


IBM Spectrum LSF and IBM Spectrum Scale User Group

What is new in Spectrum Scale and ESS?

Christopher D. Maestas
cdmaestas@us.ibm.com



Please Note

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



Notices and disclaimers

- © 2019 International Business Machines Corporation.
No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.**
IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts.
In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and disclaimers continued

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml.

Spectrum Scale Updates



QoS improvements in large clusters

QoS node collects stats X time and sends to
QoS manager node Y time

Report via **mmlsqos**

Large clusters => more communication,
performance degradation

Set new defaults based number of mounts

Allow changes

stat-slot-time : QoS collects

stat-poll-interval : QoS -> QoS manager

Table 1. Default intervals for collecting and sending statistics		
Number of nodes that have mounted the file system	Interval between collecting statistics, in milliseconds	Interval between sending statistics to the QoS manager, in seconds
< 32	1000	5
< 64	2000	10
< 128	3000	15
< 256	4000	20
< 512	6000	30
< 1024	8000	40
< 2048	10000	50
< 4096	12000	60
< 8192	12000	60
< 16384	12000	60
16384 or more	24000	120

mmchqos Device --enable [--stat-poll-interval Seconds] [--stat-slot-time Milliseconds]

Network enhancements

mmhealth checks
availability,
port state,
link state
but not connectivity

mmnetverify checks
RDMA connectivity
(between nodes)

nsdperf for “stress” testing
now opensource!

https://github.com/IBM/SpectrumScale_NETWORK_READINESS/blob/master/nsdperf.C

Attempt to reconnect socket before expel

- Only on Linux

mmchconfig *proactiveReconnect=yes*

Raise network reconnects to **mmhealth**

mmhealth node eventlog

Timestamp	Event Name	Severity	Details
2019-TIME TZ	reconnect_start	WARNING	Attempting to ...
2019-TIME TZ	reconnect_done	INFO	Reconnected to ...

mmhealth node eventlog

Timestamp	Event Name	Severity	Details
2019-TIME TZ	reconnect_start	WARNING	Attempting to ...
2019-TIME TZ	reconnect_failed	ERROR	Reconnect ... failed
2019-TIME TZ	reconnect_aborted	INFO	Reconnect ... aborted

Spectrum Scale misc.

Deprecate primary and backup server

mm[cr,ch]cluster

Grab security/encryption and/or network related data for **gpfs.snap**

Display certificate expiration

mmkeyserv server show

designate license with **mmaddnode**

mmaddnode -N *name:manager:name-a:server* --accept

Log better CCR messages

6027-4200 [E] Maximum number of retries reached

6027-4201 [B] Version mismatch on conditional put

6027-4202 [B] Version match on conditional get

6027-4203 [B] Invalid version on put

6027-4204 [E] Not enough CCR quorum nodes available

6027-4205 [E] ccr.nodes file missing or empty

6027-4206 [E] CCR is already initialized

6027-4207 [E] Unable to reach any quorum node (Check your firewall or network settings)

CCR recovery options

1) # **mmsdrrestore** -p <QNODE_WITH_GOOD_CCR_COPY>

2) # **create mmsdrbackup && mmsdrrestore -F /x/f -a**

3) # **mmsdrrestore --ccr-repair**

There is a dry-run mode
GPFS must be down

Updates mmhealth

CES with
SUDO wrapper and SE-Linux

Colorized output

Thresholds monitor, which collector

mmhealth cluster show threshold -v

Determine CES IP failover

mmhealth node eventlog | grep move_cesips

ESS monitoring of
pdisk, fan speed, new enclosures

NVME monitoring

Watchfolder monitoring

If in doubt with node state

mmhealth node show --resync

5.0.3 Spectrum Scale GUI –What's new

No default admin user

/usr/lpp/mmfs/gui/cli/mkuser ADMINUSER -g SecurityAdmin

Configure LDAP for GUI USER from GUI

Can test connectivity

Manage Quotas

user, group, fileset
capacity and inode quotas
any other setting
(scope, grace time)

Email daily quota reports

NFS client management

Support pseudo paths NFSv4

Better monitoring for NFS exports
and SMB shares

Manage NFS/SMB authentication!

