg15   /dev/sdp     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   C03  2EG74N3R            GOOD
/dev/sg16   /dev/sdq     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   C04  2EG7788R            GOOD
/dev/sg17   /dev/sdr     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   C05  2EG73YLR            GOOD
/dev/sg18   /dev/sds     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   C06  2EG74S3R            GOOD
/dev/sg19   /dev/sdt     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   C07  2EG75N2R            GOOD
/dev/sg20   /dev/sdu     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   D02  2EG6GJZR            GOOD
/dev/sg21   /dev/sdv     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   D03  2EG6M2VR            GOOD
/dev/sg22   /dev/sdw     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   D04  2EG6M65R            GOOD
/dev/sg23   /dev/sdx     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   D05  2EG6MHAR            GOOD
/dev/sg24   /dev/sdy     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   D06  2EG6GKMR            GOOD
/dev/sg25   /dev/sdz     /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   D07  2EG6M1BR            GOOD
/dev/sg26   /dev/sdaa    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   E02  2EG6YBUR            GOOD
/dev/sg27   /dev/sdab    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   E03  2EG73U4R            GOOD
/dev/sg28   /dev/sdac    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   E04  2EG7PGLR            GOOD
/dev/sg29   /dev/sdad    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   E05  2EG76VYR            GOOD
/dev/sg30   /dev/sdae    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   E06  2EG7RA1R            GOOD
/dev/sg31   /dev/sdaf    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   E07  2EG7647R            GOOD
/dev/sg32   /dev/sdag    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   F02  2EG6Z67R            GOOD
/dev/sg33   /dev/sdah    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   F03  2EG6YHGR            GOOD
/dev/sg34   /dev/sdai    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   F04  2EG6YM3R            GOOD
/dev/sg35   /dev/sdaj    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   F05  2EG6YZ2R            GOOD
/dev/sg36   /dev/sdak    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   F06  2EG6Y2VR            GOOD
/dev/sg37   /dev/sdal    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   F07  2EG712ZR            GOOD
/dev/sg38   /dev/sdam    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   G02  2EG75XNR            GOOD
/dev/sg39   /dev/sdan    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   G03  2EG73ZER            GOOD
/dev/sg40   /dev/sdao    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   G04  2EG77A1R            GOOD
/dev/sg41   /dev/sdap    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   G05  2EG7119R            GOOD
/dev/sg42   /dev/sdaq    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   G06  2EG6ZGPR            GOOD
/dev/sg43   /dev/sdar    /dev/sg44   ECS:object:qL4GBbuLTw+FkT+Sfz+FyA   G07  2EG73X2R            GOOD
```

| SCSI Device | Block Device | Enclosure | Partition Name | Slot | Serial Number | SMART | DiskSet |
|---|---|---|---|---|---|---|---|
| /dev/sg43 | /dev/sdar | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G07 | 2EG73Y2R | GOOD | |
