IBM® DeepFlash™
Elastic Storage Server

Nov 2016

Well known trends in the data storage world

Data growth continues to be explosive - driven by digital transformation of almost every industry and every business function

Challenges with the current infrastructure:

- × Low performance density (HDD throughput remains low while HDD capacity expands)
- × Performance can be insufficient
- × High power consumption
- × High floor space requirement
- × High cooling costs
- × High disk failure rates

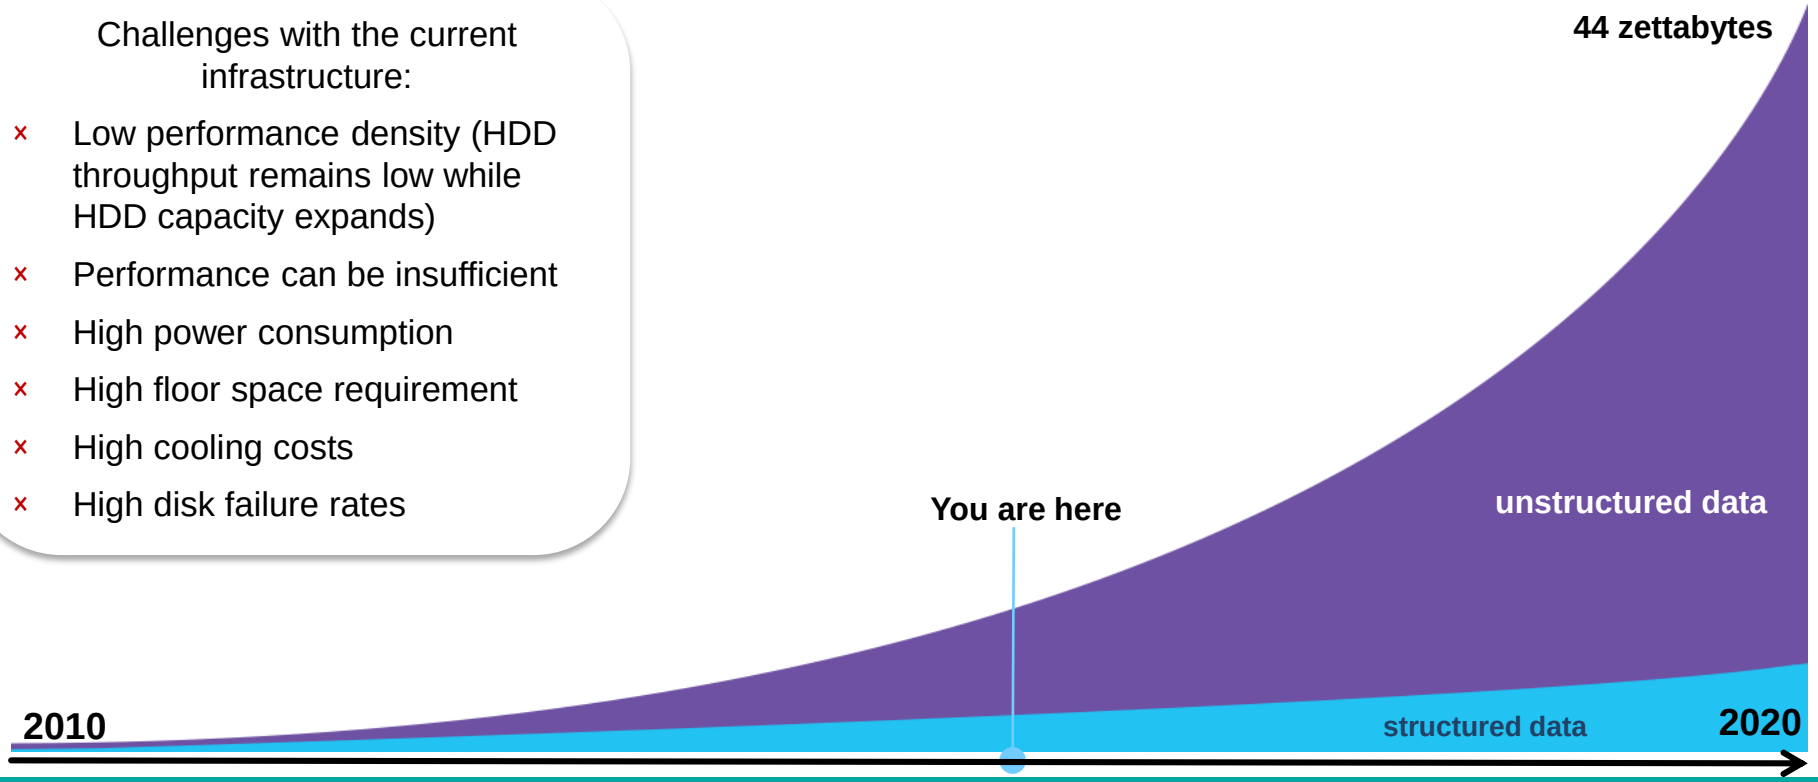

Solution: Augment with Flash

If you need speed (response time and/or throughput) instead of space.

- 8x faster application response time
- Increase throughput/rack unit by 2.8X
- Reduce MTBF
- Reduce power & cooling costs by 30%-50%

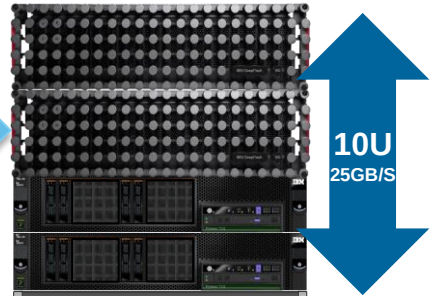

8X faster response time and same throughput as the HDD version

Introducing IBM DeepFlash™ Elastic Storage Server

ESS GF1

7U, 180TB Usable Flash

Read **13.6 GB/sec**

Write **9.3 GB/sec**

ESS GF2

10U, 360 TB Usable Flash

Read **26.6 GB/sec**

Write **16.6 GB/sec**

New Class of Flash: Big Data Flash

Scalable capacity and performance at low price points for big data

	HDD	DeepFlash ESS	Conventional Flash
Price	$	$$	$$$
performance	10's of milliseconds	Sub Milliseconds	Micro Seconds
Attributes		• High ingest rate • Low change rate • High read rate	• Extremely latency sensitive • Can justify price premium

- Performance consistently better than that of the best HDDs
- Cost comparable to that of performance optimized HDDs
