ESS update

5.3 Technical Update

Christopher D. Maestas

IBM ESS 5.3 – Announcement Overview

Highlights

Spectrum Scale 5.0 in ESS

- New standards in performance – Leveraging the Highest Performance
  Spectrum Scale System ever and deployed at Coral
- Ideal for Big Data Analytics, demanding IT workloads

New entry GL1S Model

- Entry Disk model starting at 324TB of capacity

Enhanced Install & Upgrade

- Replacement of current install with a new streamlined Menu driven process
- Deliver faster installs & upgrades

Spectrum Scale Licensing
GLxS (“new 5147/5148 ESS”) buyers, two choices

Data Access Edition*, licensed per disk

- Spectrum Scale RAID license entitlement included
- Two price tiers, HDD and SDD
- Select in eConfig

*this used to be the standard edition name, but this edition is based on capacity, not sockets. Meaning you can have unlimited clients and extra non storage server licenses

Data Management Edition, licensed per disk

- Adds Encryption, AFM-ADR, Transparent Cloud Tiering, File Audit Logging
- Two price tiers, HDD and SDD
- Select in eConfig

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

ESS 5.3 – New Entry GL1S Model

- The entry starting capacity point for disk just got lower
- GL1S with a single 5U84 storage enclosure

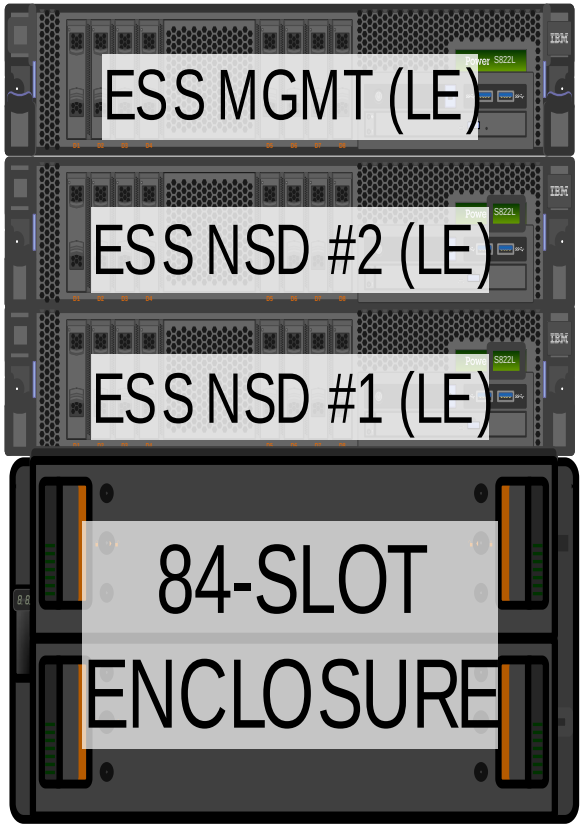

2nd Generation IBM Elastic Storage Server (ESS) Family

IBM Storage & SDI

Speed

Model GS1S
24 SSD

Model GS2S
48 SSD

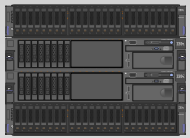

Model GS4S
96 SSD

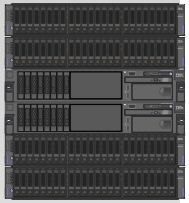

All Flash ESS
Announce: July 11, 2017
GA: Aug 25, 2017

Capacity

Model GL1S:
1 Enclosures, 9U
82 NL-SAS, 2 SSD

New!

Available Today!

Model GL2S:
2 Enclosures, 12U
166 NL-SAS, 2 SSD

Model GL4S:
4 Enclosures, 20U
334 NL-SAS, 2 SSD

Model GL6S:
6 Enclosures, 28U
502 NL-SAS, 2 SSD

Announced: April 2017
GA: June 2017

Software Changes

Software Name	Version
Spectrum Scale	5.0.0-1.1.2 (ESS 5301)
HMC (For classic only)	860 SP2
xCAT	2.13.19
System Firmware	SV860_138(FW860.42)
Red Hat Enterprise Linux	7.3 (PPC64BE and PPC64LE)
Kernel Systemd Network Manager	3.10.0-514.44.1 219-42.el7_4.10 1.8.0-11.el7_4
Open Fabrics Enterprise Distribution (Mellanox, Infiniband, some Ethernet)	MLNX_OFED_LINUX-4.1-4.1.6.1
IPR (for boot drives)	17518300
ESA	4.2.0-9

Upgrading paths to 5.3.0.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

The matrix of versions!