| /dev/sg2 | /dev/sdc | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A02 | 2EG6YWZR | GOOD | |
| /dev/sg3 | /dev/sdd | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A03 | 2EG6VEPR | GOOD | |
| /dev/sg4 | /dev/sde | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A04 | 2EG6VJ0R | GOOD | |
| /dev/sg5 | /dev/sdf | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A05 | 2EG6Z50R | GOOD | |
| /dev/sg6 | /dev/sdg | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A06 | 2EG6YTWR | GOOD | |
| /dev/sg7 | /dev/sdh | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A07 | 2EG6YT2R | GOOD | |
| /dev/sg8 | /dev/sdi | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B02 | 2EG771PR | GOOD | |
| /dev/sg9 | /dev/sdj | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B03 | 2EG73NXR | GOOD | |
| /dev/sg10 | /dev/sdk | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B04 | 2EG76Z5R | GOOD | |
| /dev/sg55 | /dev/sdbc | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B12 | 2EG7R9DR | GOOD | |
| /dev/sg56 | /dev/sdbd | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B11 | 2EG7MXBR | GOOD | |
| /dev/sg57 | /dev/sdbe | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B10 | 2EG6V5VR | GOOD | |
| /dev/sg58 | /dev/sdbf | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B09 | 2EG825VR | GOOD | |
| /dev/sg59 | /dev/sdbg | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B08 | 2EG5YUBR | GOOD | |
| /dev/sg60 | /dev/sdbh | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B01 | 2EG74VWR | GOOD | |
| /dev/sg61 | /dev/sdbi | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | C14 | 2EG742ZR | GOOD | |
| /dev/sg62 | /dev/sdbj | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | C13 | 2EG73RZR | GOOD | |
| /dev/sg63 | /dev/sdbk | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | C12 | 2EG5RD8R | GOOD | |
| /dev/sg64 | /dev/sdbl | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | C11 | 2EG74S4R | GOOD | |
| /dev/sg65 | /dev/sdbm | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | C10 | 2EG75Z3R | GOOD | |
| /dev/sg66 | /dev/sdbn | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | C09 | 2EG76P7R | GOOD | |
| /dev/sg67 | /dev/sdbo | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | C08 | 2EG74TKR | GOOD | |
| /dev/sg68 | /dev/sdbp | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | C01 | 2EG73S2R | GOOD | |
| /dev/sg69 | /dev/sdbq | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | D14 | 2EG6LPTR | GOOD | |
| /dev/sg70 | /dev/sdbr | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | D13 | 2EG42JPR | GOOD | |
| /dev/sg71 | /dev/sdbs | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | D12 | 2EG6GGBR | GOOD | |
| /dev/sg72 | /dev/sdbt | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | D11 | 2EG6EN9R | GOOD | |
| /dev/sg73 | /dev/sdbu | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | D10 | 2EG6B32R | GOOD | |
| /dev/sg74 | /dev/sdbv | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | D09 | 2EG6B1WR | GOOD | |
| /dev/sg75 | /dev/sdbw | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | D08 | 2EG5KYPR | GOOD | |
| /dev/sg76 | /dev/sdbx | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | D01 | 2EG6GS5R | GOOD | |
| /dev/sg77 | /dev/sdby | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E14 | 2EG76XZR | GOOD | |
| /dev/sg78 | /dev/sdbz | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E13 | 2EG75YZR | GOOD | |
| /dev/sg79 | /dev/sdca | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E12 | 2EG710RR | GOOD | |

