

IBM Spectrum Scale:
ESS and Spectrum Scale
monitoring using the GUI,
REST API, SNMP and eMail

—
Stefan Roth
Developer



Outline

- Introduction
- Monitor System Health in GUI
- Email Event Notification
- SNMP Event Notification
- Query System Health with REST-API
- Live Demo

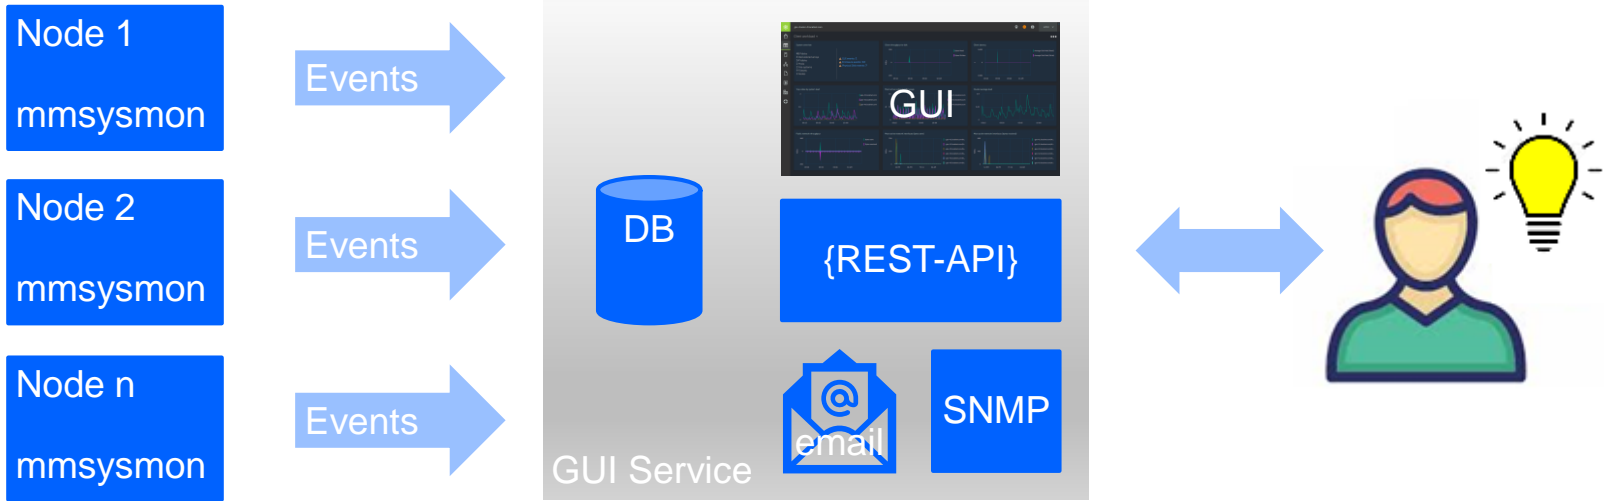


Introduction



Introduction

- Focus of this presentation is health monitoring (**mmhealth**)



- GUI service provides four interfaces:
 - GUI
 - REST-API
 - Email
 - SNMP

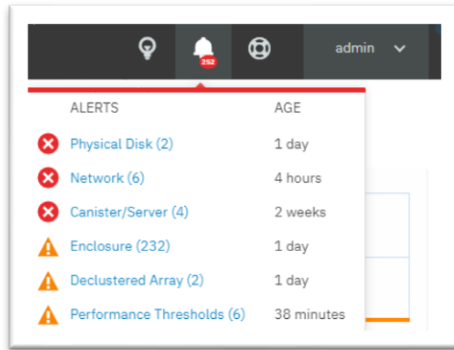
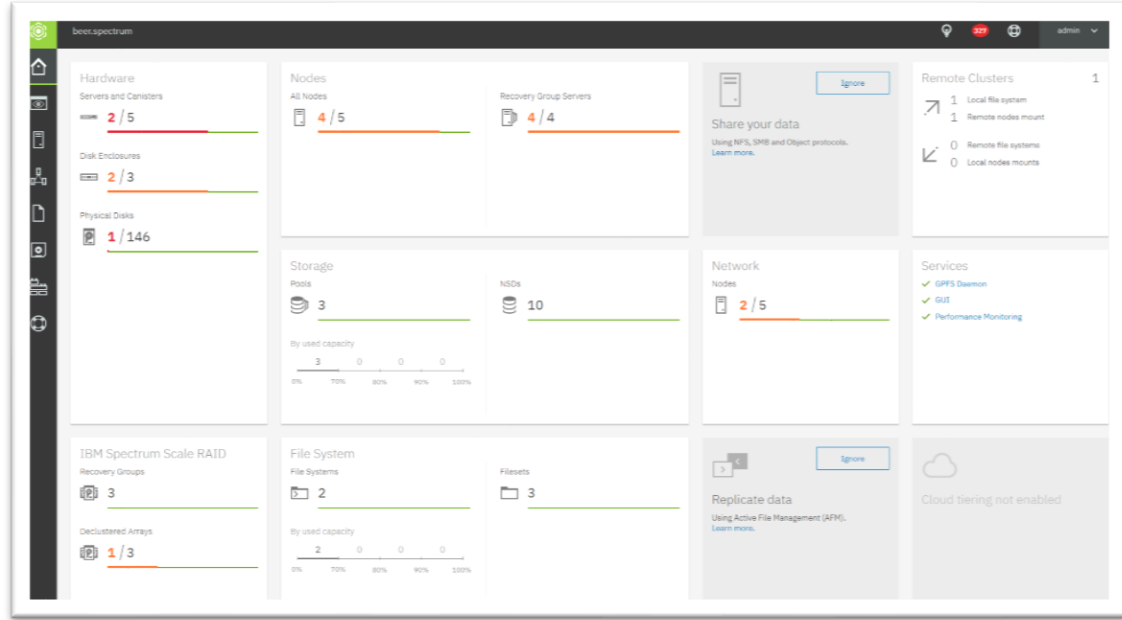
Monitor System Health in GUI



Get a rough Overview

Overview page

- Displays overall health state (healthy, degraded and failed entities)
- Click on numbers to drill down
- List of services with their health state

