

2017 IBM Elastic Storage Server - Update

Falk Steinbrueck

Program Manager, IBM Systems Storage



Disclaimer

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.

The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

The emerging era of cognitive and cloud underpinned by pervasive technical computing is upon us. IBM Spectrum Scale and ESS are well positioned to respond to this market opportunity.



▪ Largely Block Data

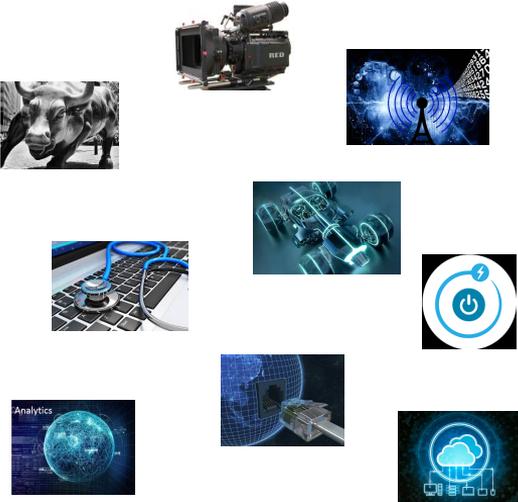
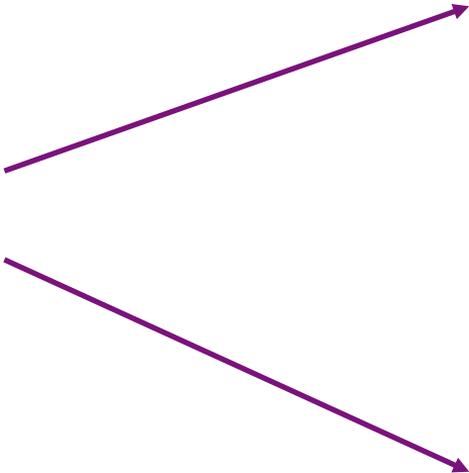
▪ Largely Unstructured Data

- Benefiting from simplified infrastructure
- Requiring cost efficiency through improved virtualization and automation
- Driving controlled data growth
- Requiring massive scale and rapid pace
- Accelerating business insights
- Relying upon data access and elasticity

The Elastic Storage Server (ESS) is positioned to meet the next market wave - the emerging era of cognitive and cloud underpinned by pervasive technical computing.



Technical Computing
Foundation



Cognitive Solutions and
Cloud Platforms Market

IBM Elastic Storage Server (ESS)

Integrated scale out data management for file and object data

IBM Storage & SDI

*The Choice for CORAL!
ORNL 250 PB of ESS Storage*

Optimal building block for high-performance, scalable, reliable enterprise storage

- Faster data access with choice to scale-up or out
- Easy to deploy clusters with unified system GUI
- Simplified storage administration with IBM Spectrum Control integration

One solution for all your data needs

- Single repository of data with unified file and object support
- Anywhere access with multi-protocol support:
NFS 4.0, SMB, OpenStack Swift, Cinder, and Manila
- Ideal for Big Data Analytics with full Hadoop transparency

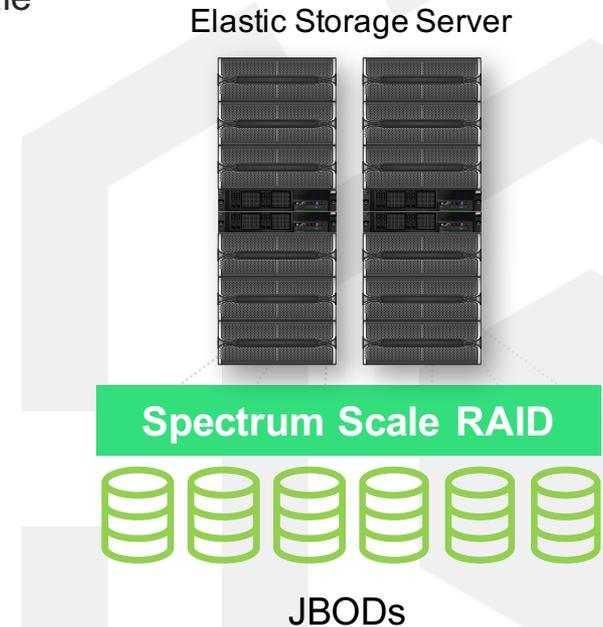
Ready for business critical data

- Disaster recovery with synchronous or asynchronous replication
- Ensure reliability and fast rebuild times using Spectrum Scale RAID's dispersed data and erasure code



Advantages of Spectrum Scale RAID

- **Use of standard and inexpensive disk drives**
 - 8+2 or 8+3 Erasure Code software implemented in Spectrum Scale
- **Sustained Data Availability**
 - Can survive full enclosure failures or full server maintenance
 - Critical Rebuilds completed in Minutes versus Days!
- **Minimal impact of rebuild on system performance**
 - Rebuild is done across all disks in system
 - Disk failure is a non-maintenance event if desired.
- **Better fault tolerance**
 - End to end checksum protects data all the way to client
 - Much higher mean-time-to-data-loss for large configs (MTTDL)
 - 8+2P: ~ 200 Years
 - 8+3P: ~ 200 Million Years