Migrate policy to external pool with
best practice excludes

.ltfsee, ,snapshots, .mmbackup
small files
recently access files
migrated files

Install Toolkit 5.0.3 New Features

Recall install toolkit introduced in 4.1.1.0

GUI installtoolkit is being deprecated

Upgrade flow changes to minimize I/O disruptions

Use to change product editions

Mixed O/S support

Pre-checks what packages must be upgrade
(and if missing dependency)

Post-checks ensure all upgrades successful

Product edition change path	Installation toolkit (cluster online or offline)	Manual node by node (cluster online)	Manual all nodes (cluster offline)
Standard Edition to Data Access Edition	Yes	Yes	No
Standard Edition to Data Management Edition	Yes	No	Yes
Standard Edition to Advanced Edition	Yes	No	Yes
Advanced Edition to Data Management Edition	Yes	Yes	No

```
# cd /usr/lpp/mmfs/VERSION/installer
```

```
# ./spectrumscale config populate
```

IBM Spectrum Scale AWS Quickstart vs Marketplace

IBM Spectrum Scale Offering	AWS Quick Start	AWS Marketplace
Product Version	4.2.3.7	5.0.3.0
RHEL Version	7.4	7.6
Offering	90 days free trial	BYOL
Support for Monitoring	No	Yes
Support for Cluster Expansion/Contraction	No	Yes
Support for Lifecycle Management	No	Yes
Support for enhanced debug data collection	No	Yes
Enhanced cluster tuning settings applied	No	Yes
Customer Support	Via email: <ol style="list-style-type: none">1. IBM Spectrum Scale forum2. IBM Spectrum Scale mailing list	Through IBM Spectrum Scale Data Management Edition Support contract: https://www.ibm.com/support/home/ Phone: 1-800-IBM-SERV (800-426-7378)

Clustered Watch

Captures file system activity

Generates an event notification for that activity

Streams the notifications to topics within the message queue

Events are consumed by a conduit

Conduit sends these events to an external 'sink'

External sink should be a Kafka message queue setup and managed independently by the customer

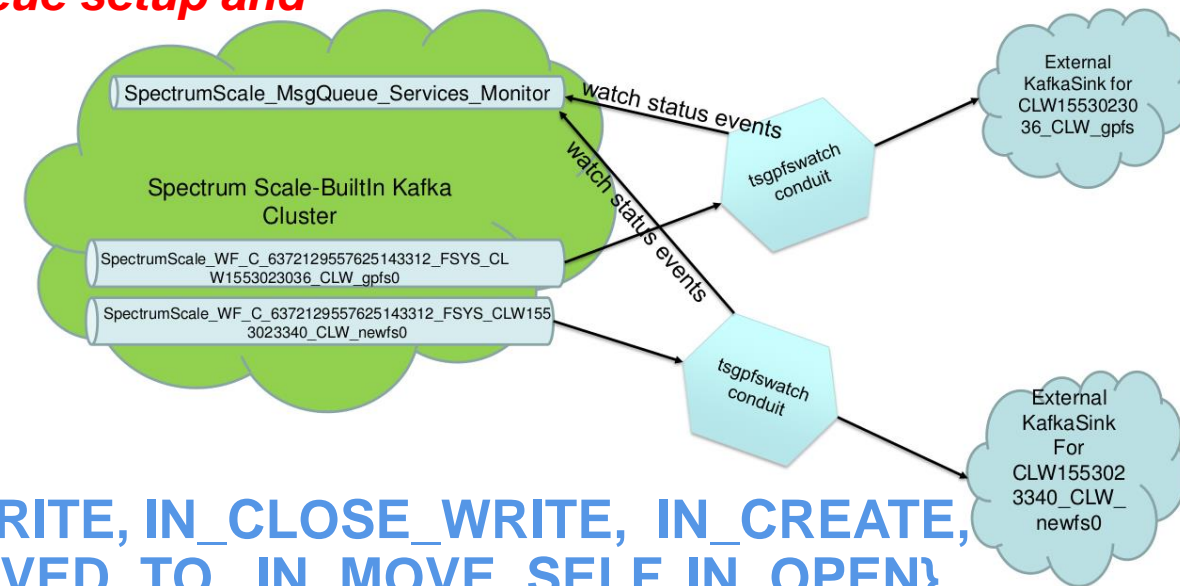
Watch file operations across cluster

mmwatch

Watch

entire file system, fileset or inode space

Events: {IN_ACCESS, IN_ATTRIB, IN_CLOSE_NOWRITE, IN_CLOSE_WRITE, IN_CREATE, IN_DELETE, IN_MODIFY, IN_MOVED_FROM, IN_MOVED_TO, IN_MOVE_SELF, IN_OPEN}




File Compression Algorithms

Genomics compression methods added in release 5.0.3 are:

Run via ILM policy or
mmchattr --compress ...

