Health Monitoring Deep Dive

Spectrum Scale German User Meeting 2023
Sindelfingen, Germany – March 22-23, 2023

Helene Wassmann, Norbert Schuld, Mathias Dietz

Disclaimer

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

IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

Agenda

Spectrum Scale

- Cluster-wide eventlog using Webhooks
- mmhealth enhancements
- AFM Monitoring

CLUSTER-WIDE EVENTLOG USING WEBHOOKS

Webhook for mmhealth events

What is a webhook (Source: Wikipedia)

Webhooks are "user-defined HTTP callbacks". They are usually triggered by some event, such as pushing code to a repository or a comment being posted to a blog. When that event occurs, the source site makes an HTTP request to the URL configured for the webhook. Users can configure them to cause events on one site to invoke behavior on another.

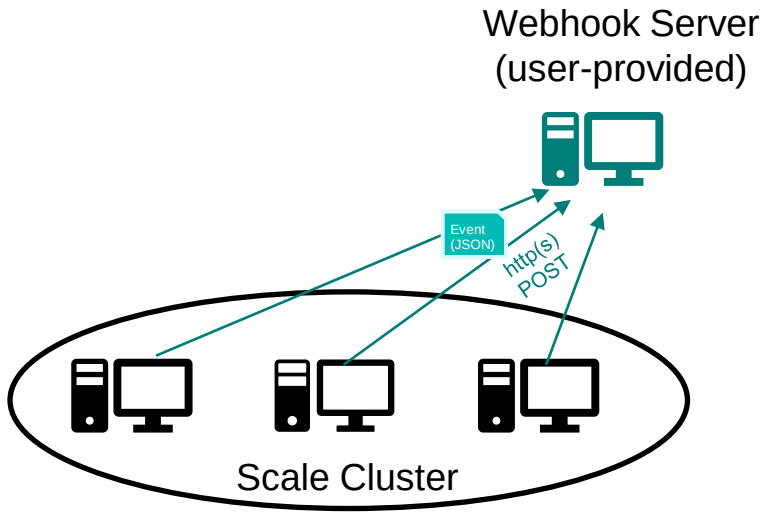

mmhealth event webhook

Users can register a webhook server (Http[s] URL) to receive mmhealth events

As mmhealth events are raised, they are posted to the registered webhook URLs

- Events are posted in JSON format. Multiple events will be bundled together
- Each cluster node is sending events directly to the registered webhook URLs
- Each Webhook URL has its own unique UUID, which is embedded in the POST header
- **A**dd, list, and remove webhooks using `mmhealth` command or REST API

https://www.ibm.com/docs/en/spectrum-scale/5.1.1?topic=reference-spectrum-scale-management-api-endpoints

Webhook Registration

- **Example:**
  - Add webhook URL via command line:

```
[root@scale01 ~]# mmhealth config webhook add http://10.21.38.26:8080/webhook/
Successfully connected to http://10.21.38.26:8080/webhook/
Webhook URL http://10.21.38.26:8080/webhook/ successfully linked to the health event monitoring system.
Webhook UUID is ef511240-a25d-489e-97b5-5befb996371c
```

  - Add webhook URL via REST API

```
[root@scale01 ~]# curl -k 'https://scale-gui:443/scalemgmt/v2/nodes/health/config/webhook/addEventWebhook'
-u admin:password -d '{ "webhookurl": "http://10.21.38.26:8080/webhook" }' --header 'Content-Type:
application/json’ --header 'accept:application/json' -X POST
```

  - List configured webhooks

```
[root@scale01 ~]# mmhealth config webhook list -Y
mmhealth:webhook:HEADER:version:reserved:reserved:url:uuid:status:
mmhealth:webhook:0:1:::https%3A//9.160.4.228%3A8090/webhook:34208ac8-ff14-4bd0-9ee2-b16a9c842c56:disabled:
mmhealth:webhook:0:1:::http%3A//10.21.38.26%3A8080/webhook/:ef511240-a25d-489e-97b5-5befb996371c:enabled:
```

Webhook Callout Details

- HTTP header example:

```
{'Remote-Addr': '9.114.205.79', 'Host': ‘scale01:9080', 'User-Agent': 'python-requests/2.25.1', 'Accept-Encoding': 'gzip,
deflate', 'Accept': 'application/json', 'Connection': 'keep-alive', 'Webhook-Uuid': '990f5044-4d2e-4a10-8127-a285bbc9e6b8',
'Content-Type': 'application/json', 'Content-Length': '1567'}
```

