SDC
STORAGE DEVELOPER CONFERENCE
SNIA ■ SANTA CLARA, 2014

Defining The Software-Defined Technology Market

Mario Blandini
HGST
mario.blandini@hgst.com
@SwiftMario

Forward Looking Statement

This presentation contains forward-looking statements that involve risks and uncertainties, including, but not limited to, the development and adoption of a new storage architecture and the potential introduction of products based on this architecture. Forward-looking statements should not be read as a guarantee of future performance or results, and will not necessarily be accurate indications of the times at, or by, which such performance or results will be achieved, if at all. Forward-looking statements are subject to risks and uncertainties that could cause actual performance or results to differ materially from those expressed in or suggested by the forward-looking statements.

Additional key risks and uncertainties include the impact of continued uncertainty and volatility in global economic conditions; actions by competitors, business conditions and growth in the various hard drive segments. More information about the other risks and uncertainties that could affect our business are listed in our filings with the Securities and Exchange Commission (the “SEC”) and available on the SEC’s website at www.sec.gov, including our Quarterly Report on Form 10-Q filed with the SEC on May 5, 2014, to which your attention is directed. We do not undertake any obligation to publicly update or revise any forward-looking statement, whether as a result of new information, future developments or otherwise, except as otherwise required by law.

Software-Defined Technology

- The new software defined disk drive solution from HGST
- This will be the first time enterprises can run distributed storage/applications directly onto storage media for next generation big data, analytics and research
- How CPU and memory resources residing on these storage devices can be leveraged to run storage services as close to the data as possible