Still have and



Spectrum Scale TCT capabilities (5.0.3)

Direct Client transparent recall support

- Any and all clients can be configured to transfer the data for transparent recall requests

Azure storage support

zLinux Spectrum Scale Cluster support

Automatic container spillover

- When a tiering container reaches the 100 million threshold a new container is automatically created and used

Security: default SigV4 support

- Move from jclouds to the Amazon SDK for S3 and S3 compatible device support

Simplified Backup / Restore

- The SOBAR script has been improved for a streamlined backup and restore of TCT

Quota support script

- Quota support script is available that allows for periodic quota reports by User / Filesystem

More performance improvements

- Further improvements in performance for both migrates and recalls

ESS Updates



IBM Elastic Storage Server (ESS)



Integrated scale-out data management for file and object data

Optimal building block for high-performance, scalable, reliable enterprise Spectrum Scale 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 Spectrum Scale data needs

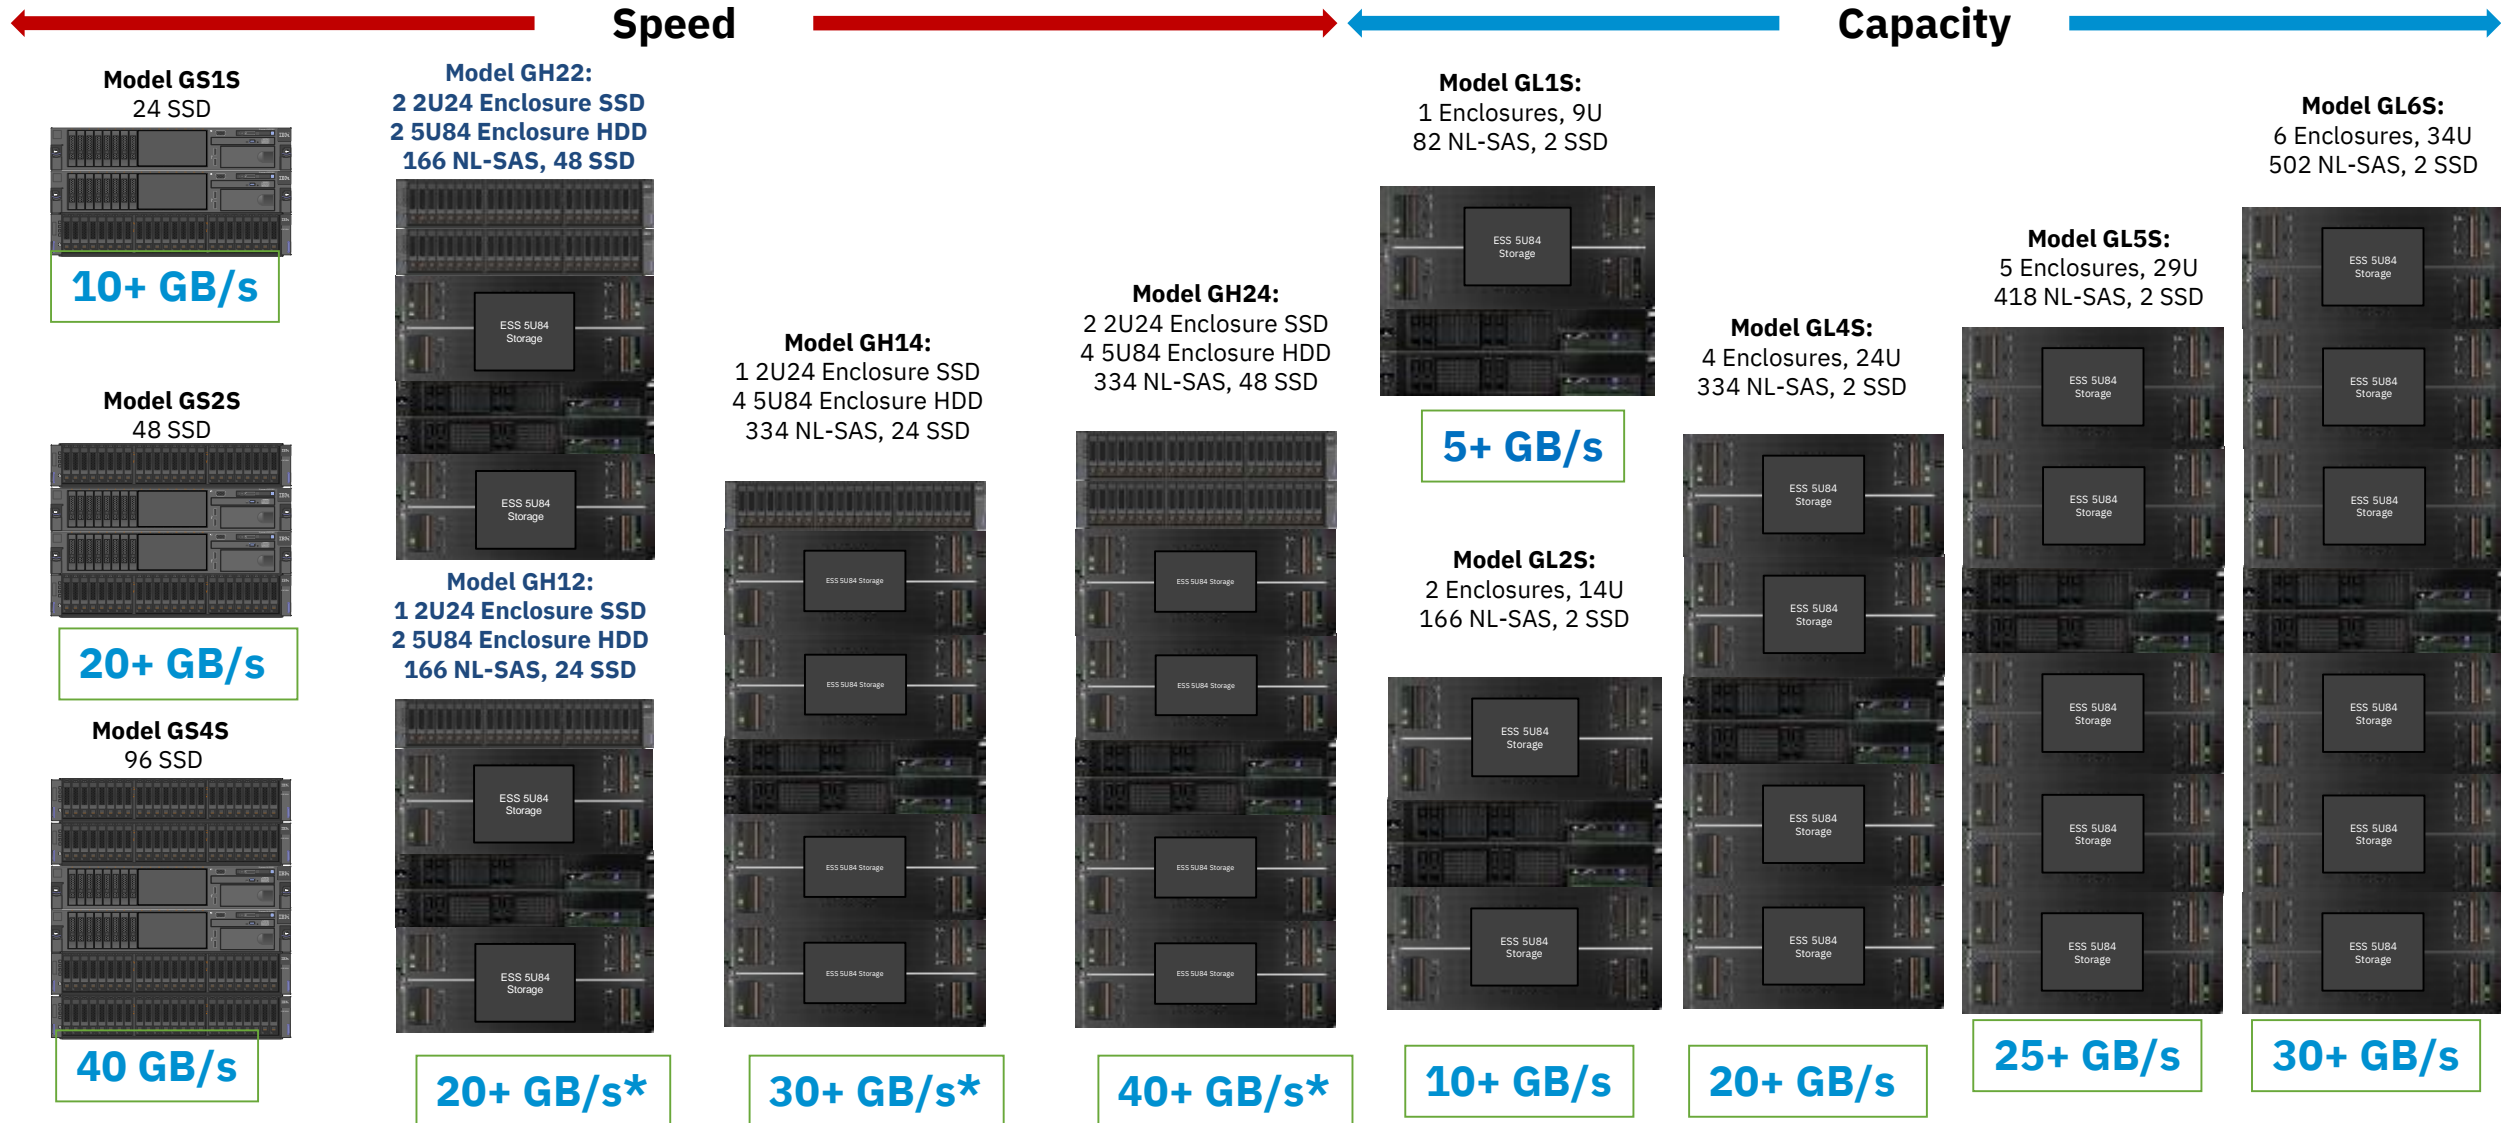
- Single repository of data with unified file and object support
- Anywhere access with multi-protocol support: NFS 4.0, SMB, OpenStack Swift, Cinder, and Manila
- Ideal for Big Data Analytics with full Hadoop transparency