ESS version	Spectrum Scale	OS	Kernel errata	OFED	Firmware	IPR	Systemd	Netmgr
5.0.2	4.2.2-3 efix11	RHEL7.2	3.10.0-327.53.1.el7.ppc64	MLNX_OFED_LINUX-3.4-2.0.0.1	FW860.10 (SV860_056)	15511300	219-30.e7_3.8	N/A
5.1.1 (LE+BE)	4.2.3.2	RHEL7.2 (LE+BE)	3.10.0-327.55.3.el7.ppc64 + ppc64le	MLNX_OFED_LINUX-4.0-2.0.0.3	FW860.30 (SV860_103)	15511800	219-30.el7_3.8	N/A
5.2 (LE+BE)	4.2.3-4	RHEL7.3 (LE+BE)	3.10.0-514.26.2.el7.ppc64 + ppc64le	MLNX_OFED_LINUX-4.1-0.1.4.1	FW860.30 (SV860_103)	16519500	219-30.el7_3.9	1.4.0-20.el7_3
5.3 (LE+BE)	5.0.0-1 (GNR efix)	RH7.3 (LE+BE)	3.10.0-514.44.1.ppc64 + ppc64le	MLNX_OFED_LINUX-4.1-4.1.6.1	SV860_138 (FW860.42)	17518300	219-42.el7_4.10	1.8.0-11.el7_4

ESS Performance, a side note

New Sizing Tool online

Re-running performance projections in POK in the next month

Scale 5.0.0 based filesystem and software

POK Benchmark center - GL6S and GS4S

Appliance Worksheet

	ESS					Additional Attributes		Workload		IOR based Sequential Performance					
	Appliance		Drive	ESS	TB	File System	Network	Profile		GB/sec (Base 10)			GiB/sec (Base 2)		
Unit ID	Model	Quantity	Model	Quantity	Marketing	Blocksize (MiB)	Link Speed	Read %	Write %	Read	Write	Total	Read	Write	Total
1	GS1S	1	SSD	24	3.84	16	FDR/EDR I	100	0						
2	GS2S	1	SSD	48	3.84	16	FDR/EDR I	100	0						
3	GS4S	1	SSD	96	3.84	16	FDR/EDR I	100	0						
4	GL1S	1	NL-SAS	82	10	16	FDR/EDR I	100	0						
5	GL2S	1	NL-SAS	166	10	16	FDR/EDR I	100	0						
6	GL4S	1	NL-SAS	334	10	16	FDR/EDR I	100	0						
7	GL6S	1	NL-SAS	502	10	16	FDR/EDR I	100	0						

How do I measure and set things?

- Magic Utility dstat
  - (watch the cut and paste of this command!)
  - dstat –noupdate –time –top-cpu –top-mem – top-io –top-bio –gpfs –gpfs-ops

```
sh-4.2# dstat --noupdate --time --top-cpu --top-mem --top-io --top-bio --gpfs --gpfs-ops
----system---- -most-expensive- --most-expensive- ----most-expensive---- ----most-expensive---- --gpfs-i/o- --------gpfs-file-operations-------
     time     |  cpu process   |  memory process  |     i/o process      |  block i/o process   | read writ| open clos read writ rdir inod
12-04 22:47:06|mmsysmon.py  0.4|mmfsd       1187M|sshd        319k  140k|sshd       1964B   32k|   0    0 |   0    0    0    0    0    0
12-04 22:47:07|migration/1  0.5|mmfsd       1187M|mmksh      2896k   14k|mmksh          0   288k|   0    0 |   0    0    0    0    0    0
12-04 22:47:08|mmsysmon.py  1.0|mmfsd       1187M|mmsysmon.py 184k 4720B|                     |   0    0 |   0    0    0    0    0    0
12-04 22:47:09|java         1.0|mmfsd       1187M|java        101k   31B|postgres: p   0    16k|   0    0 |   0    0    0    0    0    0
12-04 22:47:10|pmsensors    0.5|mmfsd       1187M|pmsensors   6170B   97B|                     |   0    0 |   0    0    0    0    0    0
```

Deployments

ESS 5.3 – Enhanced Install & Upgrade

IBM ESS clearly delivers extreme Performance and Scalability.
With this tremendous performance we recognize there is added complexity for some customers.