Elastic Storage Server (ESS) Family



SpectrumScale
IBM Power
Linux Server



"Twin Tailed" JBOD
Disk Enclosure
X TB Drives

← Capacity →

Model GL2: 2 Enclosures, 12U
116 NL-SAS, 2 SSD

Model GL4: 4 Enclosures, 20U
232 NL-SAS, 2 SSD

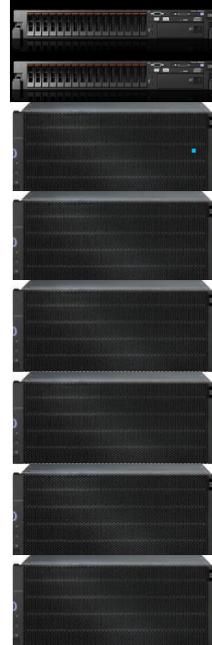
Model GL6: 6 Enclosures, 28U
348 NL-SAS, 2 SSD



6+ GB/sec*



13+ GB/sec*



25+ GB/sec*

← Speed →

Model GS1: 24 SSD

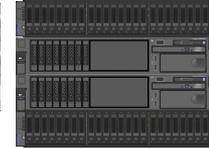
Model GS2: 46 SAS + 2 SSD
or 48 SSD Drives

Model GS4: 94 SAS + 2 SSD
or 96 SSD Drives

Model GS6: 142 SAS
+ 2 SSD Drives

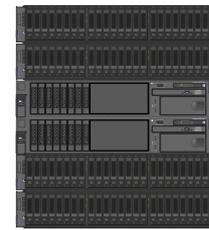


6+ GB/sec*



2 GB/sec SAS*

12 GB/sec SSD*



5 GB/sec SAS*

18+ GB/sec SSD*



7 GB/Sec*

All Performance Estimates are dependent on workload, block size and Infiniband networking connections

Guideline Usable Space includes allowance for Erasure Coding, Metadata, and Spare Disk Space. These can vary based upon customer and workload requirements.

We've Listened: Recent GNR/ESS Improvements

- Focus on quality, support and superior customer experience
 - Improved flow for Level 2 support (worldwide launch April 2017!)
 - Knowledge Center reworked to be more intuitive (https://www.ibm.com/support/knowledgecenter/SSYSP8/sts_welcome.html)
 - Simplification of performance tunables
 - Call Home Stage 1: for components in I/O and management nodes
- ESS 5.0
 - **RedHat** ppc64 Little Endian (LE) support (no need for HMC); Update to RHEL 7.2
 - **Spectrum Scale** 4.2.2 support
 - **Call Home** Stage 2: disc
 - **Performance** enhancements – Hardware and software updates (ie LSI 12Gb adapter)
 - **Installation:** 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)
 - **Security** Phase 1 hardening: Split of /home and /var from /root partition and security flags added in /etc/fstab

Overall ESS Directions

Continue to expand boundaries of high performance

- All Flash Optimized models for high File IOPS
- Deliver Petascale Flash-Only ESS Systems
- Support “Integrated Burst Buffer” configurations (Coral)
- Support high speed NICs
- Big Data Analytics validated configurations

Leveraging Flash and Server Cache

Simplify the Exascale

- *Call Home* for automated problem identification and maintenance
- Failed parts auto shipped
- Simplify troubleshooting problems by admin
- Simplify debugging and fixing Network Performance
- Enhanced Protocol Support