HGST HERITAGE innovations and firsts

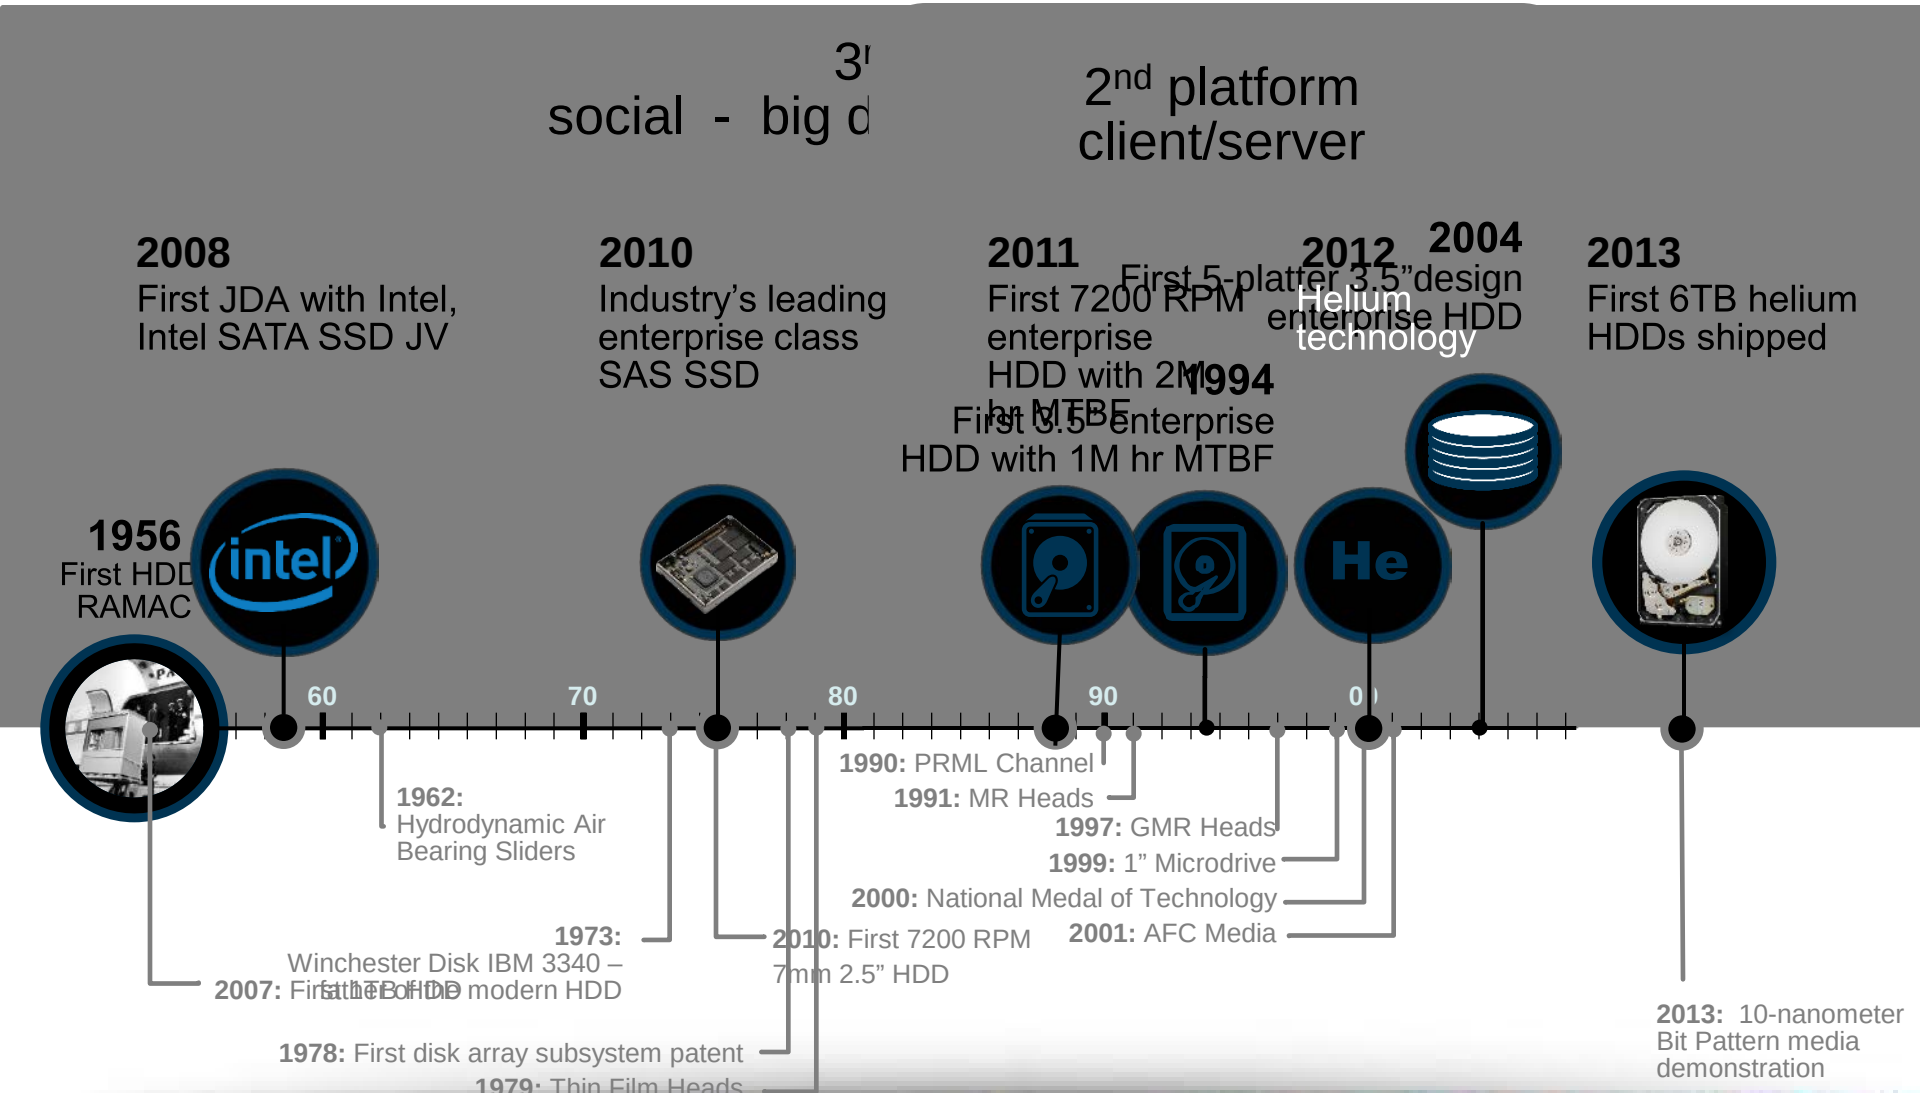

Source: Harvard Business Review, Cisco Systems, Inc.