Starting with ESS 5.3 the Install and Upgrade process has been dramatically improved

- System precheck has been improved to validate the system is ready for install
- Command line actions has been replaced by a Menu driven system
- The sequence of activities is automated behind the menu options selected
- IBM Lab Services have enhanced access to the latest RHEL Errata

This all results in faster “Time to Value” and improved customer experience.

ESS Deployment methods

Plug-N-Play mode

Unpacking and basic power connectivity completed

FSP and xCAT networks in documented ports and connected to proper vlans

SSRs have validated using **gssutils** for correct disk placement, cabling, networking, server health

Access to the EMS over ssh

Setup building block using Fusion mode with **gssutils**

Follow the manual steps but execute within **gssutils**

Fusion mode ends at network bond creation. Execute the rest of the quick deployment guide using **gssutils**

Create network bonds

Create cluster, vdisks, nsds, filesystem

Final checks

Setup the GUI, call home, connect systems to RHN

What is this gssutils that you speak of?

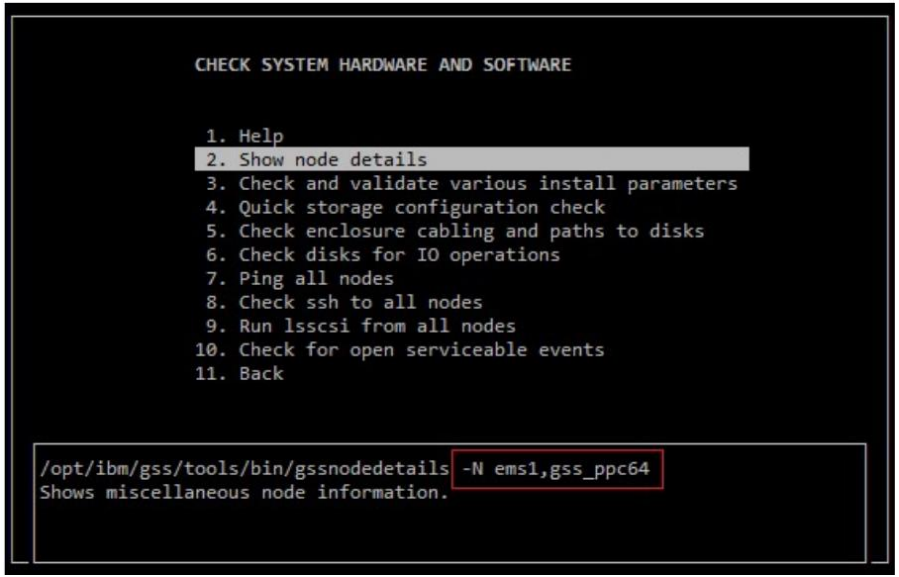

ESS Manufacturing rack configuration testing

Part of Quality Control Initiative

Should occur sometime in May

Sample order of an ESS

Run through deployment steps

Validate documentation and procedures for

- SSR
- Lab Services

IBM Storage & SDI

ESS Implementation Services and Support

Support – please look here!

ESS FAQ

https://www.ibm.com/support/knowledgecenter/SSYSP8/gnrfaq.pdf?view=kc

Scale Knowledge center

https://www.ibm.com/support/knowledgecenter/en/STXKQY_5.0.0/ibmspectrumscale500_welcome.html

ESS Redbook

http://www.redbooks.ibm.com/redpapers/pdfs/redp5253.pdf

Scale Forum

https://www.ibm.com/developerworks/community/forums/html/forum?id=11111111-0000-0000-0000-000000000479

Additional Help, Info, and Training

- "JD": J D Zeeman jdzeeman@us.ibm.com
  Global Sales Leader for Elastic Storage Server
- John Sing jmsing@us.ibm.com
  Offering Evangelist, Spectrum Scale and ESS
- Christopher Maestas: cdmaestas@us.ibm.com
  Global Architect, SDS and SDI, Spectrum Scale, ESS and Cloud Object Storage
- Indulis Bernsteins: INDULISB@uk.ibm.com
  Global Architect, SDS and SDI, Spectrum Scale, ESS and Cloud Object Storage
- Ashutosh Mate: mate@us.ibm.com
  Global Architect, SDS and SDI, Spectrum Scale, ESS and Cloud Object Storage
- Par Hettinga-Ayakannu par@nl.ibm.com
  Worldwide SDS and SDI Enablement
- Alex Q Chen aqchen@us.ibm.com
  Offering Executive, File and Object Storage, ESS
- Doug Petteway cdpettew@us.ibm.com
  Offering Manager IBM Storage, ESS
- Matt: Matthew Drahzal mdrahzal@us.ibm.com
  Offering Manager IBM Power, ESS
- Len: Leonard Accardi laccardi@us.ibm.com
  Global Sales Leader for Enterprise Storage
- Steve: Stephen Edel edel@us.ibm.com
  NA Technical Sales for Spectrum Scale SW and ESS
- David: David Cremese DCR@ch.ibm.com
  European Sales Leader for Spectrum Scale SW and ESS
- Eyal Abraham eyal.abraham@us.ibm.com
  Global Storage Solutions Sales