- Systems implementations that support massive scalability and meet enterprise requirements

Hardware overview of DeepFlash

Higher Reliability

- 1.5+ million hours MTBF (mean time between failure)
- Hot-swappable architecture - easy FRU of fans, SAS expander boards, power supplies, flash cards
- Directly samples air temp

Higher Density

- 30” x 3U chassis starting at 128TB up to 256TB (512TB on roadmap)
- 8 to 64 8TB flash cards (Board Solid State Drive) presented as SCSI targets - JBOF behind redundant SAS expanders

- 150W(idle), 750W(abs -max), typical workload 450W
- 30% to 50% lower power consumption and cooling power requirement of an equivalent HDD array

Lower Energy Consumption

Game changing performance

- Up to 2M raw IOPS
- sub 1ms latency
- 12GB/s throughput

ESS Solution Integration

- Scale-out storage software that scales Exabyte's of DeepFlash arrays under a single name space
- Unified data access including, file, object, HDFS and OpenStack
- Seamlessly tiering across Flash, Disk, Tape and Cloud Storage
- Efficient Space-Saving Compression
- Low overhead encryption

The Building Blocks of DeepFlash™ Elastic Storage Server

Spectrum Scale	• High performance Parallel File System • Seamless scaling of performance & capacity • Rich Data Management Features
Spectrum Scale RAID	• 2 and 3 parity distributed RAID • Disk Hospital for proactive disk management • End2End Data Checksum Tracking
DeepFlash 5147	• 256TB of Flash in 3U • Cost Effective High Density Flash