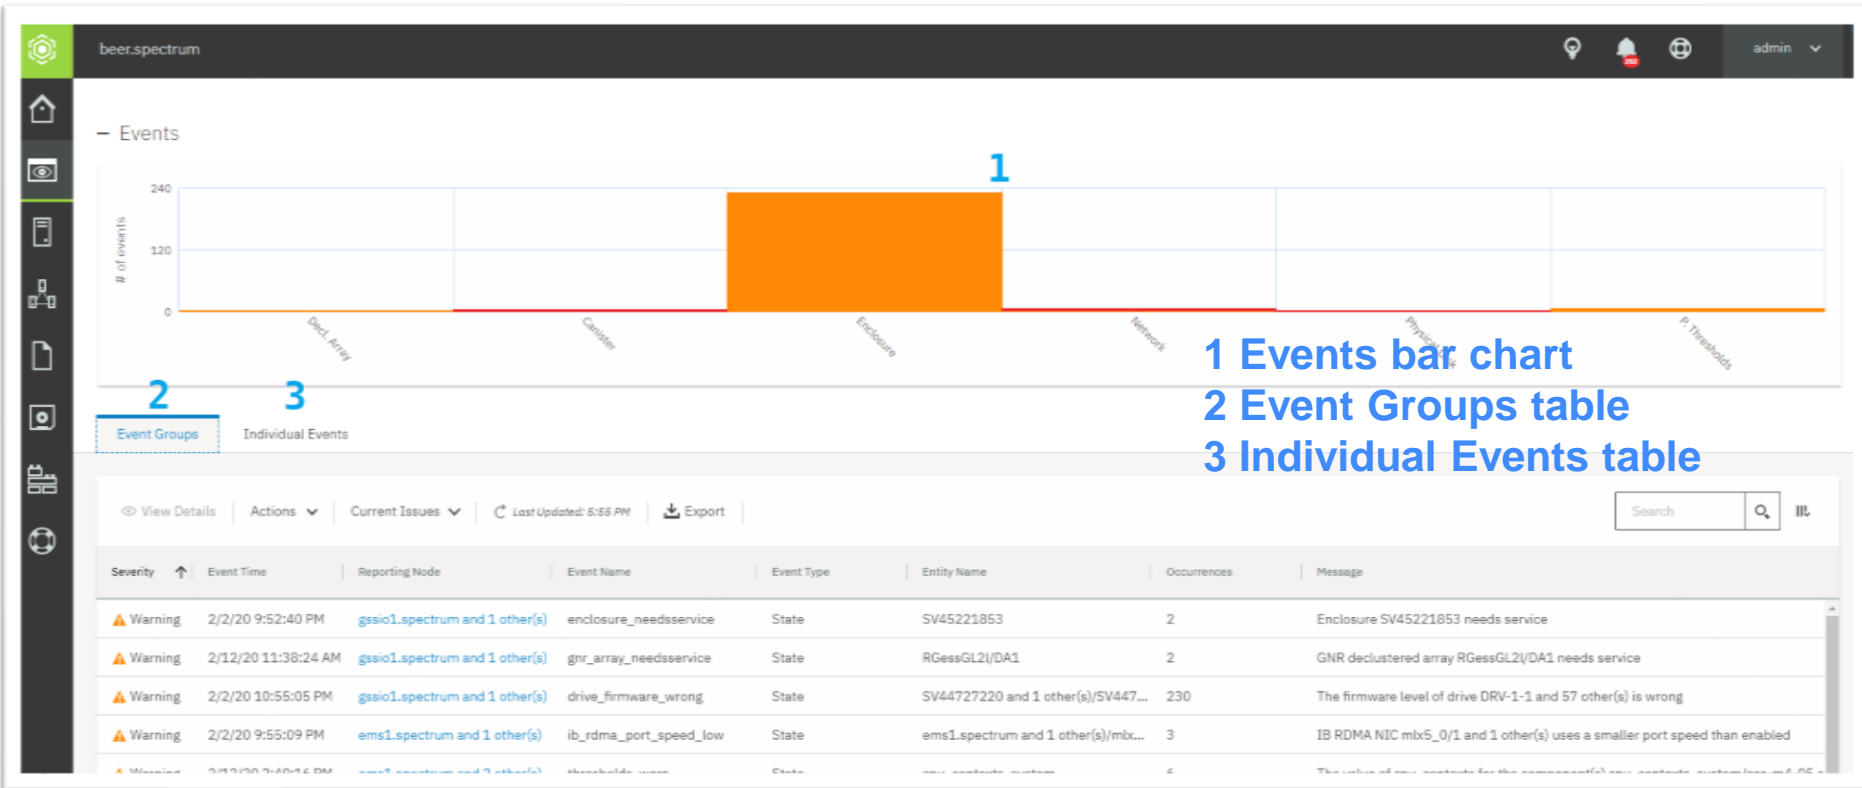


GUI header

- Displays numbers of current issues (warning and error events)
- Move mouse over: Popup displays issues separated by component
- Click: Go to events page to see details

Health Events

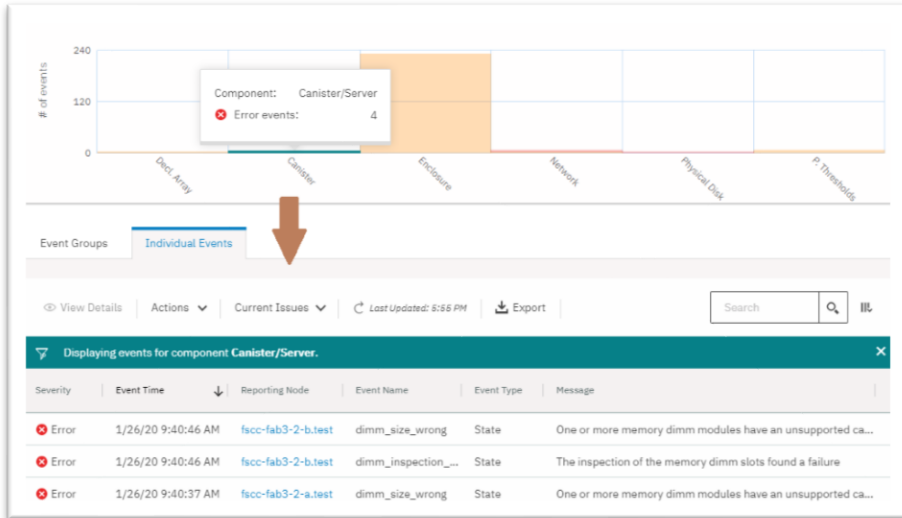
Monitoring > Events is main page for dealing with health events