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
- Five 99999s of availability



IBM Elastic Storage Server: building blocks small and large



IBM Elastic Storage Server GLxC models




Model GL1C:
1 Enclosure, 8U
104 NL-SAS, 2 SSD




1.46 PB raw

Model GL2C:
2 Enclosures, 12U
210 NL-SAS, 2 SSD



2.9 PB

Model GL4C
4 Enclosures, 16U
432 NL-SAS, 2 SSD




5.9 PB

Model GL5C
5 Enclosures, 28U
528 NL-SAS, 2 SSD




7.3 PB

Model GL6C
6 Enclosures, 28U
634 NL-SAS, 2 SSD



8.8 PB

Model GL8C
8 Enclosures, 36U
846 NL-SAS, 2 SSD



11.8 PB raw

Free! Introduction to IBM Elastic Storage Server and Spectrum Scale RAID and gssutils

(Log on with your IBM ID)

<https://www.onlinedigitallearning.com/course/view.php?id=2173>

<https://www.onlinedigitallearning.com/course/view.php?id=3570>

My Enrollments > My courses >

DL08003G

Navigation

My Enrollments

My courses

DL08003G

Support

Self Completion

You have already completed this course

Course Completion Status

Introducing the Elastic Storage Server

Basic | 30 Minutes | Self-paced

The IBM Elastic Storage Server is a big data storage system that combines Power servers, storage enclosures, and disks along with IBM Spectrum Scale and IBM Spectrum Scale RAID technology, providing analytic and technical computing storage and data services for elastic storage workloads.