INCREASING amounts of valuable data

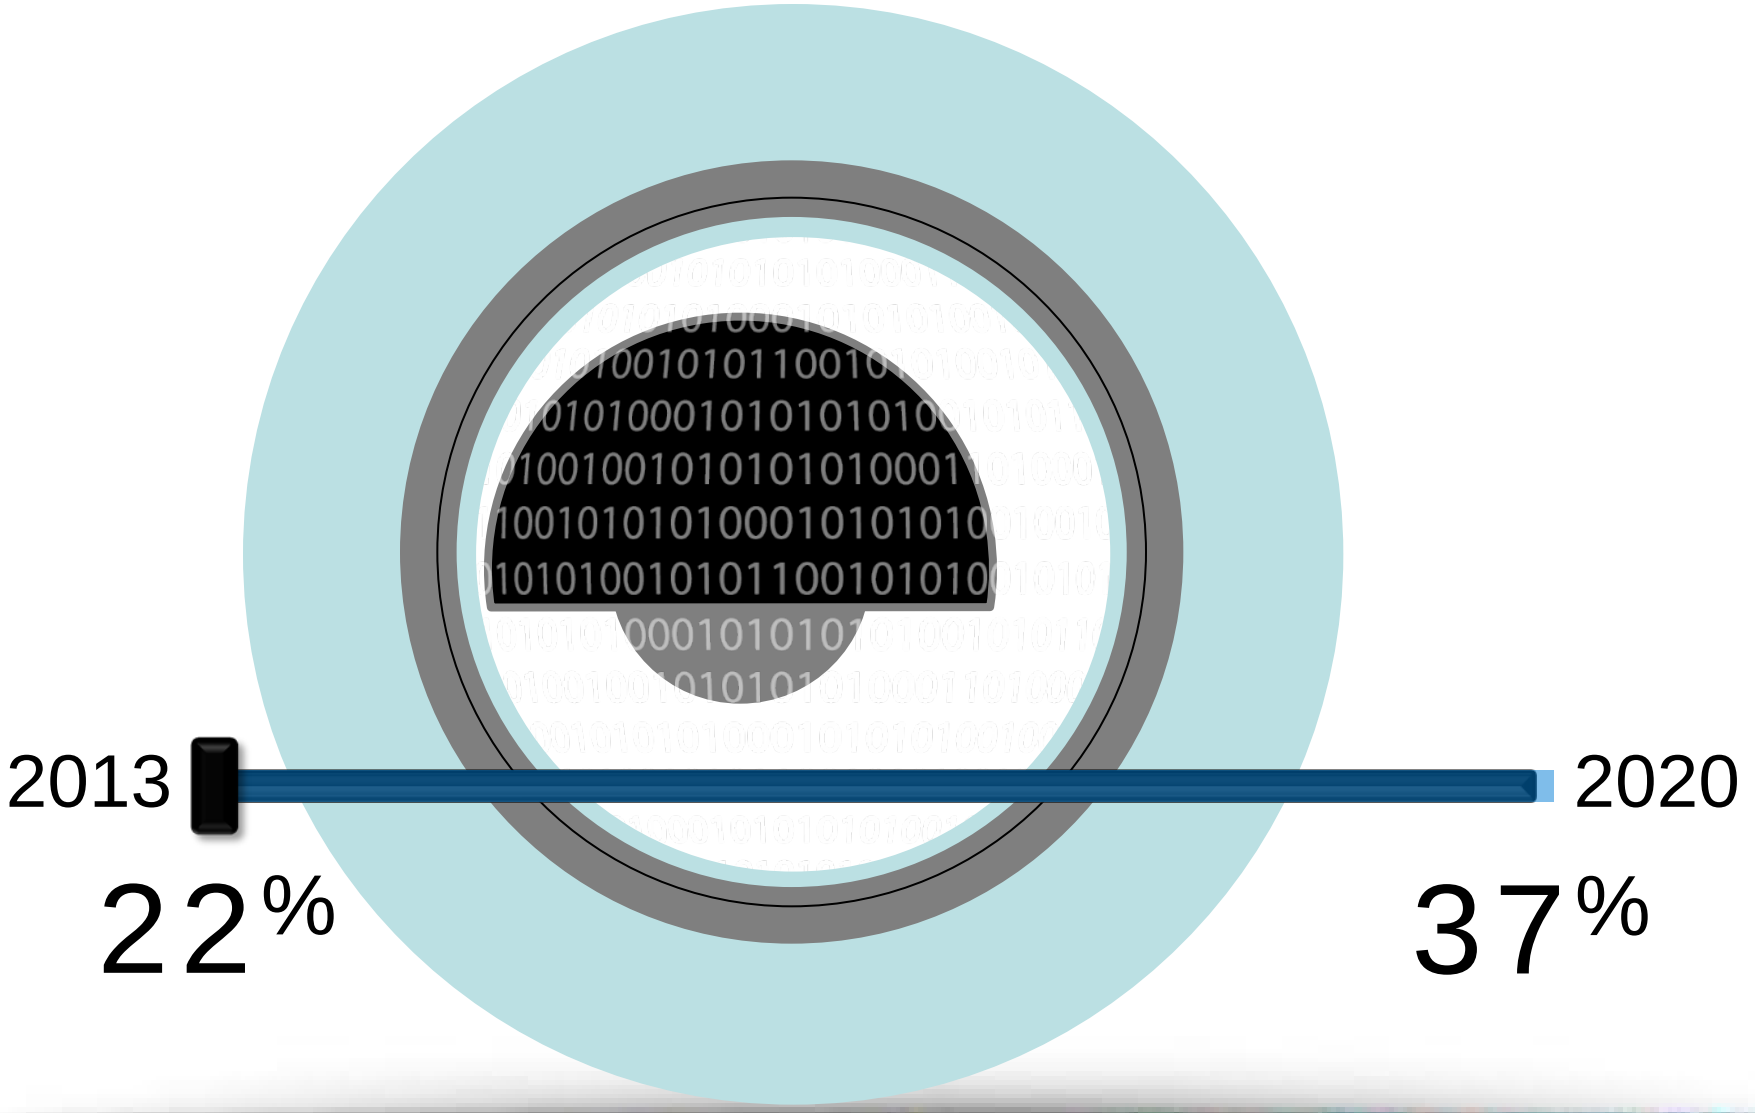

“Active” as #1 requirement in storage

- The value of mining data
  - When data is more readily accessible, it’s value increases
- Typically an analytics or compliance use case
  - Access time is ~1s
- Affordable disk media solutions offer ideal price/performance
  - Software-defined storage

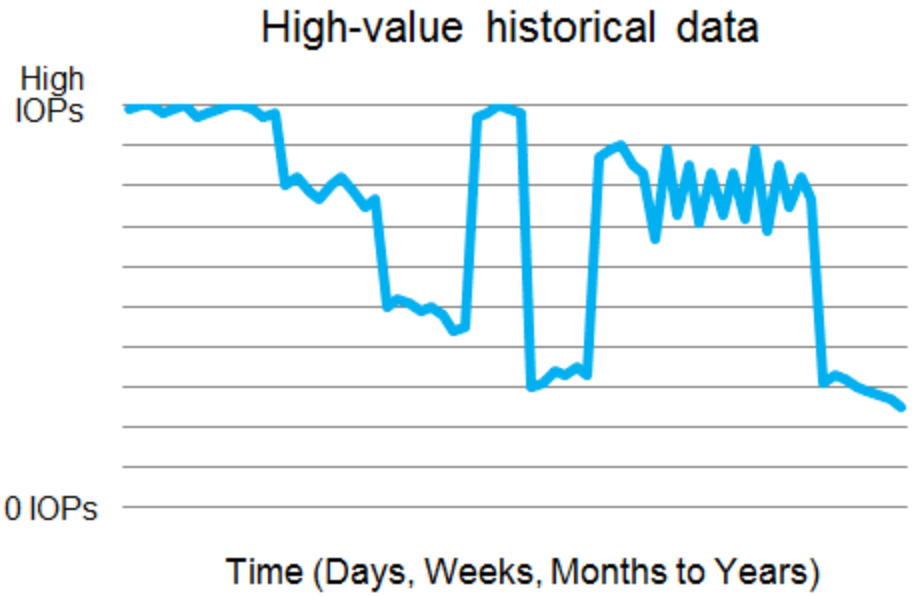

DATA CENTER CHALLENGES:
Scalability, cost, and accessibility

Traditional solution

Deep archive solution

High
SCALABILITY

High
ACCESSIBILITY

REVOLUTIONARY
cost model

INNOVATION leverages device affinity

- Optimized HDDs
- Tuned enclosure
- Device affinity innovations
- Scale out software

VENDOR-SPECIFIC ARCHITECTURES

CONFINE THE SCALE-OUT OF GROWING UNSTRUCTURED DATA

HGST
a Western Digital company

OPEN

ETHERNET
SOFTWARE
DRIVE
HARDWARE
ARCHITECTURE

A Drive Running Linux…

Leverages the Linux ecosystem

A Drive with CPU & RAM…

Runs storage services directly on the device

A Drive with Ethernet…

Connects storage directly into the data center fabric

“Under the Hood” of the Demo Drive

- Standard 3.5” HDD form factor
- Powerful, low-cost integrated SOC
- Ethernet connectivity
- Demonstrating 4TB, but technology can be applied to any HDD, SSD, etc.