Leveraging RAS at scale

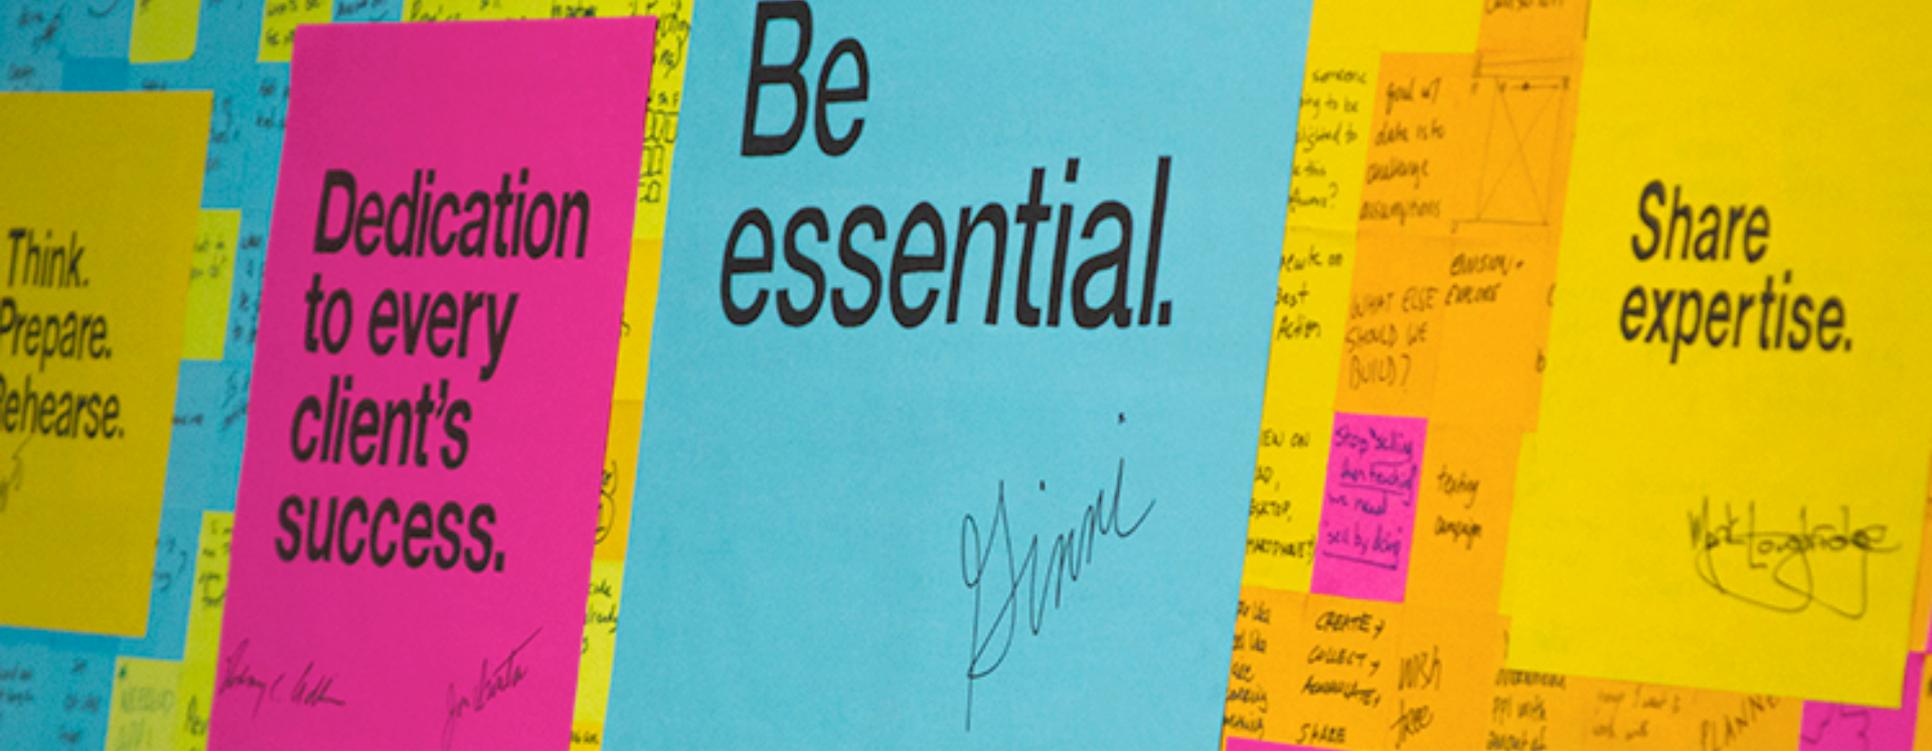
Continue to bring down TCO

- Deploy Higher Density Enclosures
- Support higher capacity drives inline with industry
- Support hybrid (Flash and Disk based) ESS models
- Flexibility in choosing HDD/SSD/Flash and RAM

Leveraging GNR Software

Welcome Request for Enhancement (RFE) Community users! Here you have an opportunity to collaborate directly with the IBM product development teams and other product users.

The screenshot shows the 'developerWorks' interface for RFEs. At the top, there's a header 'developerWorks®'. Below it, a filter section titled 'Filter the page content by brand and product' contains two dropdown menus: 'All brands' and 'All products'. The 'All products' dropdown is open, showing a search input with 'scale' and a list of results: 'Spectrum Scale (formerly known as GPFS) - Private RFEs', 'Spectrum Scale (formerly known as GPFS) - Public RFEs', and 'WebSphere eXtreme Scale'. The second option is highlighted. Below the filters, there are sorting options: 'Hot', 'Top', and 'New'. A card for an RFE is visible, showing '361 votes' and a title 'Add additional member to an existing...'. The description starts with 'The current Storawize V7000 code... increasing storage capacity, it would...'. Below the description, there's a section for 'Uncommitted Candidates' with a dropdown menu showing 'GSS' and 'GPFS Storage Server (GSS)', with the second option highlighted.



Falk Steinbrueck



Program Manager – Software Defined Storage
steinbrueck@de.ibm.com

Thank You.
IBM Storage & SDI



gssprecheck

Sample Output

```
# ./gssprecheck -G gss_ppc64l --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
```