Monitoring & Serviceability Enhancements in IBM Spectrum Scale

- User Group London 2022

Mathias Dietz (IBM) - Spectrum Scale RAS Architect
mdietz@de.ibm.com

Agenda

Monitoring Infrastructure Overview

Health Monitoring Improvements

- Stretch Cluster monitoring
- More monitoring improvements
- Outlook (5.1.4)

Call Home & Proactive Services

- Unified Callhome (ESS3500)
- Proactive Alerts / Healthchecker with Backchannel (>= 5.1.2)

Spectrum Scale – Monitoring Overview

Spectrum Scale

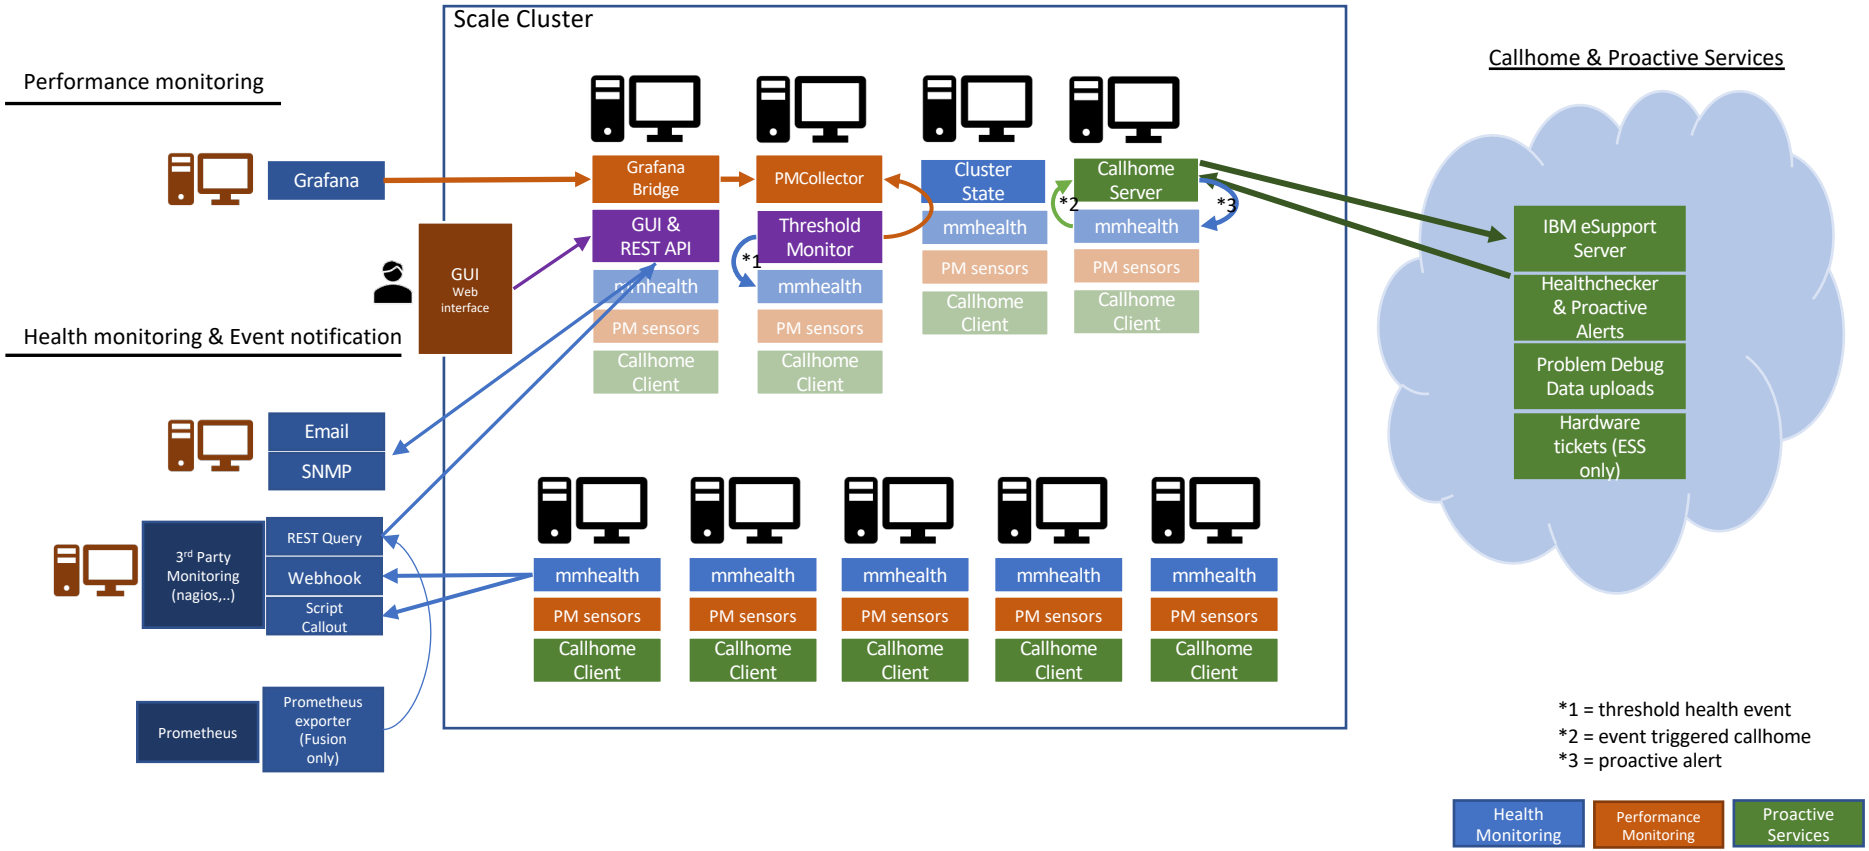

*1 = threshold health event
*2 = event triggered callhome
*3 = proactive alert

Health Monitoring | Performance Monitoring | Proactive Services

Spectrum Scale Stretch Cluster Example

Spectrum Scale

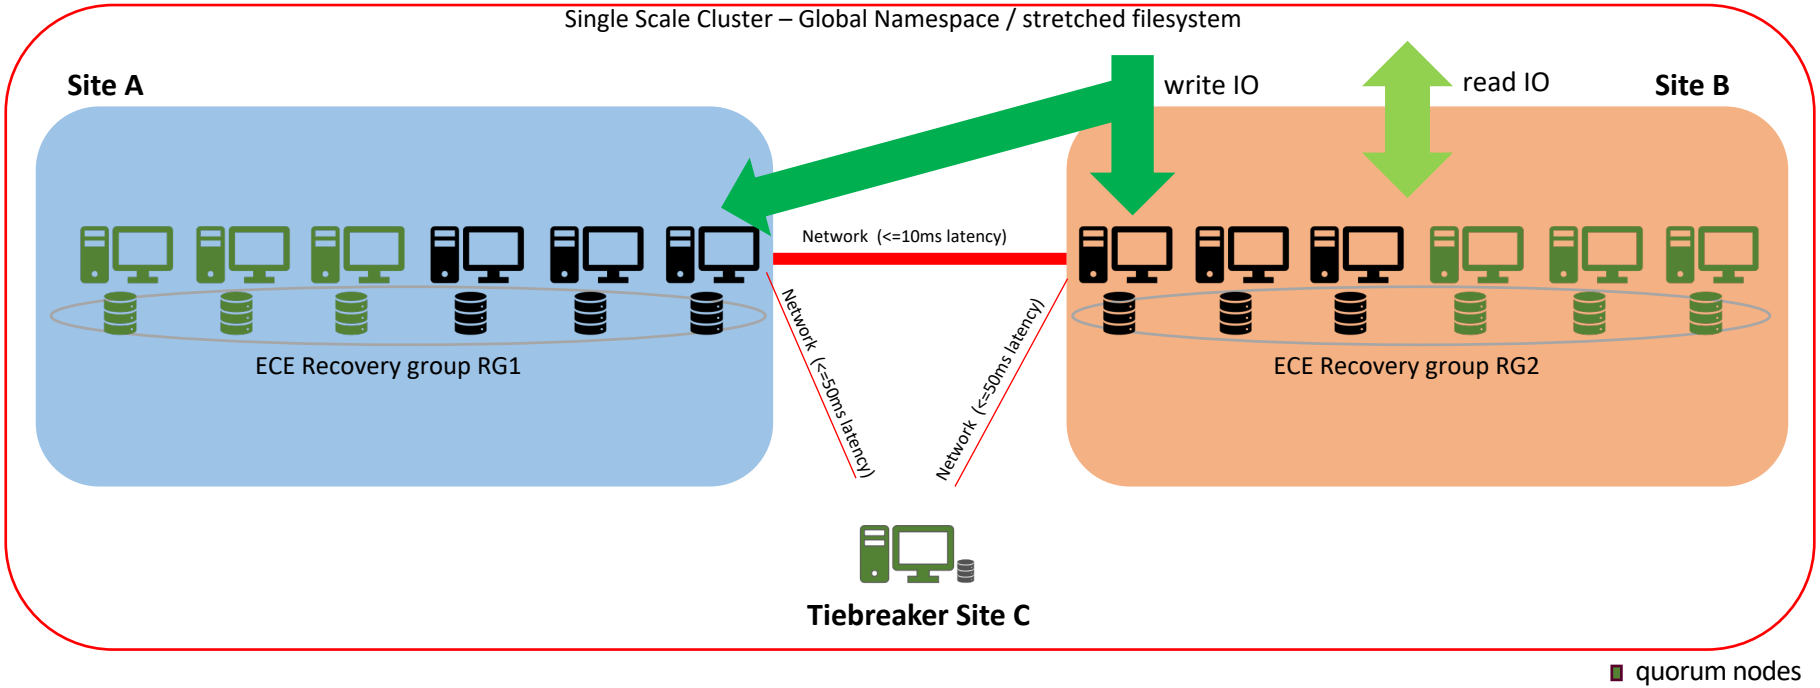

- Typical stretch cluster setup with two-way sync replication

Spectrum Scale Stretch Cluster Failures

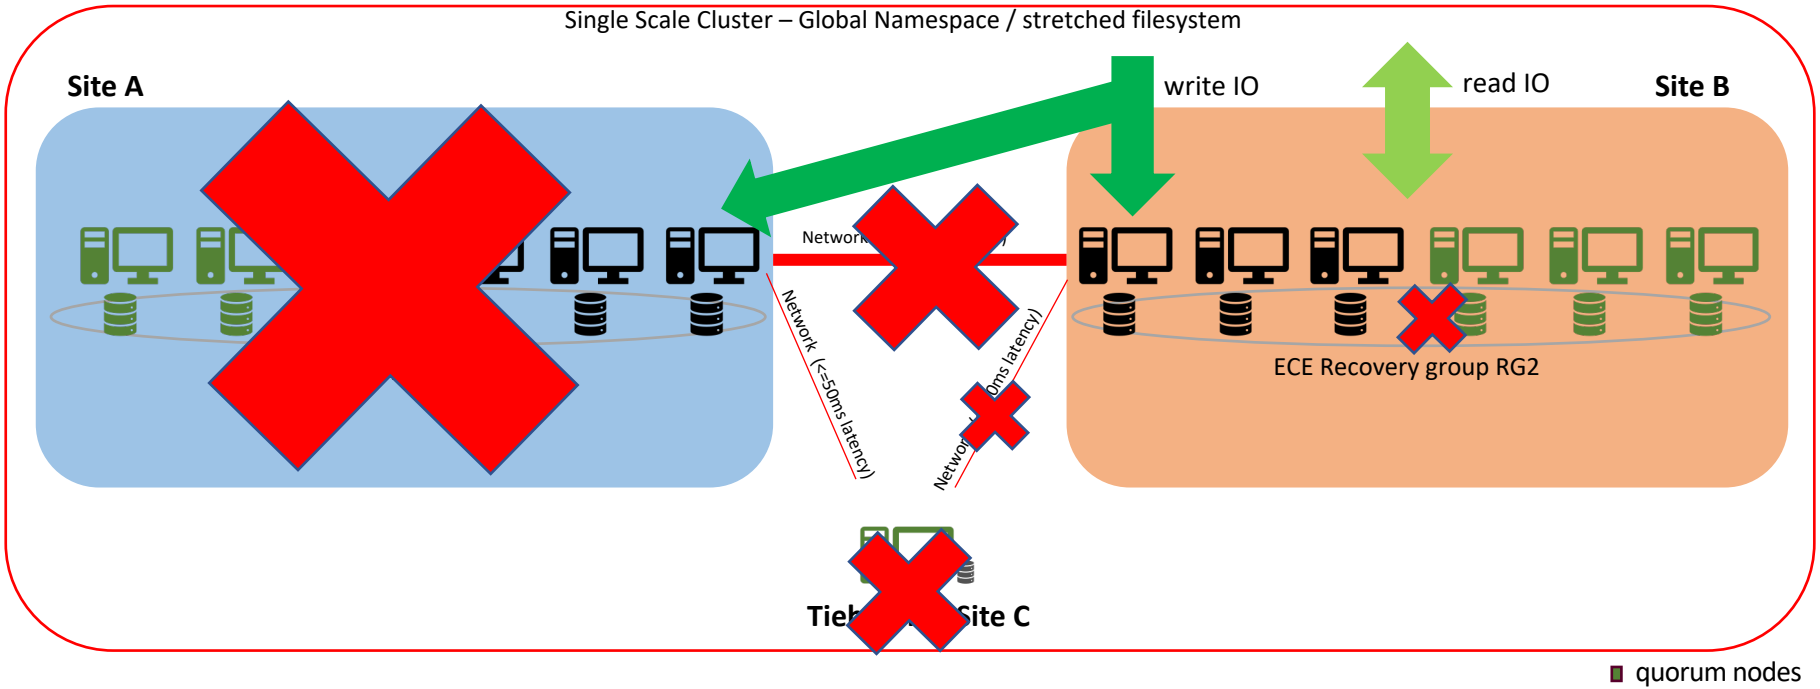

- Multiple different failure scenarios can happen
- Not always easy to understand the site status in case of failures