Integrated Demo: Enclosure Example

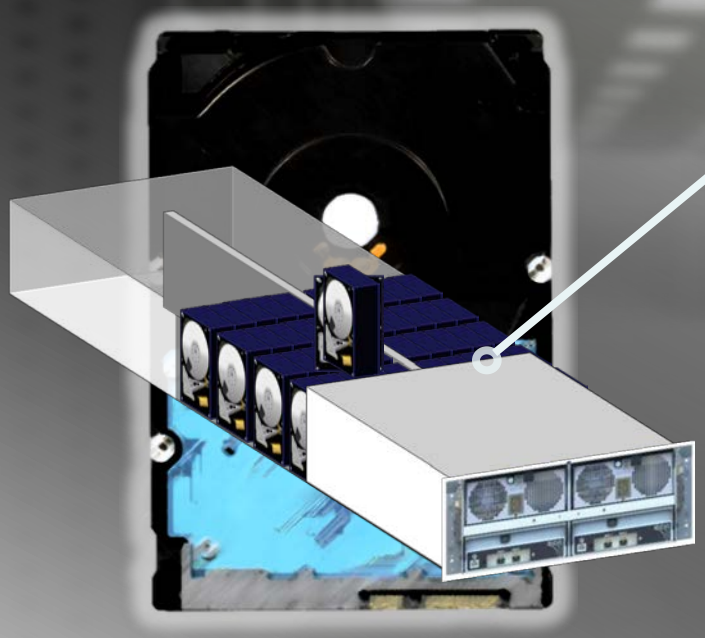

- Reference Design
  - 4U Enclosure
  - 60 drive slots
  - Embedded switched fabric
  - Hot-swap components
- Provides seamless integration into existing datacenter architectures
  - 10Gbps Ethernet ports
  - Devices appear as Linux servers

What do I see as a developer?

- Linux - Debian 7.4+ (Wheezy) on demo system
  - Working on next version of Debian (Jessie) in our labs
- CPU - 32-bit ARM, 512KB Level 2 Cache
- Memory - 2GB DRAM, DDR-3
  - 1792 MB Available to Linux
- Block storage driver
  - Drive enumerates as a SCSI disk (/dev/sda)
- Ethernet network driver
  - Enumerates as eth0 device