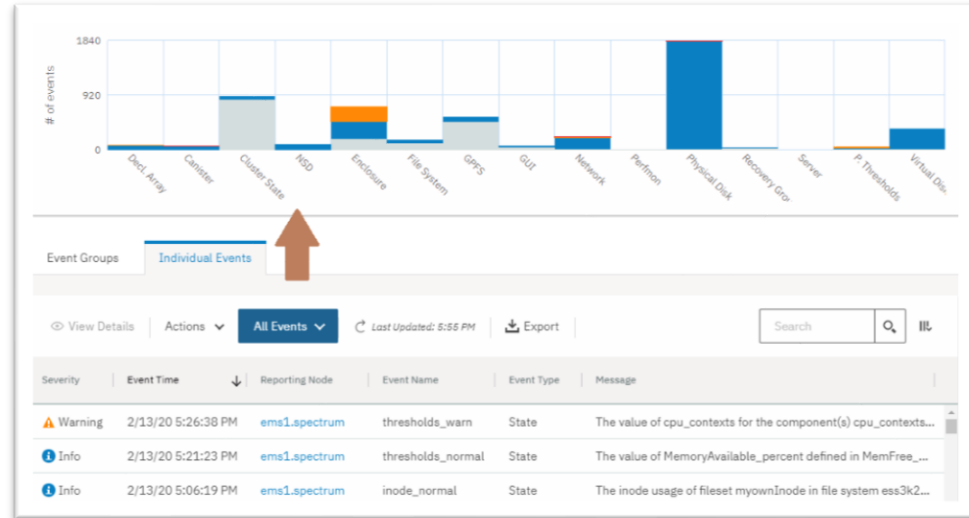
Events Bar Chart

- List numbers of events by component
- Color indicates event severities

Click on bar to filter the table by component

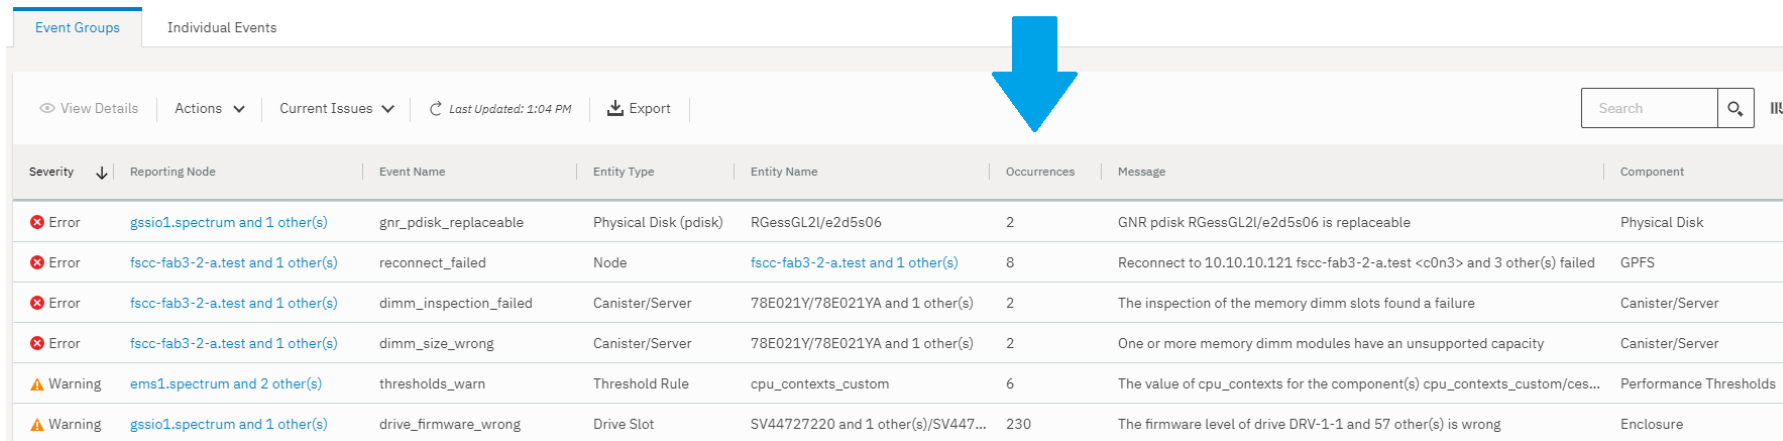


Using table filter has impact on bar chart



Event Groups

- Display many occurrences of the same event in one row
- Useful for high numbers of events



The screenshot shows the 'Event Groups' interface in IBM Spectrum Scale. The top navigation bar includes 'Event Groups' (selected), 'Individual Events', 'View Details', 'Actions', 'Current Issues', 'Last Updated: 1:04 PM', and 'Export'. A search bar is on the right. A table below lists event groups with columns for Severity, Reporting Node, Event Name, Entity Type, Entity Name, Occurrences, Message, and Component. A blue arrow points to the 'View Details' link in the top navigation bar.

| Severity | Reporting Node | Event Name | Entity Type | Entity Name | Occurrences | Message | Component |
|----------|--|------------------------|-----------------------|--|-------------|---|------------------------|
| Error | gssio1.spectrum and 1 other(s) | gnr_pdisk_replaceable | Physical Disk (pdisk) | RGessGL2/e2d5s06 | 2 | GNR pdisk RGessGL2/e2d5s06 is replaceable | Physical Disk |
| Error | fsc-fab3-2-a.test and 1 other(s) | reconnect_failed | Node | fsc-fab3-2-a.test and 1 other(s) | 8 | Reconnect to 10.10.10.121 fsc-fab3-2-a.test <c0n3> and 3 other(s) failed | GPFS |
| Error | fsc-fab3-2-a.test and 1 other(s) | dimm_inspection_failed | Canister/Server | 78E021Y/78E021YA and 1 other(s) | 2 | The inspection of the memory dimm slots found a failure | Canister/Server |
| Error | fsc-fab3-2-a.test and 1 other(s) | dimm_size_wrong | Canister/Server | 78E021Y/78E021YA and 1 other(s) | 2 | One or more memory dimm modules have an unsupported capacity | Canister/Server |
| Warning | ems1.spectrum and 2 other(s) | thresholds_warn | Threshold Rule | cpu_contexts_custom | 6 | The value of cpu_contexts for the component(s) cpu_contexts_custom/ces... | Performance Thresholds |
| Warning | gssio1.spectrum and 1 other(s) | drive_firmware_wrong | Drive Slot | SV44727220 and 1 other(s)/SV447... | 230 | The firmware level of drive DRV-1-1 and 57 other(s) is wrong | Enclosure |

- Double click (or *View Details*) on row lists individual events

Individual Events

- One row for each individual event

Event Groups **Individual Events**

View Details Actions Current Issues Last Updated: 10:52 AM Export Search

| Severity | Event Time | Reporting Node | Event Name | Event Type | Action | Entity Type | Entity Name | Message |
|----------|---------------------|---------------------|--------------------------|------------|-------------------------------|-----------------------|---------------------|--|
| Warning | 2/14/20 10:13:08 AM | gss-41.localnet.com | gui_login_attempt_failed | Notice | Mark Selected Notices as Read | Node | gss-41.localnet.com | A login attempt for user admin failed. |
| Warning | 2/13/20 11:15:36 AM | gss-43.localnet.com | unmounted_fs_check | State | Run Fix Procedure | File system | gpfs0 | The filesystem gpfs0 is probably needed, but not mounted |
| Warning | 2/13/20 11:12:23 AM | gss-43.localnet.com | unmounted_fs_check | State | Run Fix Procedure | File system | posix | The filesystem posix is probably needed, but not mounted |
| Warning | 2/13/20 11:12:16 AM | gss-42.localnet.com | unmounted_fs_check | State | Run Fix Procedure | File system | posix | The filesystem posix is probably needed, but not mounted |
| Warning | 1/16/20 1:51:36 PM | gss-41.localnet.com | gnr_pdisk_diagnosing | Notice | Mark Selected Notices as Read | Physical Disk (pdisk) | RG2/pdisk46 | GNR pdisk RG2/pdisk46 is diagnosing |
| Warning | 1/16/20 1:49:13 PM | gss-42.localnet.com | gnr_pdisk_diagnosing | Notice | Mark Selected Notices as Read | Physical Disk (pdisk) | RG2/pdisk26 | GNR pdisk RG2/pdisk26 is diagnosing |

- Fix procedures available for many issues (*Run Fix Procedure*)

Fix Procedure: Update Drive Firmware

Update Drive Firmware

The firmware level of drive DRV-24 is wrong

- Update firmware of drive ENC123000/DRV-24.
- Update firmware on all 4 drives where the firmware level is wrong.