Stretch Cluster Monitoring

Spectrum Scale

- Added the ability to **make health monitoring site aware to report a derived site status**. Mmhealth/GUI will show configured sites under “StretchCluster” component.

- Allow “Sites” to be defined based on node class definitions
  - Node classes prefixed with **“SCALE_SITE_”** will be identified as a site.
  - Value following prefix will be the site name, and nodes belonging to the definition will collectively be treated as a site.

- Nodes at each site are monitored for a subset of health events e.g. heartbeats, replication issues, quorum failures, file system mount issues, etc..

- Events at each site are aggregated to determine if the site is HEALTHY, DEGRADED or FAILED.
  e.g. a single node having heartbeat issues doesn’t report a degraded a site,
  but if ½ of the nodes report this site is marked DEGRADED.

Stretch Cluster Monitoring Example Output

Spectrum Scale

```
[root@gimli-master-1 mmsysmon]# mmhealth cluster show stretchcluster -v
Component          Node                            Status        Reasons
-------------------------------------------------------------------------------------------
STRETCHCLUSTER     gimli-master-1.fyre.ibm.com     DEGRADED      site_fs_warn
  Frankfurt                                        DEGRADED      site_fs_warn
  Ehningen                                         HEALTHY       -
  TIEBREAKER                                       HEALTHY       -


[root@gimli-master-1 mmsysmon]# mmhealth event show site_fs_warn
Event Name:              site_fs_warn
Event ID:                998418
Description:             Many nodes face file system events at the site, which indicate network, resource, or
configuration issues.
Cause:                   Many nodes face file system events at the site.
User Action:             Check the health of the file system at the site and sure that it is properly mounted on all
nodes.
Severity:                WARNING
State:                   DEGRADED
```

More Improvements (5.1.3)

Spectrum Scale

Supress individual events during maintenance - Temporary hide of events

```
[root@gimli-11 ~]# mmhealth event hide disk_down disk3 --temporary 2m
Successfully temporarily hid event disk_down for entity disk3 for 2 mins.
```

Get informed early about capacity & Inode shortage - Events for lowDiskSpace and noDiskSpace

```
Example events:
no_disk_space_warn      WARNING   File system testFs runs out of space. StoragePool=system,…
no_disk_space_inode     WARNING   Fileset root runs out of space. Filesystem=testFs,…
low_disk_space_info     INFO      Low disk space. StoragePool system in file system testFs has reached the threshold as configured in
a migration policy.
```

Get informed about replication issues - Events for ill-replicated or ill-balanced filesystems

```
Example events:
ill_replicated_fs       WARNING   The filesystem gpfs0 is not properly replicated
ill_unbalanced_fs       TIP       The filesystem gpfs0 is not properly balanced
Ill_exposed_fs          WARNING   The filesystem gpfs0 has a data exposure risk as there are files where all replica are on suspended
```