Photo Evidence

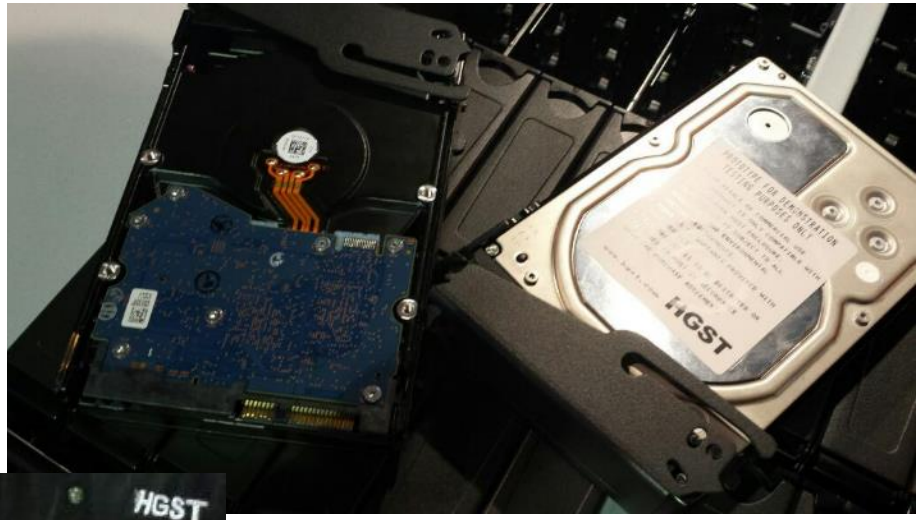

Solutions running as demonstrations

Open Management Framework

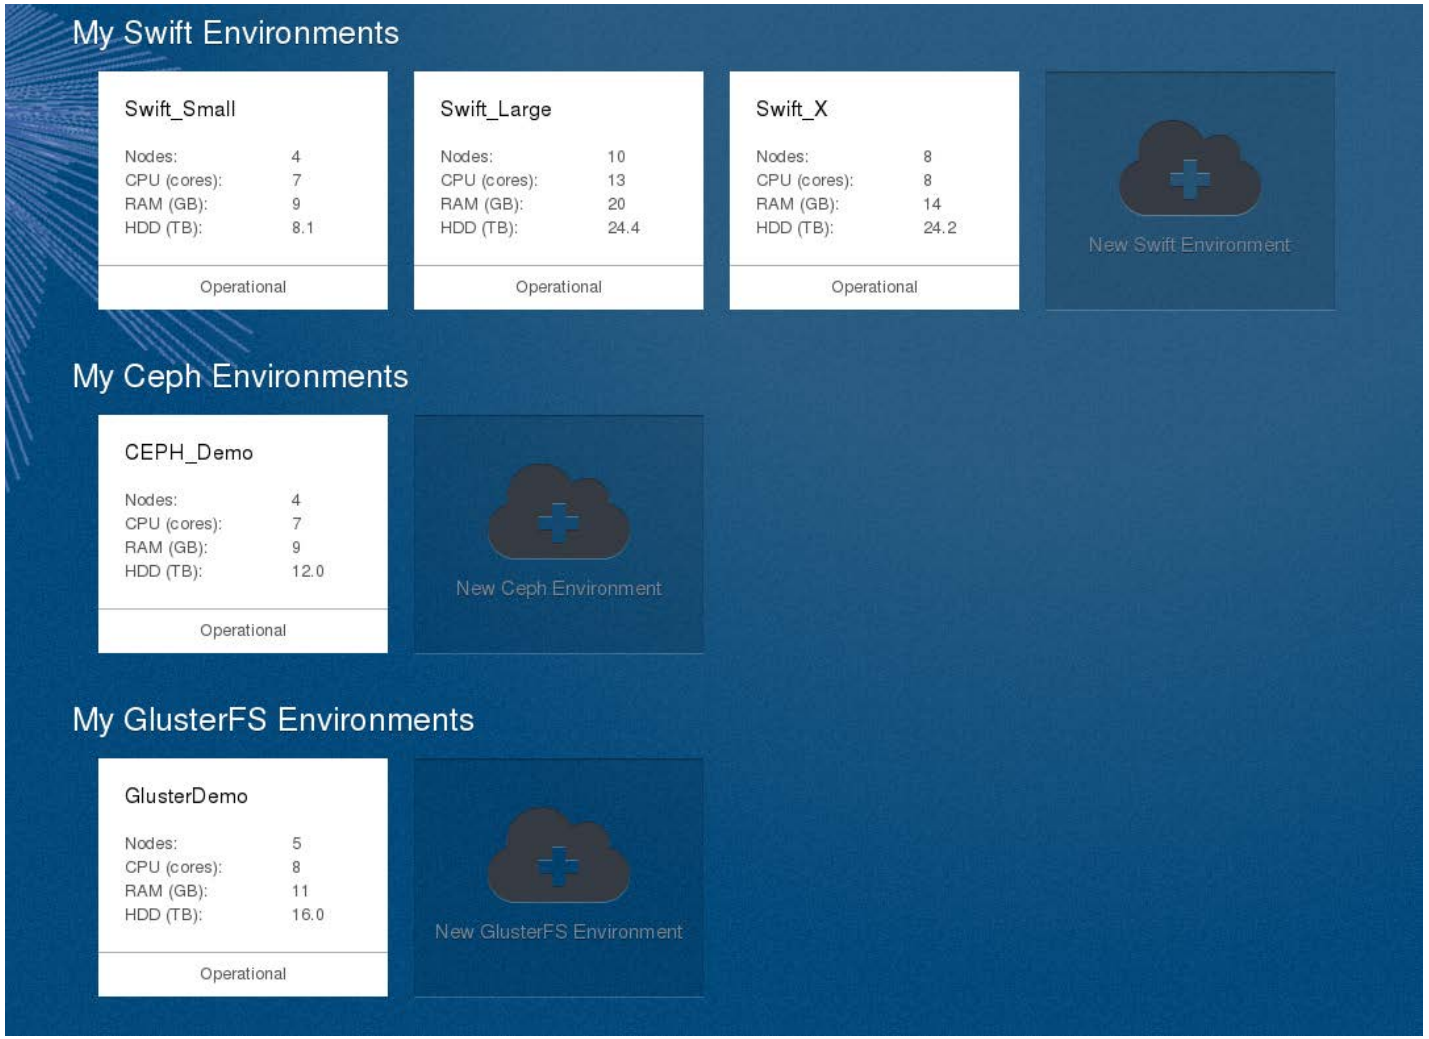

Intermixing Intel & HGST nodes


```
root@clientmachine:~# ssh 10.20.0.166
Last login: Wed May 14 15:39:56 2014 from 10.20.0.10
[root@intel-17-525400880101 ~]# swift-get-nodes /etc/swift/object.ring.gz AUTH_admin 3_30con smallfile

Account         AUTH_admin
Container       3_30con
Object          smallfile


Partition       7540
Hash            eba03c550dfb936e9d68653dff116ece

Server:Port Device      10.2.0.4:6000 sda4
Server:Port Device      10.2.0.3:6000 sda4
Server:Port Device      10.2.0.5:6000 sdb


curl -I -XHEAD "http://10.2.0.4:6000/sda4/7540/AUTH_admin/3_30con/smallfile"
curl -I -XHEAD "http://10.2.0.3:6000/sda4/7540/AUTH_admin/3_30con/smallfile"
curl -I -XHEAD "http://10.2.0.5:6000/sdb/7540/AUTH_admin/3_30con/smallfile"


ssh 10.2.0.4 "ls -lah /srv/node/sda4/objects/7540/ece/eba03c550dfb936e9d68653dff116ece/"
ssh 10.2.0.3 "ls -lah /srv/node/sda4/objects/7540/ece/eba03c550dfb936e9d68653dff116ece/"
ssh 10.2.0.5 "ls -lah /srv/node/sdb/objects/7540/ece/eba03c550dfb936e9d68653dff116ece/"
[root@intel-17-525400880101 ~]# ssh 10.2.0.3
```