IBM SPECTRUM SCALE

Support update, common issue and best practice

Guanglei LI
liguangl@cn.ibm.com
March 2018

Follow the sun support – Aligning support staff to customer time zone

- Spectrum Scale Support is growing to better meet customer needs.
- Beginning late 2016 we substantially grew the support team in Beijing, China, with experienced Spectrum Scale staff.
- Improved response time on severity 1 production outages; reducing customer waiting time before L2 is engaged as well as time to resolution.
- Positive impact to timely client L2 communication for severity 2, 3, and 4 PMRs within our customer time zone.
- Setup and grew EMEA support team in Germany in late 2017
- 3 major sites: North America, China, Germany
- PagerDuty was introduced this year for better PMR monitor

IBM Spectrum Scale Level 2 Support Global Time Zone Coverage

Global team locations

- North America
  - ✓ *Poughkeepsie, NY USA
  - ✓ Toronto, ON Canada
- AP
  - ✓ *Beijing, China
  - ✓ India
- Europe
  - ✓ *Germany

* Major sites

Support Delivery: Managers

1st Level: Bob Simon: ragonese@us.ibm.com; 1-845-433-7285

1st Level: Jun Hui Bu: bujunhui@cn.ibm.com; 86-10-8245-4113

1st Level: Dennis Kunkel: Dennis.Kunkel@de.ibm.com; 49-170-3387365

WW 2nd Level: Wenwei Liu: wliu@ca.ibm.com; 1-905-316-2623

Support Executive

Andrew Giblon: agiblon@ca.ibm.com; 1-905-316-2582

Thank You.

IBM Storage & SDI

Backup

COMMON FIELD ISSUE AND BEST PRACTICES

DATA COLLECTION: GPFS.SNAP

1) Use the "--limit-large-files" flag to limit the amount of 'large files' collected. The 'large files' are defined to be the internal dumps, traces, and log dump files that are known to be some of the biggest consumers of space in gpfs.snap (these are files typically found in /tmp/mmfs of the form internaldump.*.*, trcrpt.*.*, logdump*.*.*). Added in version 4.1.1

--limit-large-files: YYYY:MM:DD:HH:MM | Num_Days_back | 0

2) Limit the nodes on which data is collected using the '-N' flag to gpfs.snap. By default data will be collected on all nodes, with additional master data (cluster aware commands) being collected from the initiating node.

- Note: Please avoid using the –z flag on gpfs.snap unless supplementing an existing master snap or you are unable to run a master snap.

3) To clean up old data over time, it's recommended that gpfs.snap be run occasionally with the '--purge-files' flag to clean up 'large debug files' that are over the specified number of days old. added in version 4.2.0

--purge-files: KeepNumberOfDaysBack | 0

FIRST TIME DATA COLLECTION FOR PERF/HANG

1. Gather waiters and create working collective. It can be good to get multiple looks at what the waiters are and how they have changed, so doing the first mmlsnode command (with the -L) numerous times as you proceed through the steps below might be helpful (specially if issue is pure performance, no hangs).

   ```
   mmlsnode -N waiters > /tmp/waiters.wcoll
   mmdsh -N /tmp/waiters.wcoll "mkdir /tmp/mmfs 2>/dev/null“
   mmlsnode -N waiters -L | sort -nk 4,4 > /tmp/mmfs/service.allwaiters.$(date +"%m%d%H%M%S")
   ```

2. View allwaiters and waiters.wcoll files to verify that these files are not empty.If either (or both) file(s) are empty, this indicates that the issues seen are not GPFS waiting on any of it's threads. Data to be gathered in this case will vary. Do not continue with steps. Tell Service person and they will determine the best course of action and what docs will be needed.