Simplify replace of Call Home node

New option to change/update call home server node

```
[root@gimli-11 ~]# mmcallhome group change autoGroup_1 --server myclusternode2
Call home group autoGroup_1 has been changed
```

Automatic call home cleanup when cluster node is removed

Outlook (5.1.4 and beyond)

Spectrum Scale

Get fine granular disk error events - Enhanced Disk monitoring.

```
disk_down_io              ERROR     Disk nsd1 is reported as down because of an I/O issue. FS=gpfs0,..
disk_down_unexpected      ERROR     Disk nsd1 is reported as unexpected down. FS=gpfs0
disk_down_change/del/rpl ERROR     Disk nsd1 is reported as down because the configuration changed. FS=gpfs0
```

Know when your AFM gateway nodes run out of queue memory – AFM queue memory monitoring

```
New default threshold rule:
Warning level  = 80% of afmHardMemThreshold (default 5G)
Error level    = 90% of afmHardMemThreshold (default 5G)
```

Get informed when you filesystem is at risk because of missing desc disks - Stretch Cluster descriptor disk monitoring

```
site_fs_desc_fail         ERROR      Site FRANKFURT has no descriptor disks for all of the defined file systems
site_fs_desc_warn         WARNING    Site HAMBURG file system gpfs0 has no descriptor disks in failure groups 1
```

Avoid running with risk of data loss due to write cache – ESS Write Cache Enabled (WCE) check

```
gnr_pdisk_wcache_enabled WARNING    GNR pdisk n001p003 has write cache enabled
```

Get informed about SCSI passthrough IO hangs - passthrough hang detection

```
passthrough_query_hang          ERROR     A SCSI pass-through query request hang has been detected on disk n001p004 affecting file system gpfs0. Reason: …
```

ESS Unified Call Home

Spectrum Scale

ESS used to have 2 separate call home stacks:

- Hardware Call Home (ESA) -> automatic open tickets for hardware failures
- Software Call Home

Disadvantages: more complex installation and maintenance, inconsistent configurations, low usability

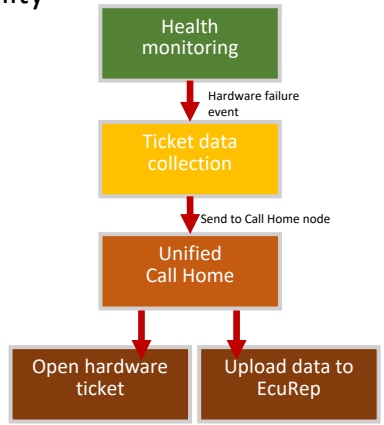

ESS 3500 comes with a single “Call Home” stack

Simplified configuration, no inconsistencies , higher usability

- Software Call Home now has the capability to open hardware tickets in Salesforce
- Single management command (mmcallhome)
- High level flow:
  - Hardware failures are detected by mmhealth and failure events will be raised
  - Ticket data collection is started (collecting event specific debug data)
  - Hardware ticket is created and related ticket data is uploaded to EcuRep
- ESA will still be used for older ESS systems (e.g. ESS 5000, ESS 3000, etc.)

Proactive Alerts / Healthchecker

Spectrum Scale

Spectrum Scale Health Checker is an IBM internal Cloud application which is used to **detect best practice violations, common misconfigurations and known problems** by running rules against incoming call home data and sending proactive alerts back to your system (>=5.1.2)

Get informed proactively if your system…

- is impacted by a new issue
  - e.g. when we learn about a new problem from customer A, we can identify which other customers are impacted by the same problem and send them a proactive alert
- is impacted by a Spectrum Scale Flash
- is not following “best practices” and could be tuned to improve performance, etc. (Tip)
- is misconfigured which might lead to problems in the future

- Fully integrated with the Spectrum Scale monitoring infrastructure (GUI, mmhealth, event notification, webhook,..)

Benefit from continuous enhancements and new proactive alerts without having to update to a new Spectrum Scale release !