OpenStack Swift services running on HGST node


```
[root@intel-17-525400880101 ~]# ssh 10.2.0.3
Last login: Wed May 14 14:56:43 2014 from 10.2.0.2
0 ethdrive-63-000cca00424f:~# ps -ef | grep swift
swift      590     1  0 12:16 ?        00:00:15 /usr/bin/python /usr/bin/swift-account-replicator /etc/swift/account-server.conf
swift      591     1  0 12:16 ?        00:00:09 /usr/bin/python /usr/bin/swift-container-updater /etc/swift/container-server.conf
swift      592     1 10 12:16 ?        00:24:32 /usr/bin/python /usr/bin/swift-object-replicator /etc/swift/object-server.conf
swift      593     1  0 12:16 ?        00:00:00 /usr/bin/python /usr/bin/swift-container-server /etc/swift/container-server.conf
swift      594     1  0 12:16 ?        00:00:00 /usr/bin/python /usr/bin/swift-object-updater /etc/swift/object-server.conf
swift      595     1  0 12:16 ?        00:00:04 /usr/bin/python /usr/bin/swift-container-auditor /etc/swift/container-server.conf
swift      597     1  0 12:16 ?        00:00:00 /usr/bin/python /usr/bin/swift-object-server /etc/swift/object-server.conf
swift      598     1  0 12:16 ?        00:00:00 /usr/bin/python /usr/bin/swift-account-server /etc/swift/account-server.conf
swift      599     1  0 12:16 ?        00:00:00 /usr/bin/python /usr/bin/swift-account-auditor /etc/swift/account-server.conf
swift      600     1  4 12:16 ?        00:09:58 /usr/bin/python /usr/bin/swift-container-replicator /etc/swift/container-server.conf
swift      601     1  0 12:16 ?        00:00:01 /usr/bin/python /usr/bin/swift-object-auditor /etc/swift/object-server.conf
swift      777   601  0 12:16 ?        00:01:48 /usr/bin/python /usr/bin/swift-object-auditor /etc/swift/object-server.conf
swift      851   598  0 12:16 ?        00:00:07 /usr/bin/python /usr/bin/swift-account-server /etc/swift/account-server.conf
swift      857   593  2 12:16 ?        00:06:01 /usr/bin/python /usr/bin/swift-container-server /etc/swift/container-server.conf
swift      859   597  6 12:16 ?        00:15:18 /usr/bin/python /usr/bin/swift-object-server /etc/swift/object-server.conf
swift    16905   591  0 16:03 ?        00:00:00 /usr/bin/python /usr/bin/swift-container-updater /etc/swift/container-server.conf
root     16908 16797  0 16:03 pts/0    00:00:00 grep swift
0 ethdrive-63-000cca00424f:~#
```