3. Gather internaldump from all nodes in the working collective

   ```
   mmdsh -N /tmp/waiters.wcoll "/usr/lpp/mmfs/bin/mmfsadm dump all > /tmp/mmfs/service.\$(hostname -s).dumpall.\$(date +"%m%d%H%M%S")"
   ```


FIRST TIME DATA COLLECTION FOR PERF/HANG CONT.

4. Gather kthreads from all nodes in the working collective

   mmdsh -N /tmp/waiters.wcoll "/usr/lpp/mmfs/bin/mmfsadm dump kthreads > /tmp/mmfs/service.\$(hostname -s).kthreads.\$(date +"%m%d%H%M%S")“

   *note:

   If running Linux OS on SpectrumScale (formerly GPFS) 4.1 or higher - this step could be skipped.

5. If this is a performance problem, get 60 seconds mmfs trace from the nodes in the working collective.

   If AIX …

   mmtracectl --start --aix-trace-buffer-size=256M --trace-file-size=512M -N /tmp/waiters.wcoll ; sleep 60; mmtracectl --stop -N /tmp/waiters.wcoll

   If Linux ..

   mmtracectl --start --trace-file-size=512M -N /tmp/waiters.wcoll ; sleep 60; mmtracectl --stop -N /tmp/waiters.wcoll

6. Run gpfs.snap to collect all the data generated

   gpfs.snap -N /tmp/waiters.wcoll

PERFORMANCE TUNING

1) pagepool - cache user file data and file system metadata
   Needs to understand the IO pattern on client nodes when tuning pagepool: Sequential IO, Random IO, Direct IO

2) maxFilesToCache - controls how many file descriptors each node can cache.
   - Needs large value if there will be many files opened concurrently, e.g., 1M for NFS & Samba service. Large value can improve the performance of user interactive operations like running "ls"
   - Small value with many files being accessed will cause high CPU usage
   - Increasing maxFilesToCache in a large cluster with hundreds of nodes increases the number of tokens a token manager needs to store. Ensure that the manager node has enough memory and tokenMemLimit is increased when running GPFS version 4.1.1 and earlier.

3) workerThreads - controls an integrated group of variables that tune the file system performance
   - New in GPFS 4.2.0.3 to simplify tuning. Some variables are auto-calculated when WorkerThreads is enabled. e.g, worker1Threads, worker3Threads
   - You can manually adjust external variables to avoid auto-tuned by workerThreads when Spectrum Scale computed from WorkerThreads are not suitable for your workload
   - Default 48. Increaset o 512 or 1024 if there will be many threads access GPFS file system on that node. E.g., running NFS and Samba service on that node

PERFORMANCE TUNING CONT.

1. defaultHelperNodes – Specify the nodes to be used for distributed commands
   - Command list: mmadddisk, mmapplypolicy,mmbackup, mmchdisk, mmcheckquota, mmdefragfs, mmdeldisk, mmdelsnapshot, mmfileid,mmfsck, mmimgbackup, mmimgrestore, mmrestorefs, mmrestripefs,mmrpldisk
   - Example: runningmmrestripefs on limited nodes including NSD servers
2. maxMBps - indicates the maximum throughput in megabytes per second that GPFS can submit into or out of a single node
   - It's a hint GPFS uses to calculate how many prefetch/writebehind threads should be scheduled
   - Set client nodes maxMBpS based on IO throughput. 2x of total IO throughput divided by # of client nodes

FS CORRUPTION

1) MMFS_FSSTUCT error
   - It will be printed into system log if GPFS detect FS corruption when access the file system.
   - fsstructlx.awk(Linux) fsstruct.awk(AIX) under /lpp/mmfs/samples/debugtools/ to decode the MMFS_FSSTRUCT message in system log:
     fsstructlx.awk /var/log/messages > fsstruct.message
   - mmhealth will report FS corruptions
2) Offline mmfsck to check file system and generate report
   - GPFS file system needs to be unmounted from all nodes.
   - Use patch file option (from ver 4.1.1) to avoid two rounds of long running mmfsck:

     mmfsck -nV --patch-file /tmp/fsck.patch
   - Online mmfsck
     - run mmfsck with –o option while FS is mounted
     - Can only fix the lost blocks – data block marked as used but not referenced by any file/dir

FS CORRUPTION CONT.