Updating firmware on many drives can take several minutes up to hours, depending on the number of drives.
Click *Finish* to perform the selected action.

Status: No task running

Cancel Back Finish

Fix Procedure: Replace Disks

Prerequisite

Prerequisite

Replace Disk

Before you get started

Do not replace the disk yet! You are about to replace disk named e2d5e06, located at Rack rack U09-12, Enclosure 1818-80E-SV45221853 Drawer 5 Slot 6. Make sure you have a replacement disk with the FRU 46W6911 before proceeding.

- Use a replacement disk with a different FRU than 46W6911

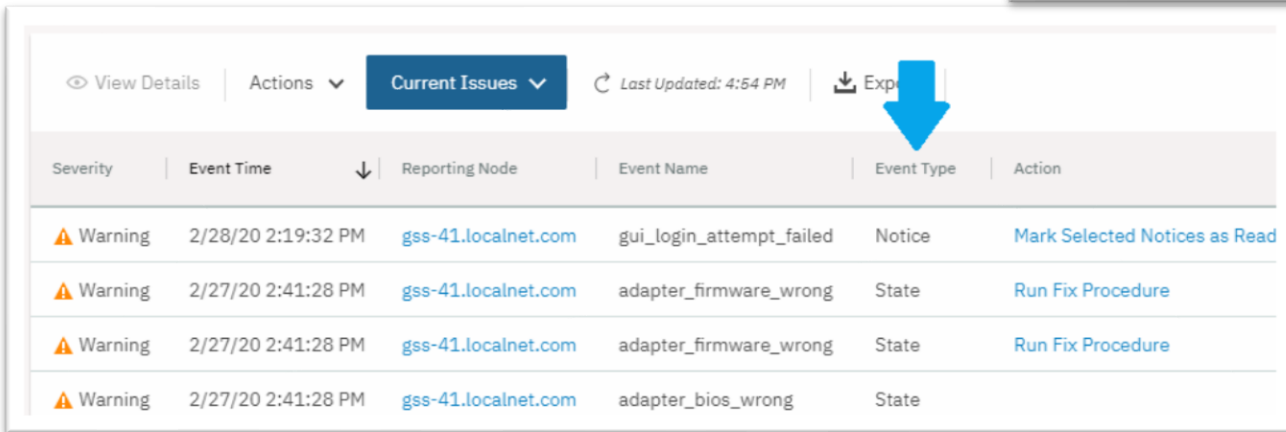
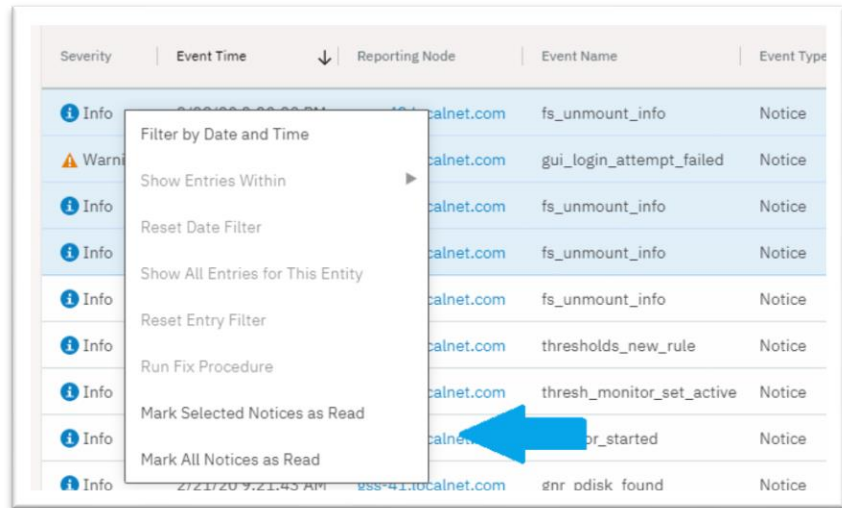
If you have a replacement disk, click Next to prepare the disk for removal. This suspends the I/O to the disk and turns off its power.

Status: No task running

Cancel Back Next

Event Types

| Event Type | Description |
|------------|---|
| State | <ul style="list-style-type: none"> Events that are generated because of state changes Impact on overall state of entities Source: mmhealth node show <component> Example: <i>thresholds_warn</i> and <i>thresholds_normal</i> |
| Notice | <ul style="list-style-type: none"> Events that are not caused by a state change Cannot become inactive on it's own Mark notices as Read to make notices historical Source: mmhealth node eventlog Example: <i>gui_login_attempt_failed</i> |