Proactive Alert – mmhealth GUI & CLI

Spectrum Scale

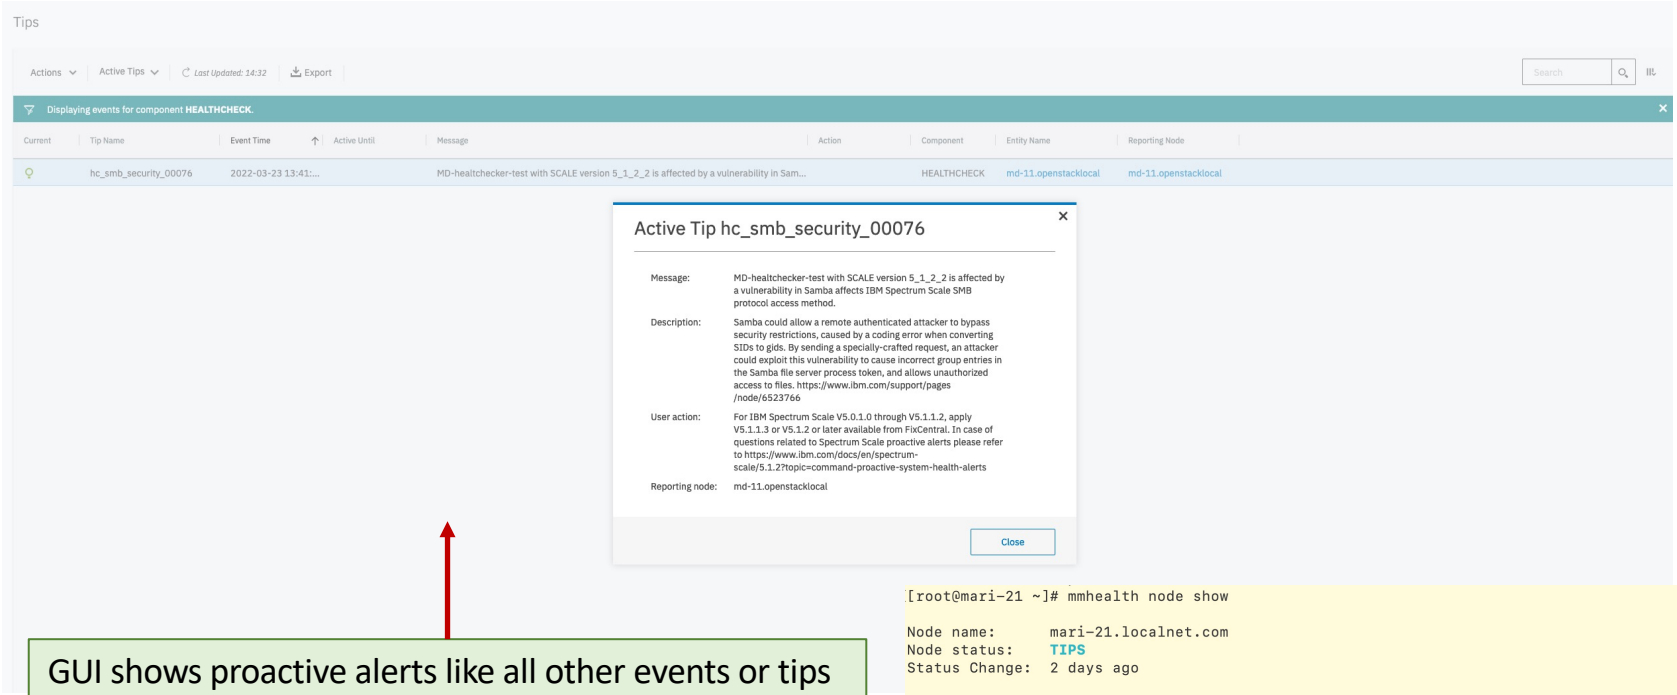

GUI shows proactive alerts like all other events or tips

New **HEALTHCHECK** mmhealth component on Call Home server node

```
[root@mari-21 ~]# mmhealth node show

Node name:      mari-21.localnet.com
Node status:    TIPS
Status Change:  2 days ago

Component       Status        Status Change     Reasons & Notices
------------------------------------------------------------------------------------------
GPFS            TIPS          2 days ago        gpfs_pagepool_small, gpfs_maxstatcache_low
NETWORK         HEALTHY       2 days ago        -
FILESYSTEM      HEALTHY       2 days ago        -
DISK            HEALTHY       2 days ago        -
CALLHOME        HEALTHY       2 days ago        -
CESIP           HEALTHY       2 days ago        -
FILESYSMGR      HEALTHY       2 days ago        -
GUI             HEALTHY       2 days ago        -
HEALTHCHECK     FAILED        1 day ago         hc_alert_scale_00067_notok
PERFMON         HEALTHY       2 days ago        -
THRESHOLD       HEALTHY       2 days ago        -
[root@mari-21 ~]#
```