Use cases for IBM DeepFlash™ Elastic Storage Server

DeepFlash Elastic Storage Server is targeted at clients who value faster application response times and/or throughput per rack unit

Best suited for data intensive workloads that have:

- High ingest rate
- High and random reads
- Low change rate

Use-cases with the above workload characteristics are in every industry:

- Big Data Analytics across industries
- Video Editing and Streaming
- Genome Sequencing analysis
- Health care data analytics
- Mobile "Humvee" data centers for portable data collection and analysis

We've Listened: Recent GNR/ESS Improvements

- Focus on quality, support and superior customer experience
  - Improved flow for Level 2 support
  - Knowledge Center reworked to be more intuitive
  - Simplification of performance tunables
  - Phase 1 hardware callhome: components in I/O and management nodes (working on callhome for disks/enclosures)
- ESS 5.0
  - RedHat ppc64 Little Endian (LE) support (no need for HMC)
  - Supports Spectrum Scale 4.2.2
  - Performance enhancements – Hardware and software updates (ie LSI 12Gb adapter)
  - gssprecheck – catches common installation and upgrade issues prior to start, decreases chance of deployment errors and further reduces the time required for system bring up (found in sample directory)
  - Phase 1 Security hardening: Split of /home and /var from /root partition and security flags added in /etc/fstab

Backup

gssprecheck Sample Output

```
#  ./gssprecheck -G gss_ppc64 --upgrade --file ./gssdeploy.cfg
2016-10-28T13:23:43.402218 >>>ESS500_BETA_2<<< Start of pre-install check
2016-10-28T13:23:43.402285 This may take a few minutes. Please be patient
2016-10-28T13:23:45.232631 nodelist:    c55f04n04 c55f04n05
================== Summary of EMS node ============================
[OK] Parsing configuraton file
[OK] Checking xCAT version
[OK] Checking xCAT site table
[OK] Checking xdsh connectivity
[ERROR] Bonded link check
[ERROR] Spectrum Scale lock check
[OK] Checking deploy iface
[OK] Timezone consistency check
[OK] Universal time consistency check
[OK] Quorum node check
[OK] long waiters check
[ERROR] mmhealth health check
[ERROR] mmhealth eventlog check
      >>Running gnrhealthcheck...This will take a few moments<<
[ERROR] GNR health check
[OK] Manifest check
[OK] Checking /etc/hosts exists
[OK] /etc/hosts advanced checks
[OK] Checking for general repo errors
[OK] Checking for enabled external subscriptions
[OK] Checking kernel repo
[ERROR] Checking correct redhat version
[OK] Checking correct Endian type
[ERROR] High CPU % process found
[OK] Checking for servicable events
[OK] Root FS space check
[ERROR] Checking that tracing is disabled
[OK] Checking resolv.conf
[OK] Active Node Check
[OK] Checking for deadlocks
```