| SCSI Device | Block Device | Enclosure | Partition Name | Slot | Serial Number | SMART | DiskSet |
|---|---|---|---|---|---|---|---|
| /dev/sg79 | /dev/sdca | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E12 | 2EG710RR | GOOD | |
| /dev/sg80 | /dev/sdcb | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E11 | 2EG73XLR | GOOD | |
| /dev/sg81 | /dev/sdcc | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E10 | 2EG73WAR | GOOD | |
| /dev/sg82 | /dev/sdcd | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E09 | 2EG6ZD9R | GOOD | |
| /dev/sg83 | /dev/sdce | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E08 | 2EG7102R | GOOD | |
| /dev/sg84 | /dev/sdcf | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | E01 | 2EG76TXR | GOOD | |
| /dev/sg85 | /dev/sdcg | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | F14 | 2EG6YNRR | GOOD | |
| /dev/sg86 | /dev/sdch | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | F13 | 2EG6YXZR | GOOD | |
| /dev/sg87 | /dev/sdci | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | F12 | 2EG6Z77R | GOOD | |
| /dev/sg88 | /dev/sdcj | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | F11 | 2EG6YHTR | GOOD | |
| /dev/sg89 | /dev/sdck | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | F10 | 2EG6YKAR | GOOD | |
| /dev/sg90 | /dev/sdcl | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | F09 | 2EG5B24R | GOOD | |
| /dev/sg91 | /dev/sdcm | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | F08 | 2EG7103R | GOOD | |
| /dev/sg92 | /dev/sdcn | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | F01 | 2EG743HR | GOOD | |
| /dev/sg93 | /dev/sdco | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G14 | 2EG73W8R | GOOD | |
| /dev/sg94 | /dev/sdcp | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G13 | 2EG5M2WR | GOOD | |
| /dev/sg95 | /dev/sdcq | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G12 | 2EG76RER | GOOD | |
| /dev/sg96 | /dev/sdcr | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G11 | 2EG5JAUR | GOOD | |
| /dev/sg97 | /dev/sdcs | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G10 | 2EG599ER | GOOD | |
| /dev/sg98 | /dev/sdct | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G09 | 2EG70XGR | GOOD | |
| /dev/sg99 | /dev/sdcu | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G08 | 2EG75YRR | GOOD | |
| /dev/sg100 | /dev/sdcv | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | G01 | 2EG775SR | GOOD | |
| /dev/sg45 | /dev/sdas | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A14 | 2EG6GB8R | GOOD | |
| /dev/sg46 | /dev/sdat | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A13 | 2EG6BMDR | GOOD | |
| /dev/sg47 | /dev/sdau | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A12 | 2EG5P5PR | GOOD | |
| /dev/sg48 | /dev/sdav | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A11 | 2EG65VPR | GOOD | |
| /dev/sg49 | /dev/sdaw | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A10 | 2EG6ZM7R | GOOD | |
| /dev/sg50 | /dev/sdax | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A09 | 2EG7138R | GOOD | |
| /dev/sg51 | /dev/sday | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A08 | 2EG6VHBR | GOOD | |
| /dev/sg52 | /dev/sdaz | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | A01 | 2EG6GJDR | GOOD | |
| /dev/sg53 | /dev/sdba | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B14 | 2EG6Y11R | GOOD | |
| /dev/sg54 | /dev/sdbb | /dev/sg44 | ECS:object:qL4GBbuLTw+FkT+Sfz+FyA | B13 | 2EG6ZTER | GOOD | |

total: 98

```