Ceph OSD running on HGST node

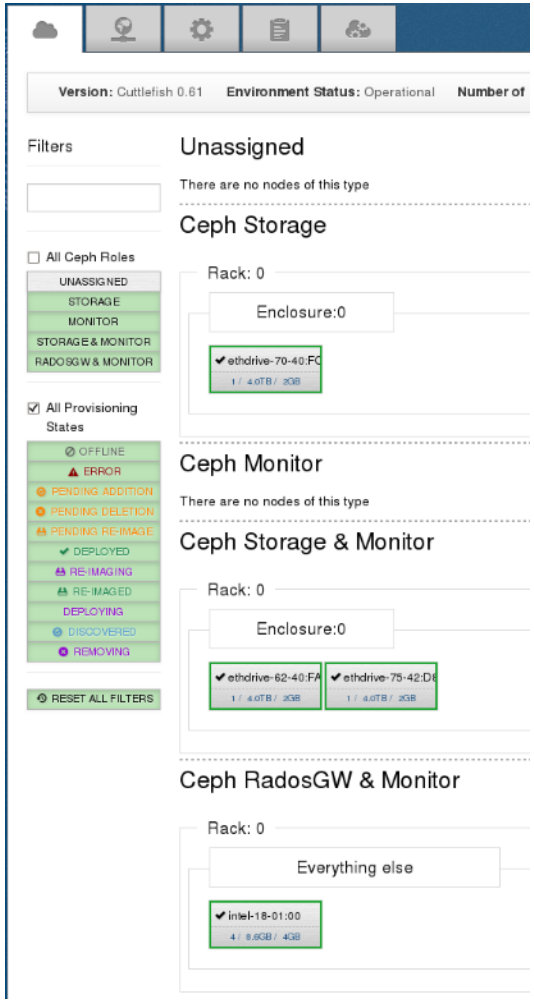

Bricks services running on HGST node

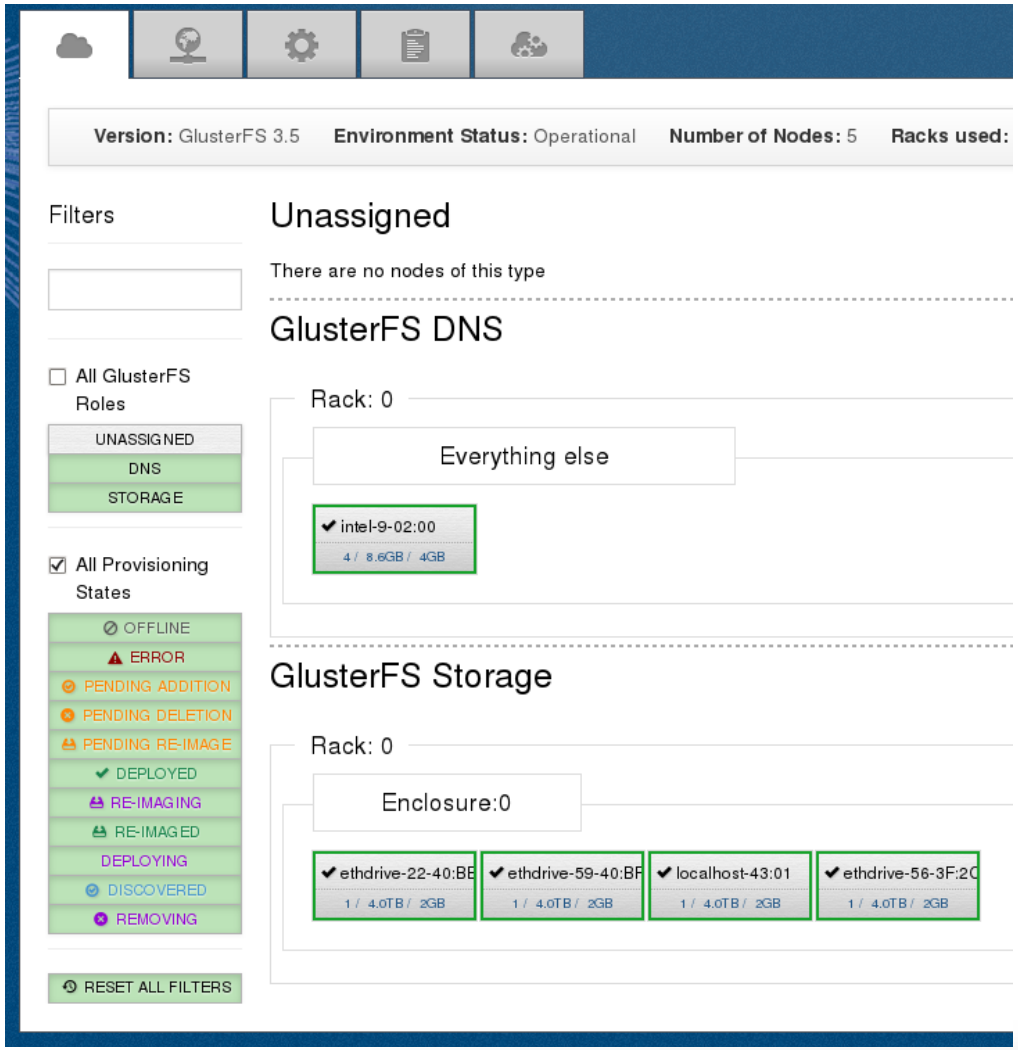

Want to Develop with HGST?

Contact me to learn more!

HGST Goal = Work with of cloud software to (re)define the data center of the future

Or, learn how to become a developer at http://www.hgst.com/opendev

Software-Defined Technology

- The new software defined disk drive solution from HGST
- This will be the first time enterprises can run distributed storage/applications directly onto storage media for next generation big data, analytics and research
- How CPU and memory resources residing on these storage devices can be leveraged to run storage services as close to the data as possible

Thank You

Mario Blandini
HGST
mario.blandini@hgst.com
@SwiftMario