Inside the IBM DeepFlash 150

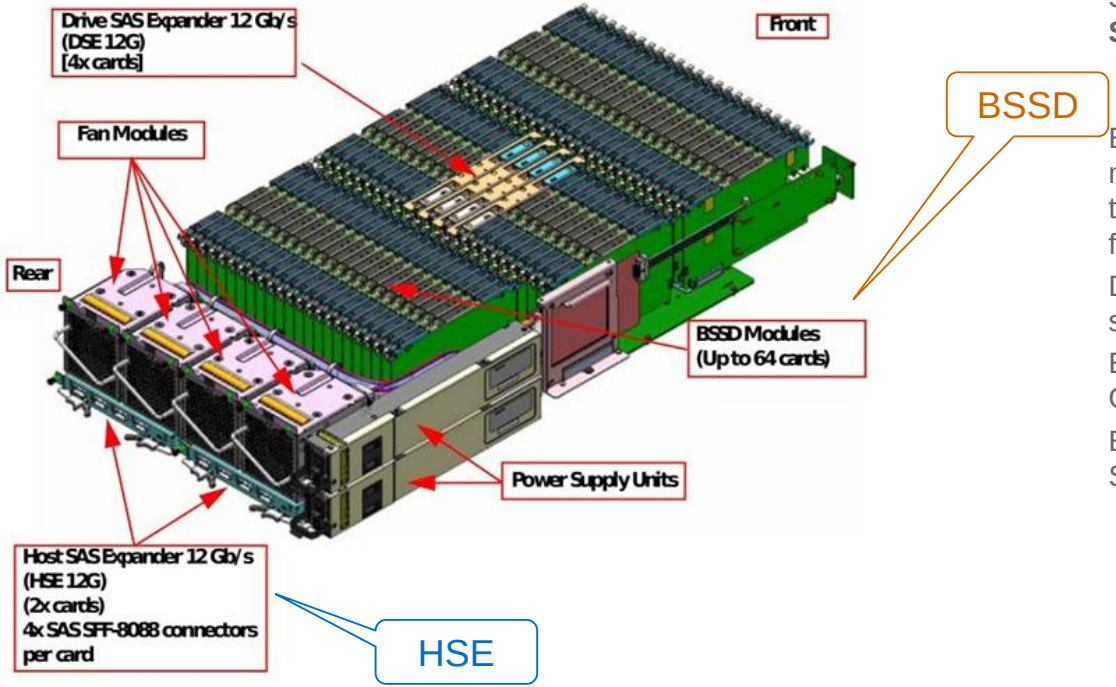

Storage elements are called a **Board Solid State Drive** (BSSD)

- The BSSD is a cost-optimized implementation of an SSD in a different form factor.

Each BSSD provides 8TB of raw flash, with the maximum 64 BSSDs in 3U chassis delivering a total capacity of 512TB (initial offering of 256TB for ESS 5.0).

DeepFlash 150 enclosures are available in three sizes: 128TB, 256TB and 512TB.

External connectivity with DeepFlash 150 is 12 Gb/s SAS (SAS 3.0).

Each DeepFlash 150 enclosure has two Host SAS Expanders (HSEs)

- Each with four mini-SAS connections for interfacing to:
- One or more external servers through SAS 6Gbit links using SAS cables.

Existing HDD/SSD based Elastic Storage Server Building Blocks

GS models use 2U 24x2.5” JBODs; GL models use 4U 60x3.5” JBODs

Supported drives: 1.8TB SAS, 400GB, 800GB, 1.6TB SSD 2.5” and 2TB,4TB,6TB and 8TB NL-SAS 3.5” HDDs