Healthchecker / Proactive Alerts Flow

Spectrum Scale

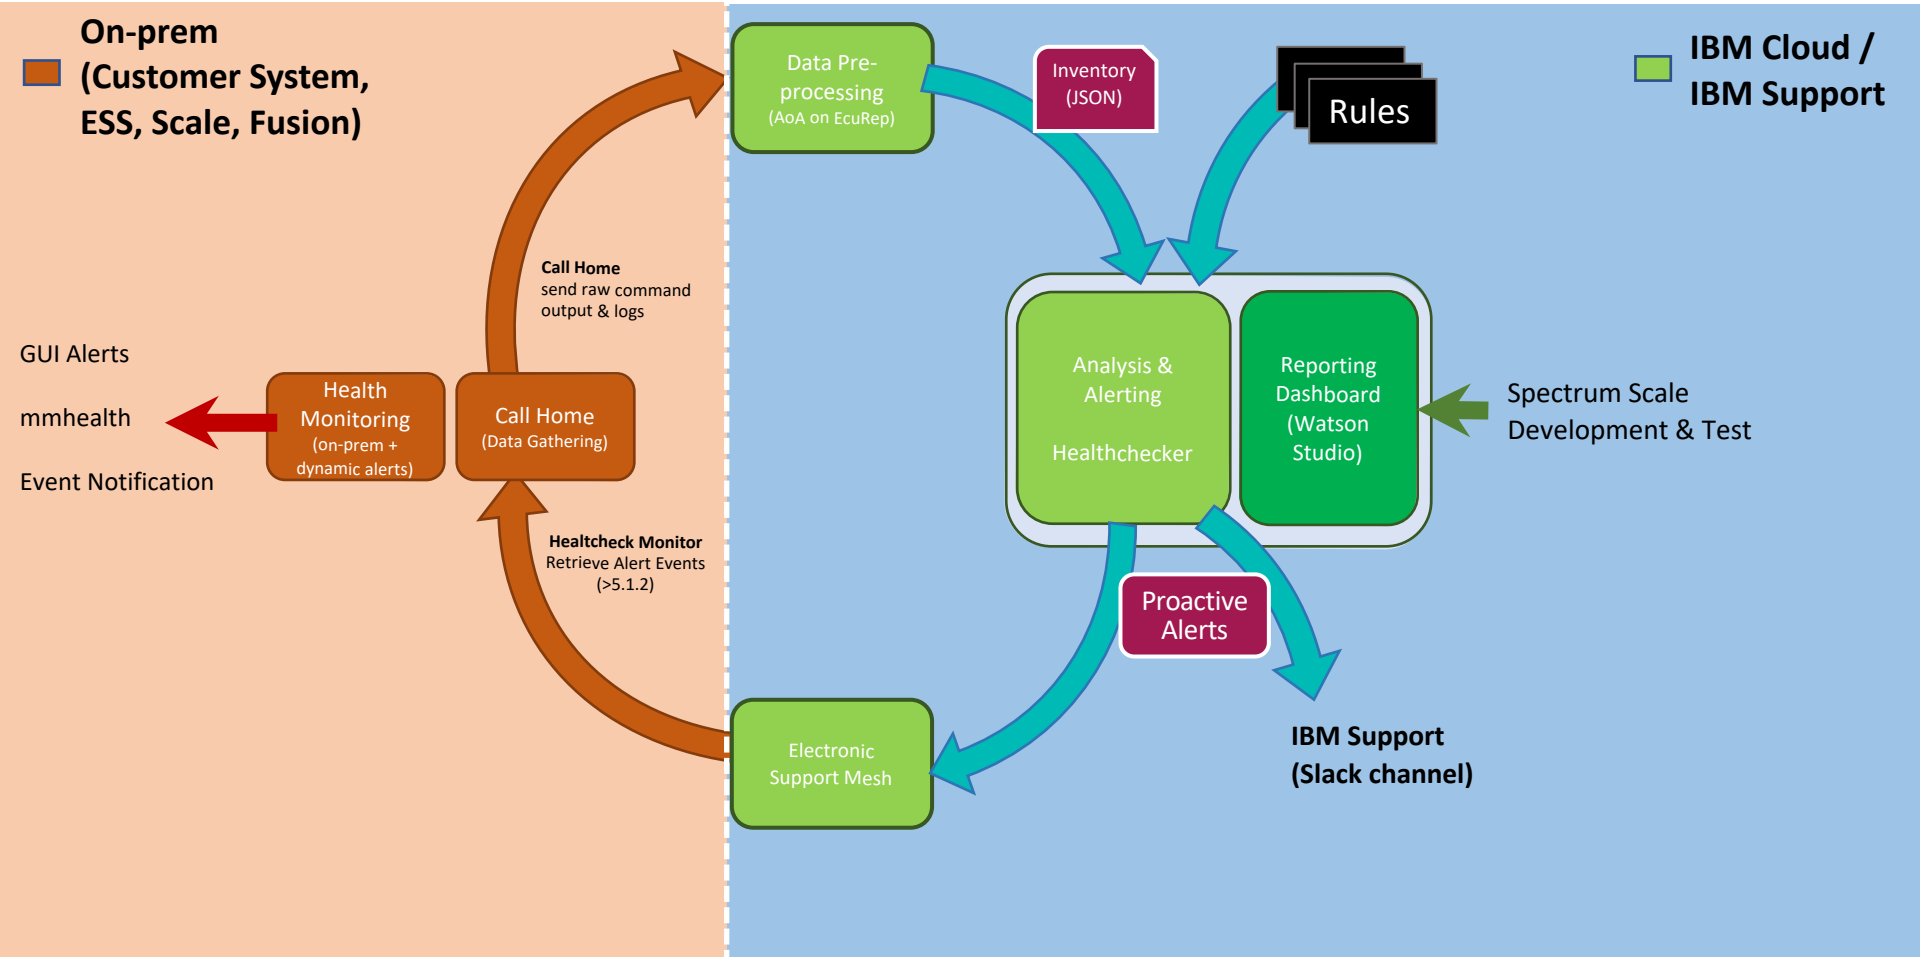

Reminder: Please Enable Call Home

Spectrum Scale

Enable Call Home on your Spectrum Scale clusters !

- Get the benefits now! and help our service team
- Help to improve the product development

Call Home Client Value:

- Improved service response time and quicker resolution due to
  - Data immediately available to analyse a ticket (Call Home , FTDC)
  - Tools and Salesforce linkage to give support team a quick overview on customer system
  - Automatically open PMRs for hardware failures
- Better customer support experience (consumability/ease of use)
  - Easy and fast sharing of diagnostics data with support