Event Filters

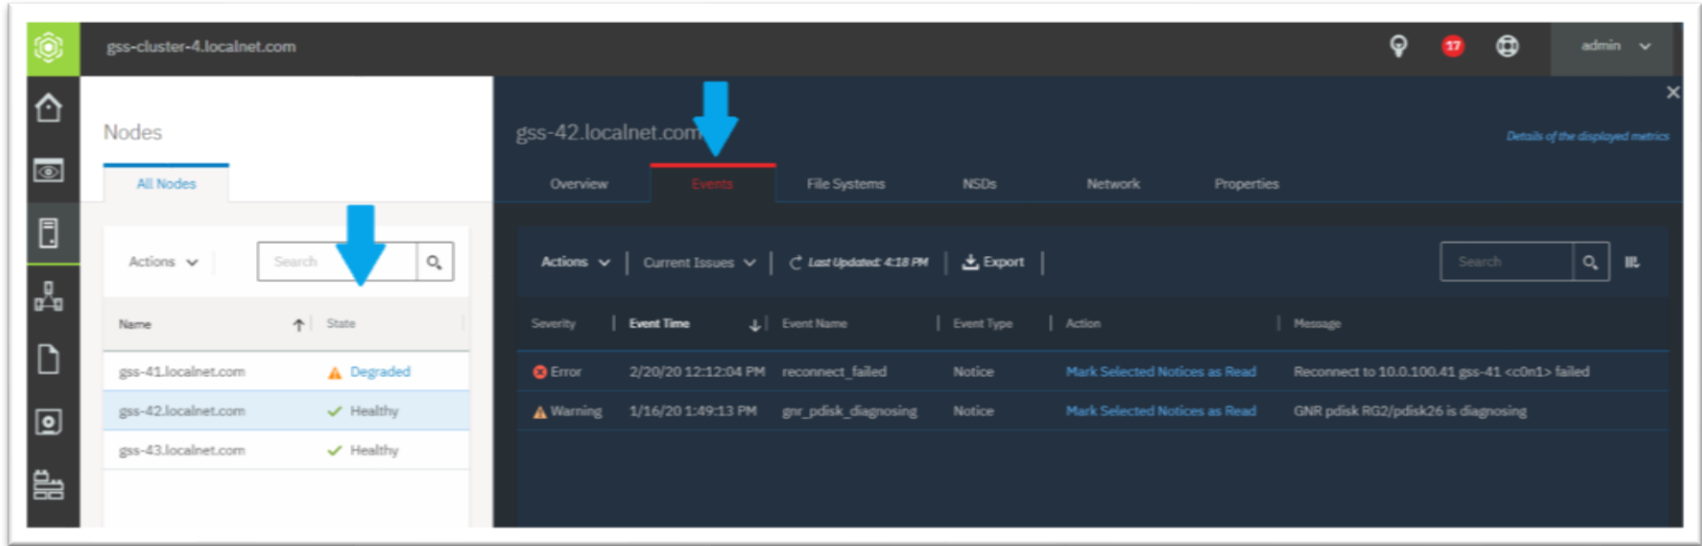
| Filter | Description |
|----------------|--|
| Current Issues | List active events with severity error and warning |
| Current State | List active events that are generated because of a state change |
| Notices | List active notices |
| All Events | List all events, irrespective of severity and type This shows both active and historical events |

- The filter is available in both tables (*Event Groups* and *Individual Events*)
- Historical events have a grey icon
- GUI keeps the last 10.000 historical events

The screenshot shows the 'Individual Events' view in the IBM Spectrum Scale GUI. At the top, there are tabs for 'Event Groups' and 'Individual Events'. Below the tabs, there is a toolbar with 'View Details', 'Actions', a filter dropdown menu, 'Last Updated: 2:20 PM', and 'Export'. The filter dropdown menu is open, showing options: 'All Events', 'Current Issues', 'Current State', 'Notices', and 'All Events'. The table below displays a list of events with columns for Severity, Event Time, Event Name, Event Type, Action, Entity Type, and Entity Name.

| Severity | Event Time | Event Name | Event Type | Action | Entity Type | Entity Name |
|----------|--------------------|--|------------|-------------------------------|-------------|-----------------|
| Info | 2/28/20 2:21:50 PM | mounted_fs_check | State | | File system | gpfs0 |
| Warning | 2/28/20 2:20:20 PM | unmounted_fs_check | State | | File system | gpfs0 |
| Info | 2/28/20 2:20:20 PM | gss-43.localnet.com fs_unmount_info | Notice | Mark Selected Notices as Read | File system | gpfs0 |
| Warning | 2/28/20 2:19:32 PM | gss-41.localnet.com gui_login_attempt_failed | Notice | Mark Selected Notices as Read | Node | gss-41.local... |

Health monitoring in other GUI pages



- Some GUI pages list entities (nodes, filesystems, SMB shares, AFM relationships, ...)
- Display health state of each entity
- List events for selected entity

Monitoring ESS Hardware

- *Monitoring > Hardware and Monitoring > Hardware Details* are useful to monitor ESS

Hardware

Disk Enclosure 5141-AF8-78E021Y

Enclosure | Top Canister/Server | Bottom Canister/Server

Rack | Location

rack | 14

Right Power Supply Unit
State: Healthy

Health state: ✓ Healthy [See Details](#)

Serial number: 78E021Y

Machine type and model: 5141-AF8

Rack: rack

Rack slot: U17-18

Building block: group2

Building block type: SN0

Vendor: IBM-ESS

Drive slots: 24

Description: IBM 5141-AF8 NVMe/PCI switch enclosure

Hardware Details

[Refresh](#)

Filter

Display "unhealthy" devices

- beer.spectrum
 - group2
 - Bottom Canister/Server
 - Boot Drives
 - ATA-2**
 - ATA-3
 - CPU's
 - Memory (DIMM)
 - Fans

Boot Drive ATA-2

Building Block group2 > Bottom Canister/Server > Boot Drive ATA-2

Serial number: 174167401563

Port: ATA-2

Boot drive: ✓ The bootdrive attached to port ata-2 is available [Show Detail](#)

Boot drive smart: ✓ The smart assessment of bootdrive 174167401563 attached

Email Event Notification



Introduction to Email Event Notification

- GUI can send emails for health events
- **Email for each event**
 - Consolodate many events in one email if they appear in short time frame
- **Daily email for all events** of last 24 hours (sent at 4:15)
- **Daily quota reports** (sent at 2:15)
 - List user, group and fileset quotas with exceeded quota limits

Email Server

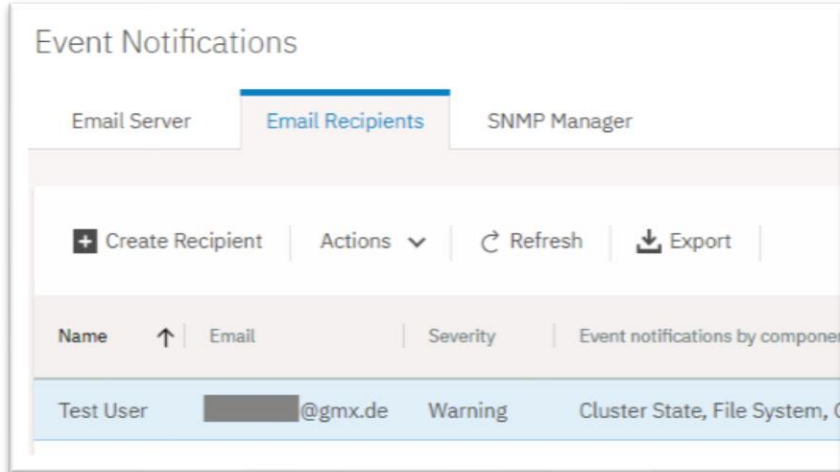
- The email server is configured in *Monitoring > Event Notifications > Email Server*
- *Test Email* can be used to send a test email

The screenshot shows the 'Email Server' configuration page. At the top, there are three tabs: 'Email Server' (selected), 'Email Recipients', and 'SNMP Manager'. Below the tabs, there is a toggle switch for 'Email notifications enabled' which is turned on. The configuration fields are as follows:

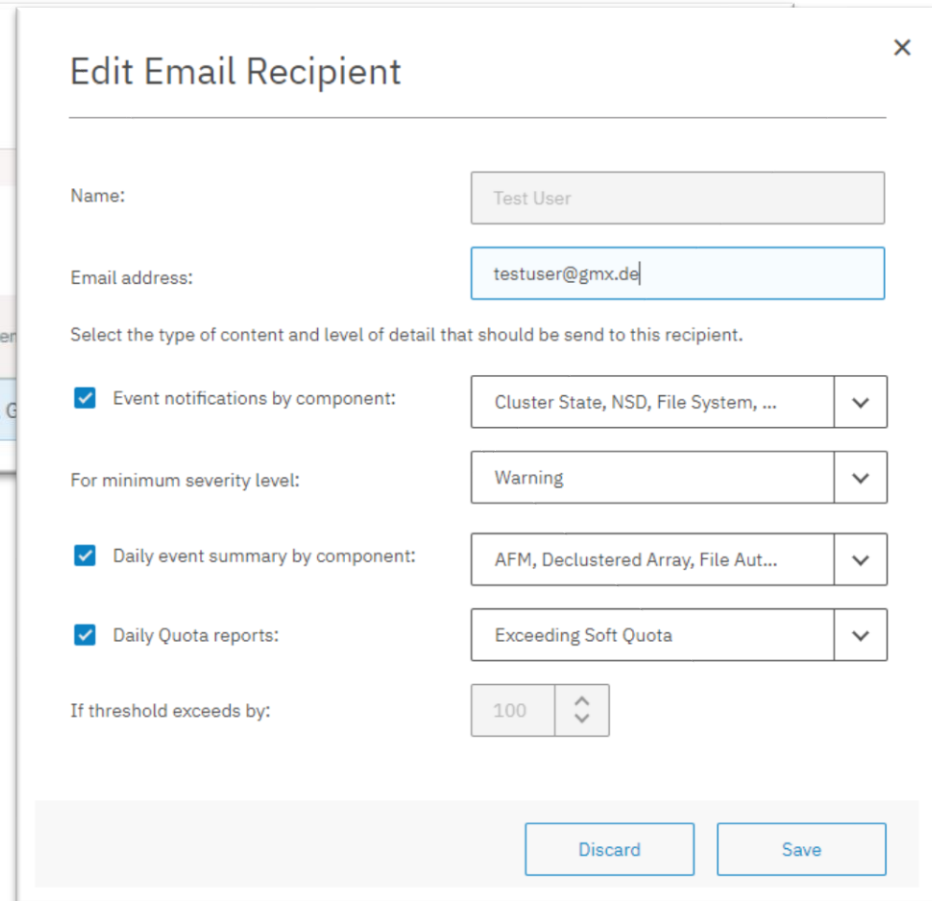
- IP address or host name:** mail.gmx.net
- Port:** 587
- Sender's email address:** testuser@gmx.de
- Password:** (masked with dots)
- Use different login: (empty field)
- Sender's name:** Spectrum Scale GUI
- Subject:** &messageId &message (with a 'Variable' button next to it)
- Header:** (empty text area)
- Footer:** (empty text area)
- Test email address:** (empty text area)

At the bottom right, there are two buttons: 'Discard' and 'Configure'.

Email Recipients



- Three types of emails
 - For each individual event
 - Daily event summary
 - Daily quota report
- Select events components
- Select minimum severity



Email Notifications

- Tailor the email
 - Sender's name
 - Subject (text and variables)
 - Header and footer

Sender's name: Subject: [Variable](#)

Header:

Footer:

- Email lists events in table

unmounted_fs_check The filesystem gpfs0 is probably needed, but not mounted

From: IBM Spectrum Scale

This email has been automatically sent by IBM Spectrum Scale.

| Time | Cluster Name | Reporting Node | Event Name | Entity Type | Entity Name | Severity | Message |
|-------------------------|----------------------------|---------------------|--------------------|-------------|-------------|----------|--|
| 12.02.2020 16:00:23.520 | gss-cluster-4.localnet.com | gss-43.localnet.com | unmounted_fs_check | FILESYSTEM | gpfs0 | WARNING | The filesystem gpfs0 is probably needed, but not mounted |

Cheers.

Daily Quota Reports (Example: Fileset quota, soft quota exceeded)

+ Create Fileset Quota | Actions ▾ | Refresh | Export

| Name | File System | Used Capacity | Soft Limit (Capacity) | Soft Limit (% Capacity Used) ↓ | Hard Limit (Capacity) | Hard Limit (% Capacity Used) |
|-------|-------------|---------------|-----------------------|--------------------------------|-----------------------|------------------------------|
| fset2 | gpfs0 | 890.24 MiB | 500.00 MiB | 178% | 1.00 GiB | 86 % |



Daily Quota reports: Exceeding Soft Quota ▾



Digest Mail for quotas.
Von: IBM Spectrum Scale

This email has been automatically sent by IBM Spectrum Scale.

| Name ID | Device | Fileset | Hard limit | Soft limit | Usage | Grace time | Hard limit (inodes) | Soft limit (inodes) | Usage (inodes) | Grace time (inodes) |
|---------|--------|---------|------------|------------|-------|------------|---------------------|---------------------|----------------|---------------------|
| fset2 | FID: 2 | gpfs0 | 1.00 GiB | 500.00 MiB | 87 % | 7days | 200 | 100 | 7 % | none |

Cheers.

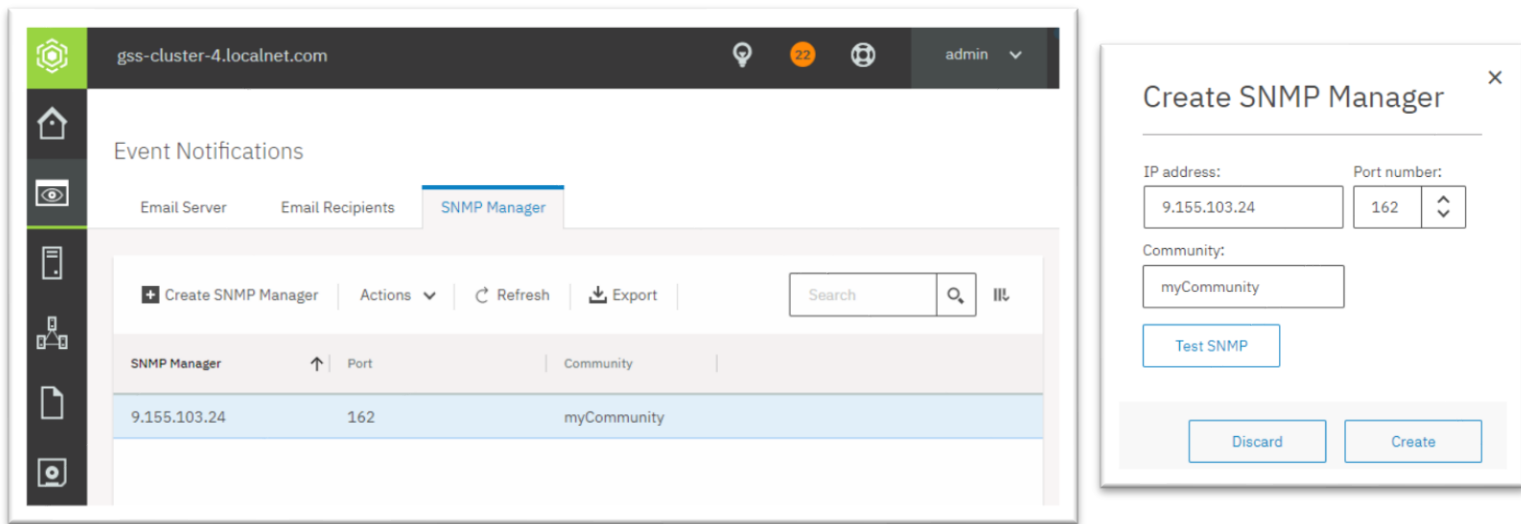
Note: Usage and Usage (inodes) in email always relate to hard quota limits