dino-black:~ % cs_hal list fs
Volume(s):
SCSI Device Block Device FS UUID                                Type      Slot  Label SMART    Mount Point
----------- ------------ ------------------------------------- --------- ----- ----- ------   ------------
/dev/sg2    /dev/sda     0ddb9635-ff27-4cd3-8c2f-58a6f5226d30  ext3            BOOT  GOOD     /boot
/dev/sg2    /dev/sda     2192b3ef-2a44-4450-9b04-327c00215454  xfs                   GOOD     /root2
/dev/sg2    /dev/sda     ffa9607a-4b6f-4218-9266-c083fb1989a1  xfs                   GOOD     /var
/dev/sg2    /dev/sda     746b09d4-f07a-49dc-8b40-86220dfc7edc  xfs                   GOOD     /
/dev/sg2    /dev/sda     f7c37c92-4bc5-4abf-95a5-efa51c46f6bc  swap v1               GOOD     -
/dev/sg3    /dev/sdb     90a52650-e0f3-49e4-810b-a505cdcadb51  xfs       0           GOOD     /data-disks/ss-90a52650-e0f3-49e4-810b-a505cdcadb51
/dev/sg4    /dev/sdc     173aef8b-80e9-4be2-a510-3b88d3343f8a  xfs       1           GOOD     /data-disks/ss-173aef8b-80e9-4be2-a510-3b88d3343f8a
/dev/sg5    /dev/sdd     bcfb1897-152b-482b-bde6-de9665ad7c51  xfs       2           GOOD     /data-disks/ss-bcfb1897-152b-482b-bde6-de9665ad7c51
/dev/sg6    /dev/sde     bc6946ae-770f-4621-9ea5-f2d1e5ec0f28  xfs       3         SUSPECT    /data-disks/ss-bc6946ae-770f-4621-9ea5-f2d1e5ec0f28
/dev/sg7    /dev/sdf     52446742-a566-4036-8b0c-5cd7901474f0  xfs       4           GOOD     /data-disks/ss-52446742-a566-4036-8b0c-5cd7901474f0
/dev/sg8    /dev/sdg     c9ee0971-d8dc-4621-8958-d79890d0f590  xfs       5           GOOD     /data-disks/ss-c9ee0971-d8dc-4621-8958-d79890d0f590
/dev/sg9    /dev/sdh     294bcd25-ab19-40ee-8c03-cd71e94e9e06  xfs       6           GOOD     /meta/294bcd25-ab19-40ee-8c03-cd71e94e9e06
/dev/sg10   /dev/sdi     cb5cac6c-1cdf-49ec-8754-a475db3d4afd  xfs       7           GOOD     /data-disks/ss-cb5cac6c-1cdf-49ec-8754-a475db3d4afd
/dev/sg11   /dev/sdj     91739495-2a46-47d2-8676-d8b4b3f8fd76  xfs       8           GOOD     /data-disks/ss-91739495-2a46-47d2-8676-d8b4b3f8fd76
/dev/sg12   /dev/sdk     9f2a0ae1-d97b-4fb1-873e-6a9bfb2c3254  xfs       9           GOOD     /data-disks/ss-9f2a0ae1-d97b-4fb1-873e-6a9bfb2c3254
/dev/sg13   /dev/sdl     404a8c5a-19c0-4949-bd33-edd83ca4ee8f  xfs       10          GOOD     /meta/404a8c5a-19c0-4949-bd33-edd83ca4ee8f
```

```
dino-black:~ % cs_hal list fs
Volume(s):
SCSI Device Block Device FS UUID                                Type      Slot Label SMART      Mount
----------- ------------ ------------------------------------- --------- ---- ----- ------     ------
/dev/sg2    /dev/sda     0ddb9635-ff27-4cd3-8c2f-58a6f5226d30  ext3           BOOT  GOOD /boot
/dev/sg2    /dev/sda     2192b3ef-2a44-4450-9b04-327c00215454  xfs                  GOOD /root2
/dev/sg2    /dev/sda     ffa9607a-4b6f-4218-9266-c083fb1989a1  xfs                  GOOD /var
/dev/sg2    /dev/sda     746b09d4-f07a-49dc-8b40-86220dfc7edc  xfs                  GOOD /
/dev/sg2    /dev/sda     f7c37c92-4bc5-4abf-95a5-efa51c46f6bc  swap v1              GOOD -
```

volumes (file systems)

```
dino-black:~ % cs_hal list disks
Disks(s):
SCSI Device Block Device Enclosure   Slot Serial Number       SMART
----------- ------------ ----------- ---- ------------------ -----
n/a         /dev/sda     RAID vol    n/a  not supported      n/a
/dev/sg0    n/a          RAID array  0    9QE801ME           GOOD
/dev/sg1    n/a          RAID array  1    9QE834TG           GOOD
/dev/sg3    /dev/sdb     /dev/sg18   0    9WM0R49P           GOOD