Launch

gssutils simulation for SSR Tools s

My Enrollments / My courses / DL08010G: SSR gssutils

- This content is for installers or administrators for the IBM Elastic Storage Server.
- You need 15 minutes or less to go through this content.
- This simulation is for the **gssutils** tool that is pre-installed as part of the IBM Elastic Storage Server. It is used to verify the installation of an ESS.

```
ESS INSTALLATION AND DEPLOYMENT TOOLKIT

1. Help
2. SSR Tools
3. Plug n Play and Hybrid
4. Install
5. Upgrade
6. Validation checks
7. View/Collect service data (snaps)
8. Exit

man gssutils_panel_1
Help
```

```
CHECK SYSTEM HARDWARE AND SOFTWARE

1. Help
2. Show node details
3. Check and validate various install parameters
4. Quick storage configuration check
5. Check enclosure cabling and paths to disks
6. Check disks for IO operations
7. Ping all nodes
8. Check ssh to all nodes
9. Run lsscsi from all nodes
10. Check for open serviceable events
11. Back

/opt/ibm/gss/tools/bin/gssnodedetails -N ems1,gss_ppc64
Shows miscellaneous node information.
```

Spectrum Scale Licensing



new ESS choices, unlimited clients and servers (non-NSD)!

License Type	Options
Data Access Edition	Default with ESS
Data Management Edition	Adds Encryption, AFM-DR, Transparent Cloud Tiering, File Audit Logging
Storage Utility Offering	<ol style="list-style-type: none">1. https://www.ibm.com/us-en/marketplace/storage-utility-offering2. Deliver 3 year capacity on day 13. May use Spectrum Control Storage Insights to look at growth4. Pay for what you use , Calculated Monthly

All nodes in a single cluster must be on compatible licenses
All nodes on Standard Edition --OR-- Data Management Edition

Software Changes

Software Name	Previous Version 5.3.3	Current Version 5.3.4
Spectrum Scale	5.0.2.3 + efix 5.0.2.3.0.1	5.0.3.1
HMC (For classic only)	860 SP3	860 SP3
xCAT	2.14.5	2.14.6
System Firmware	FW860.60 (SV860_180)	FW860.60 (SV860_180)
Red Hat Enterprise Linux (PPC64BE and PPC64LE)	7.5	7.6
Kernel Systemd Network Manager	3.10.0-862.27.1 219-57.el7_5.5 1.12.0-8.el7_6	3.10.0-957.12.2 219-62.el7_6.6 1.12.0-10.el7_6
Open Fabrics Enterprise Distribution (Mellanox, Infiniband, some Ethernet)	MOFED 4.5-2.2.0.1	MOFED 4.6-1.0.1.2
IPR (for boot drives)	19511400	19512200
ES AGENT	4.5.0.0	4.5.1

5.3.4 – Online upgrade from ESS 5.1.X

ESS version	3.5.5 (or earlier)	4.0.x	4.5/4.6	5.0.x	5.1.x	5.2.0	5.3.0
3.5.5 (or earlier)	Yes	Yes	Yes	No	No	No	NO
4.0.x	N/A	Yes	Yes	Yes	No	No	NO
4.5/4.6	N/A	N/A	Yes	Yes	Yes	No	NO
5.0.x	N/A	N/A	N/A	Yes	Yes	Yes	NO
5.1.x	N/A	N/A	N/A	N/A	Yes	Yes	YES
5.2.0	N/A	N/A	N/A	N/A	N/A	Yes	YES
5.3.0	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Non-disruptive upgrades




Simple expansion of Storage Capacity

- Spectrum Scale will automatically rebalance data in the background
- System automatically puts the new capacity to use
- No need to Archive & Restore data
- No System disruption*

Model GS1S
With 24 SSDs

Install additional
drawer with 24 SSDs

Model GS2S
With 48 SSDs



Non Disruptive Upgrades	
From	To
GS1S	GS2S
GS2S	GS4S
GL1S	GL2S
GL2S	GL4S
GL4S	GL6S
GL1C	GL2C
GL2C	GL4C
GL4C	GL5C

Example

*Requires space available in the rack

New features

- “mmvdisk” is default – also used in ECE
- Security models on how to run with:
 - “sudo” wrappers
 - “selinux”
 - “admin mode central”
 - “firewall”

Thank You.
IBM Storage & SDI

A decorative graphic consisting of numerous blue diagonal bars of varying lengths and orientations, creating a sense of movement and depth in the bottom right corner of the slide.