SNMP Event Notification



Introduction to SNMP Event Notification

- GUI can send an SNMP notification for each health event
- SNMP notification receivers can be configured in GUI page *Monitoring > Event Notifications > SNMP Manager*



The screenshot displays the 'Event Notifications' section of the GUI, specifically the 'SNMP Manager' tab. The main interface shows a table with one entry:

| SNMP Manager | Port | Community |
|--------------|------|-------------|
| 9.155.103.24 | 162 | myCommunity |

Below the table, a modal dialog titled 'Create SNMP Manager' is open. It contains the following fields and buttons:

- IP address: 9.155.103.24
- Port number: 162
- Community: myCommunity
- Buttons: Test SNMP, Discard, Create

- Applications like Nagios can be configured to receive GUI SNMP notifications

GUI SNMP Event Notification Details

- The SNMP Event Receiver receives notifications for all health events
 - No filter to filter events to send as SNMP (severity, component, event name, etc.)
- SNMP Version is SNMP v2
- SMMP MIB is located on each GUI node
`/usr/lpp/mmfs/gui/IBM-SPECTRUM-SCALE-GUI-MIB.txt`
- GUI does not support to query information via SNMP
 - Use REST API to query informaton

| OID Range | Description |
|-----------------------------|----------------------------|
| .1.3.6.1.4.1.2.6.212.10.0.1 | GUI Event Notification OID |
| .1.3.6.1.4.1.2.6.212.10.1.1 | Cluster ID |
| .1.3.6.1.4.1.2.6.212.10.1.2 | Entity Type |
| .1.3.6.1.4.1.2.6.212.10.1.3 | Entity Name |
| .1.3.6.1.4.1.2.6.212.10.1.4 | Component |
| .1.3.6.1.4.1.2.6.212.10.1.5 | Severity |
| .1.3.6.1.4.1.2.6.212.10.1.6 | Event Time |
| .1.3.6.1.4.1.2.6.212.10.1.7 | Event Name |
| .1.3.6.1.4.1.2.6.212.10.1.8 | Message |
| .1.3.6.1.4.1.2.6.212.10.1.9 | Reporting Node |

GUI SNMP versus GPFS Core SNMP

- GPFS core also provides an SNMP interface
 - This is out of scope for this presentation
 - Do not confuse this with GUI SNMP
 - The MIB file is located on every GPFS node:
/usr/lpp/mmfs/data/GPFS-MIB.txt
 - More information:
https://www.ibm.com/support/knowledgecenter/STXKQY_5.0.4/com.ibm.spectrum.scale.v5r04.doc/bl1adv_snmp.htm

Configure snmptrapd to call a Custom Script

snmptrapd can

- receive SNMP notifications from GUI and
- call a custom script that handles the event in the SNMP notification

Instructions

1. Install the **snmptrapd** application

```
yum install net-snmp
```

2. Configure **/etc/snmp/snmptrapd.conf** to call a script that for each SNMP event

```
disableAuthorization yes
```

```
traphandle default /usr/local/bin/traphandler
```

3. Open port 162 for UDP in firewall

```
iptables -I INPUT -p udp -m udp --dport 162 -j ACCEPT
```

```
iptables-save
```

Configure snmptrapd to call a Custom Script (continued)

4. Check in which directory the MIB has to be placed

```
net-snmp-config --default-mibdirs  
/root/.snmp/mibs:/usr/share/snmp/mibs
```

5. Place the GUI MIB to one of those directories *)

```
cp IBM-SPECTRUM-SCALE-GUI-MIB.txt /usr/share/snmp/mibs/
```

6. Add the following line to **/etc/snmp/snmp.conf**

```
mibs +IBM-SPECTRUM-SCALE-GUI-MIB
```

7. Start the SNMP event receiver

```
systemctl start snmptrapd
```

*) Note: The MIB file can be found on each GUI node in `/usr/lpp/mmfs/gui/`

Create a SNMP Handler Script for GUI notifications

- Create the sample script `/usr/local/bin/traphandler`

```
#!/bin/sh

read host
echo ===== START =====
echo Spectrum Scale Event SNMP Notification received from $host:

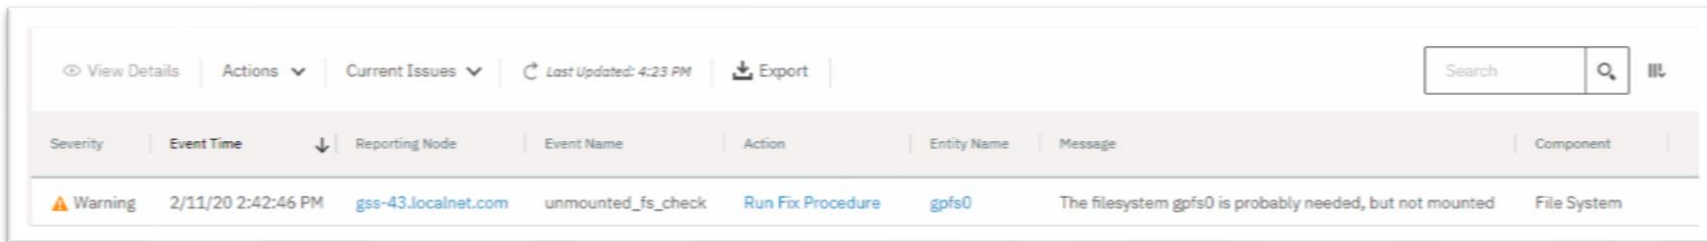
while read oid val
do
  if [[ "$oid" == IBM-SPECTRUM-SCALE-GUI-MIB::ibmSpectrumScaleGuiEvent* ]]
  then
    key="${oid:52}"
    echo $key = $val
  fi
done
echo ===== END =====
```