```

# COMMUNITY

# Kinetic Open Storage Project, hosted @ Linux Foundation


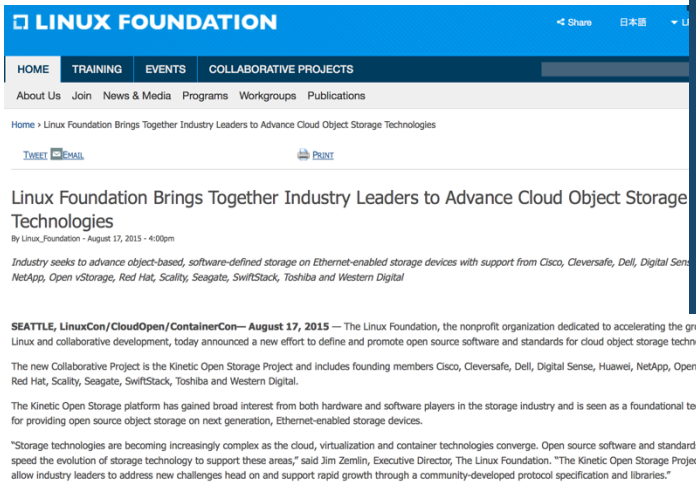

LINUX FOUNDATION

HOME | TRAINING | EVENTS | COLLABORATIVE PROJECTS

About Us | Join | News & Media | Programs | Workgroups | Publications

Home › Linux Foundation Brings Together Industry Leaders to Advance Cloud Object Storage Technologies

TWEET | EMAIL | PRINT

## Linux Foundation Brings Together Industry Leaders to Advance Cloud Object Storage Technologies

By Linux_Foundation - August 17, 2015 - 4:00pm

*Industry seeks to advance object-based, software-defined storage on Ethernet-enabled storage devices with support from Cisco, Cleversafe, Dell, Digital Sens… NetApp, Open vStorage, Red Hat, Scality, Seagate, SwiftStack, Toshiba and Western Digital*

**SEATTLE, LinuxCon/CloudOpen/ContainerCon— August 17, 2015** — The Linux Foundation, the nonprofit organization dedicated to accelerating the growth of Linux and collaborative development, today announced a new effort to define and promote open source software and standards for cloud object storage technologies.

The new Collaborative Project is the Kinetic Open Storage Project and includes founding members Cisco, Cleversafe, Dell, Digital Sense, Huawei, NetApp, Open vStorage, Red Hat, Scality, Seagate, SwiftStack, Toshiba and Western Digital.

The Kinetic Open Storage platform has gained broad interest from both hardware and software players in the storage industry and is seen as a foundational technology for providing open source object storage on next generation, Ethernet-enabled storage devices.

"Storage technologies are becoming increasingly complex as the cloud, virtualization and container technologies converge. Open source software and standards can speed the evolution of storage technology to support these areas," said Jim Zemlin, Executive Director, The Linux Foundation. "The Kinetic Open Storage Project will allow industry leaders to address new challenges head on and support rapid growth through a community-developed protocol specification and libraries."

https://github.com/Kinetic

# Plugfests – Participating Companies

- Seagate
- Toshiba
- Western Digital/HGST
- Sanmina/Newisys
- EMC
- Exablox
- Red Hat
- Swiftstack
- Huawei
- NetApp
- AOL

KINETIC
Open Storage Project

- Plugfest Activity
  - **Multiple** Kinetic drives
  - **Multiple** Kinetic chassis
  - **Multiple** Kinetic client software
- End-to-end operation
  - **All drives present successfully operated in all enclosures present**
  - **Ceph, Exablox, Swiftstack software deployed to successfully read/write objects using all drives present**
- Events
  - December 2015 in Cambridge, MA
  - April 2016 at Vault in Raleigh, NC
  - **20 September 2016 at the SNIA Storage Developer Conference in Santa Clara, CA**
- Join us!

# SUMMARY

# Software-Defined Storage

- Scale-out storage is all about density (PB/rack) and cost ($/TB)
  - achieved by simplicity
  - less components, less cables
  - less code, less layers
- Many deployments need flexibility
  - start small, grow large
  - adjustable compute/storage ratios
  - purchase-time choice *good*; dynamic choice **better**
- Ethernet drives offer this flexibility & scalability

# Learn More

- Join the next Kinetic Plugfest & Open House
  - Tuesday, 20 September 2016
  - SNIA Storage Developer Conference
  - Santa Clara, CA

- 28 February 2017 at FAST in Santa Clara, CA
- 22 March 2017 at Vault in Cambridge, MA

# QUESTIONS?