  - Easy way to see history of system configuration changes (config diff)
- Provide customer insights reports for support and development
- Proactively detect issues (Healthchecker)

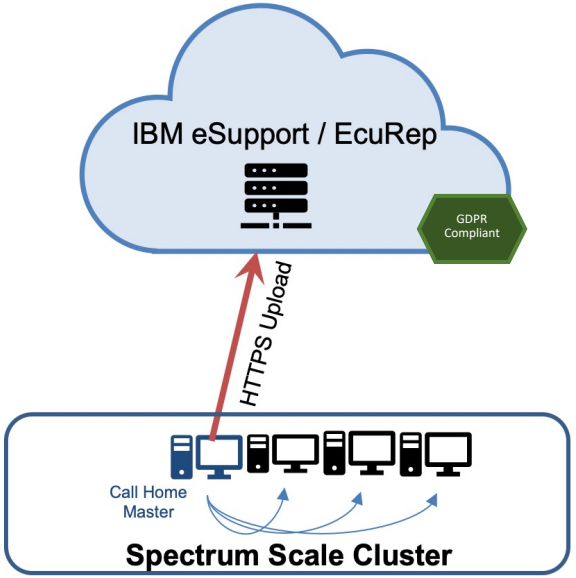

Recap:

Call Home is an integrated feature of Spectrum Scale software (available on ESS, Fusion and CNSA)

1. Collects & upload configuration, health, performance/usage data (no PI) from the cluster nodes
   - daily/weekly data collection and regular heartbeats*
2. Event-triggered data collection and upload (event specific data)*
3. Automatically open PMRs for hardware failures (ESS 3500 only)

*See Spectrum Scale Knowledge Center for details of collected and transferred configuration data and logs

Thank you for using IBM Spectrum Scale!