- Make the script executable
`chmod +x /usr/local/bin/traphandler`

Sample Script output for a health event

- The script by default logs (echo) to `/var/log/messages`
- Sample output for a `unmounted_fs_check` event:

```
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: ===== START =====  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: Spectrum Scale Event SNMP Notification received from gss-41.localnet.com:  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: Cluster = "10583479681613060679"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: EntityType = "FILESYSTEM"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: EntityName = "gpfs0"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: Component = "FILESYSTEM"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: Severity = "WARNING"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: Time = "11.02.2020 14:42:46.890"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: Name = "unmounted_fs_check"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: Message = "The filesystem gpfs0 is probably needed, but not mounted"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: ReportingNode = "gss-43.localnet.com"  
Feb 11 14:42:55 sonas-cli-test08 snmptrapd: ===== END =====
```



The screenshot shows a GUI interface for monitoring events. At the top, there are navigation options: 'View Details', 'Actions', 'Current Issues', 'Last Updated: 4:23 PM', and 'Export'. A search bar is located on the right. Below this is a table of events with the following columns: Severity, Event Time, Reporting Node, Event Name, Action, Entity Name, Message, and Component.

| Severity | Event Time | Reporting Node | Event Name | Action | Entity Name | Message | Component |
|----------|--------------------|---------------------|--------------------|-------------------|-------------|--|-------------|
| Warning | 2/11/20 2:42:46 PM | gss-43.localnet.com | unmounted_fs_check | Run Fix Procedure | gpfs0 | The filesystem gpfs0 is probably needed, but not mounted | File System |

System Health in REST API



Introduction to REST API calls for health monitoring

- There are two REST calls to query health events and health states

| | | |
|-----|--|--------------------------|
| GET | /scalemgmt/v2/nodes/{name}/health/events | Get System Health events |
| GET | /scalemgmt/v2/nodes/{name}/health/states | Get System Health states |

- REST calls allow to select event by fields (components, nodes, severity, state, event name, ...)
- Provides paging if more than 1000 results
 - https://www.ibm.com/support/knowledgecenter/STXKQY_5.0.4/com.ibm.spectrum.scale.v5r04.doc/bl1adm_apiv2paging.htm

Query health events using the REST-API

- Example: Query events of component FILESYSTEM that are not severity info

```
https://<gui>/scalemgmt/v2/nodes/all/health/events?filter=severity%21%3DINFO%2Ccomponent%3DFILESYSTEM'
```

- Request URL without encoding (easier to read)

```
https://<gui>/scalemgmt/v2/nodes/all/health/events?filter=severity!=INFO,component=FILESYSTEM'
```

- REST Response:

```
{
  "events": [
    {
      "activeSince": "2020-02-13 11:12:16,770",
      "component": "FILESYSTEM",
      "description": "An internally mounted or a declared but not mounted filesystem was detected",
      "entityName": "posix",
      "entityType": "FILESYSTEM",
      "message": "The filesystem posix is probably needed, but not mounted",
      "name": "unmounted_fs_check",
      "reportingNode": "gss-42.localnet.com",
      "severity": "WARNING",
      "state": "DEGRADED",
      "type": "STATE_CHANGE",
      "userAction": "Run mmismount all_local to verify that all expected
    },
    {
      "activeSince": "2020-02-13 11:12:23,819",
      "component": "FILESYSTEM",
      ...
  ]
}
```

| Severity | Event Time | Event Name | Entity Type | Entity Name | Message | Component |
|----------|---------------------|--------------------|-------------|-------------|--|-------------|
| Warning | 2/13/20 11:15:36 AM | unmounted_fs_check | File system | gpts0 | The filesystem gpts0 is probably needed, but not mounted | File System |
| Warning | 2/13/20 11:12:23 AM | unmounted_fs_check | File system | posix | The filesystem posix is probably needed, but not mounted | File System |
| Warning | 2/13/20 11:12:16 AM | unmounted_fs_check | File system | posix | The filesystem posix is probably needed, but not mounted | File System |

Query health states using the REST-API

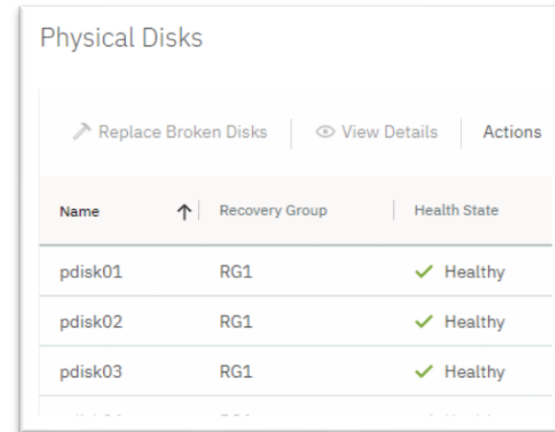
- Example: Query state for all pdisks
- Limit output to a few event fields (entityName, parentName, state)

```
https://<gui>/scalemgmt/v2/nodes/all/health/states?fields=entityName%2CparentName%2Cstate&filter=entityType%3DPDISK
```

```
https://<gui>/scalemgmt/v2/nodes/all/health/states?fields=entityName,parentName,state&filter=entityType=PDISK
```

- REST Response:

```
{
  "states": [
    {
      "entityName": "pdisk01",
      "entityType": "PDISK",
      "parentName": "RG1",
      "state": "HEALTHY"
    },
    {
      "entityName": "pdisk02",
      "entityType": "PDISK",
      "parentName": "RG1",
      "state": "HEALTHY"
    },
    ...
  ]
}
```



The screenshot shows a web interface titled "Physical Disks". At the top, there are three buttons: "Replace Broken Disks" (with a right-pointing arrow), "View Details" (with an eye icon), and "Actions". Below the buttons is a table with the following columns: "Name", "Recovery Group", and "Health State". The table contains three rows of data, each representing a physical disk. The first row shows "pdisk01" in "RG1" with a "Healthy" status (indicated by a green checkmark). The second row shows "pdisk02" in "RG1" with a "Healthy" status. The third row shows "pdisk03" in "RG1" with a "Healthy" status. There are ellipses at the end of each row, indicating that there are more disks in the system.


| Name | Recovery Group | Health State |
|---------|----------------|--------------|
| pdisk01 | RG1 | ✓ Healthy |
| pdisk02 | RG1 | ✓ Healthy |
| pdisk03 | RG1 | ✓ Healthy |

Live Demo



Thank you!

Provide Feedback ×



Tell IBM What You Think

Let us know what you think about IBM Spectrum Scale. It takes only a couple of minutes for you to help us improve our service. [IBM Privacy Policy](#)

Please help us to improve Spectrum Scale with your feedback

- If you get a survey in email or a popup from the GUI, please respond
- We read every single reply