1) Upload mmfsck output and patch file for IBM to review. Additional output may be required:
   - tsfindinode to identify the pathname for corrupted inodes. Needs to mount FS
   - tsdbfs output for inode dumps
2) Run offline mmfsck fix under guidance of IBM support
   - If patch is used, run it with:

     mmfsck <fs> -V --patch-file /tmp/mmfs/fsck.patch –patch
3) Log recovery failure
   - mmfsck <fs > -xk
     - Needs to unmount FS
     - Supported in ver >=4.2
     - Run it after confirmed with IBM support.

BEST PRACTICE: NSD MISSING

1) Disk Missing
    1) Use "mmlsnsd –X" to check if any disk reported as "(not found)"
    2) Use "tspreparedisk –s" on each node to check if a NSD could be identified.
    3) mmnsddiscover –a –N all
    4) User exit of /var/mmfs/etc/nsddevices could affect NSD discovery
    5) Disk type mismatch: mmchconfig updateNsdType=<nsd_type_file>
2) Disk Header Missing
    1) There are 3 parts in NSD header: NSD desc, Disk desc, FS desc.
    2) "mmfsadm test readdescraw /dev/dev_name" could be used to show headers.
    3) Use tspreparedisk & dd command to restore NSD header. Do this under guidance of IBM support, and not able to restore in some cases.
    4) A common cause for header missing: disk header erased by UEFI driver update link

BEST PRACTICE: EXPEL

1) Network
   - GPFS will send out pings before expel a node:
     … is being expelled because of an expired lease. Pings sent: 60. Replies received: 0
   - Common causes
     - Mis-matched MTU size: Jumbo Frames enabled on some or all nodes but not on the network switch.
     - Old adapter firmware levels and/or incorrect OFED software are utilized
     - OS specific (TCP/IP, Memory) tuning has not been re-applied.
     - verbsRdmaSend is enabled for SS ver < 5.0. It has scaling issue in GPFS 3.x and 4.x link1 link2
     - Node A can't talk with Node B. Node A will ask Cluster Manager to expel Node B. Node A or Node B will be expelled.
2) Node load
   - GPFS cluster manager is too busy to handle incoming lease request. Avoid overloading cluster manager on large scale cluster
   - GPFS >= 4.2.3 support Prioritization of critical RPCs including lease request
   - Increase failure detection time for node expel:
     mmchconfig minMissedPingTimeout=120 (default is 3)
     mmchconfig maxMissedPingTimeout=120 (default is 60)
     mmchconfig leaseRecoveryWait=120 (default is 35)

BEST PRACTICE: EXPEL CONT.

1) Expel auto data collection from 4.1.1
   - **When a node is about to be expelled for unknown reasons, debug data is collected automatically to help find the root cause**
   - **Controled by config parameter: expelDataCollectionDailyLimit, expelDataCollectionMinInterval**
   - **Expel debug data will be collected on cluster manager and involved nodes.**

2) Auto data collection for unhealthy TCP connections from 4.2.3.
   - GPFS log(var/adm/ras/mmfs.log.laest):
     The TCP connection to IP address 192.168.38.52 c38f2bc1n02 <c0n4> (socket 45) state is unexpected: ca_state=0 unacked=46 rto=25856000
   - Controlled by expel Data collection parameters.

SPECTRUM SCALE ANNOUNCE FORUMS

Monitor the Announce forums for news on the latest problems fixed, technotes, security bulletins and Flash advisories.

https://www.ibm.com/developerworks/community/forums/html/forum?id=11111111-0000-0000-0000-000000001606&ps=25

Subscribe to IBM notifications (for PTF availability, Flashes/Alerts): https://www-947.ibm.com/systems/support/myview/subscription/css.wss/subscriptions

ADDITIONAL RESOURCES

Tuning parameters change history:

https://www.ibm.com/support/knowledgecenter/STXKQY_4.2.2/com.ibm.spectrum.scale.v4r22.doc/bl1adm_changehistory.htm?cp=STXKQY

ESS best practices:

https://www.ibm.com/support/knowledgecenter/en/SSYSP8_3.5.0/com.ibm.spectrum.scale.raid.v4r11.adm.doc/bl1adv_planning.htm

Tuning Parameters:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20(GPFS)/page/Tuning%20Parameters

Share Nothing Environment Tuning Parameters:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20%28GPFS%29/page/IBM%20Spectrum%20Scale%20Tuning%20Recommendations%20for%20Shared%20Nothing%20Environments

Further Linux System Tuning:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Welcome%20to%20High%20Performance%20Computing%20(HPC)%20Central/page/Linux%20System%20Tuning%20Recommendations