Supported NICs: 10GbE, 40GbE Ethernet and EDR Infiniband

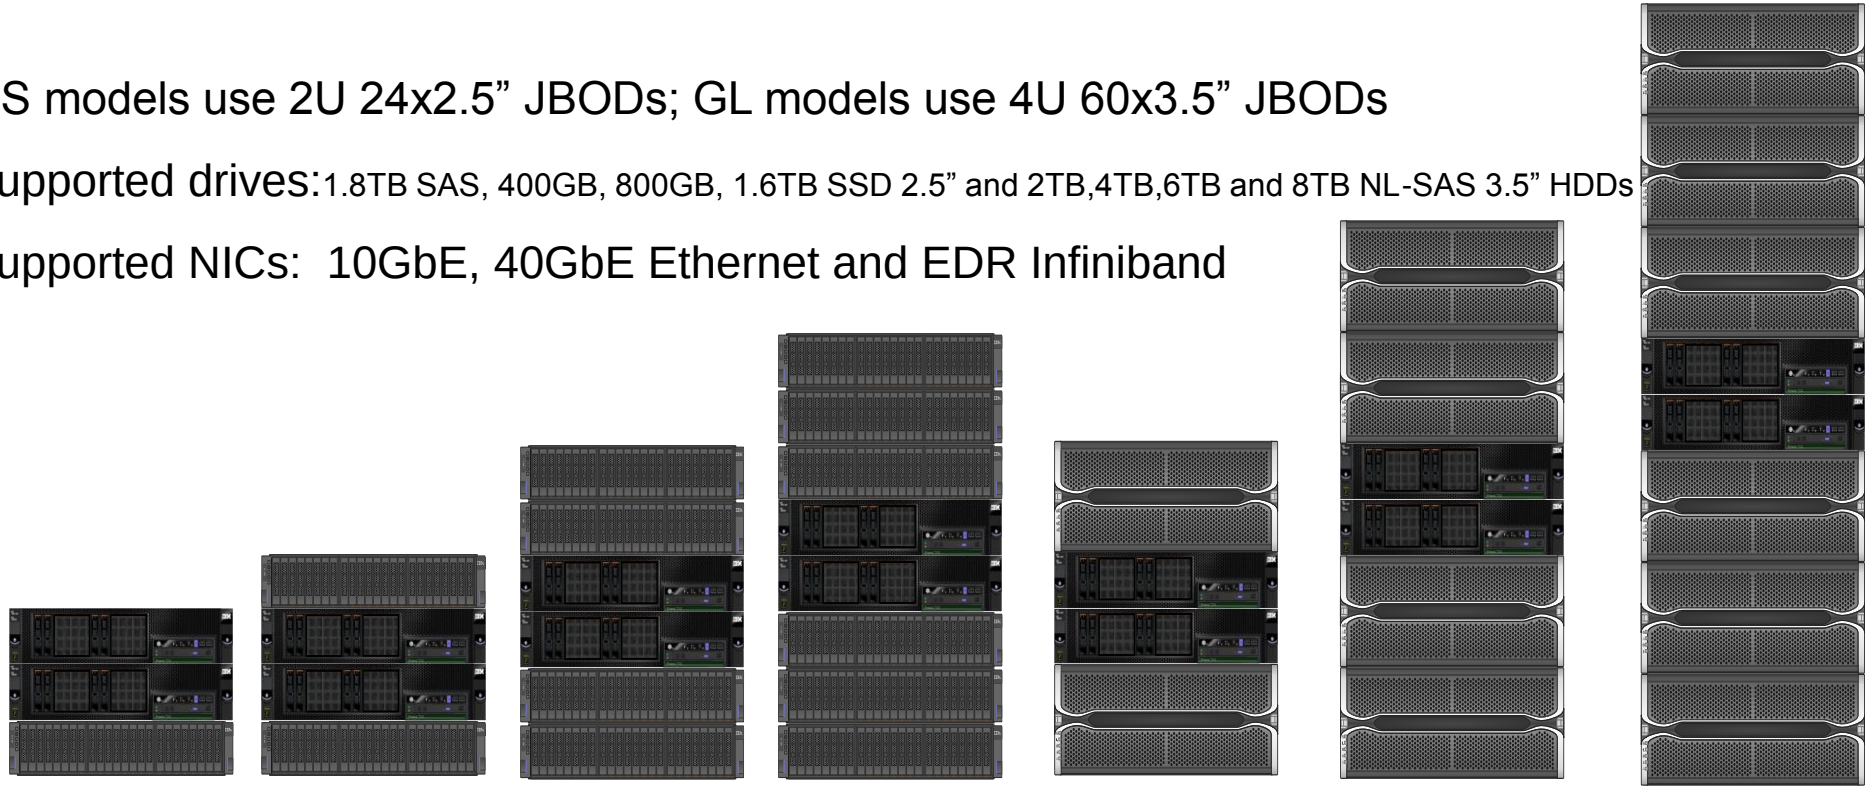

GS1 building block | **GS2 building block** | **GS4 building block** | **GS6 building block** | **GL2 building block** | **GL4 building block** | **GL6 building block**