- Snippet of an event message:

```
{   "version": "1",
    "reportingController": "spectrum-scale",
    "reportingInstance": "afm-21.vmlocal",
    "events": [{
            "cause": "The GUI can connect to the pmcollector.",
            "component": "gui",
            "description": "The GUI checks the connection to the pmcollector.",
            "entity_name": "",
            "entity_type": "NODE",
            "event": "gui_pmcollector_connection_ok",
            "event_type": "STATE_CHANGE",
            "identifier": "afm-21.vmlocal",
            "internalComponent": "guipmcollector",
            "message": "The GUI can connect to the pmcollector that is running on afm-21.vmlocal using port 9980.",
            "node": "1",
            "requireUnique": true,
            "scope": "NODE",
            "severity": "INFO",
            "state": "HEALTHY",
            "time": "2023-03-21T15:17:32+01:00",
            "user_action": "N/A“
            ……
```

Cluster-wide event log using webhook

Implement a cluster-wide event log

- With minimal effort it is easily possible to implement a webhook server which writes each event to a log file

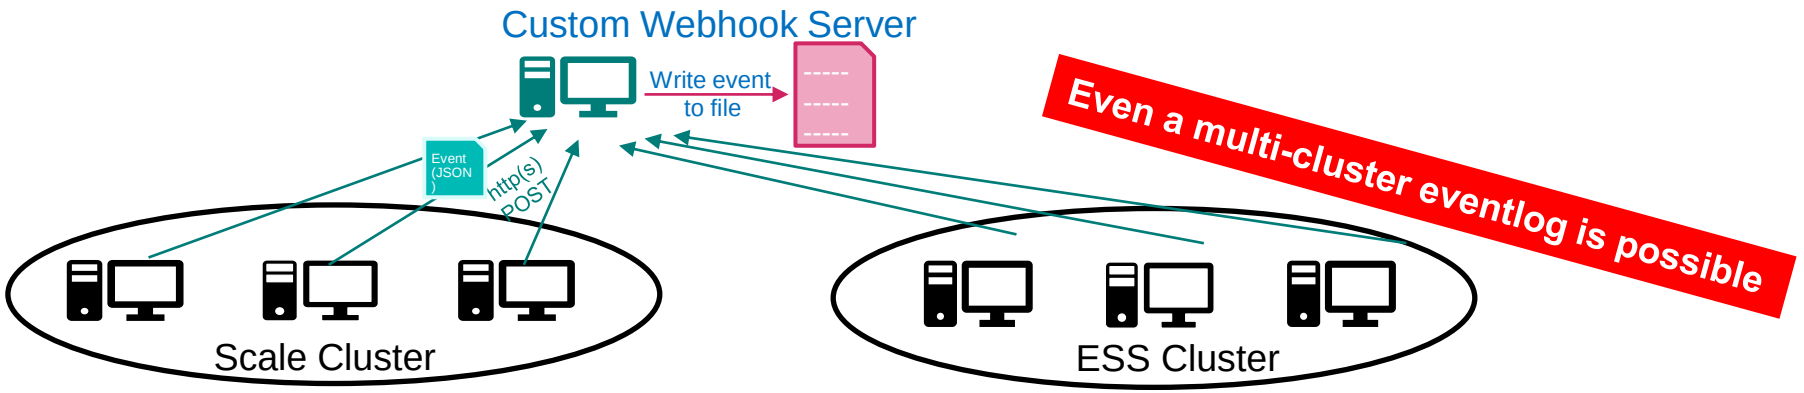

- Webhook server example code (Go, Rust, Python) is available in the Scale documentation: https://www.ibm.com/docs/en/spectrum-scale/5.1.5?topic=command-configuring-webhook-by-using-mmhealth
- Extend webhook server by adding a logEvent() function which writes out received events to an “eventlog” file

```
# Log all events to file (eventlog)
# @param node - which node is reporting the event
# @param ev - mmhealth event structure with event details
def logEvent(self, node, ev):
    with open("eventlog.txt","a") as f:
        f.write(f"{ev['time']} {node} {ev['component']} {ev['severity']} {ev['state']} {ev['event']} {ev['message']} \n" )
```

- Call the logEvent(..) function on every POST for each event

Webhook Outlook

Add option to send individual http requests for each event

- Some webhooks servers (e.g. Tivoli Netcool) do not support bulk events

Optimize event structure for tracking component states (Webhook version 2)
e.g. to automate ticket creation / deletion for health events

- Introduce a “full identifier” to be able to map corresponding events
- Ensure essential fields are filled to remove the need for parsing the event message
  - Entity_name field populated with visible mmhealth enity names (e.g. ens0)
- Consistent use of scope field (Node vs cluster)

Improve security and reliability

- Add bearer token argument which is send in authentication header
- Improve retry logic when webhook server is temporarily not available

MMHEALTH ENHANCEMENTS

MMHEALTH command improvements

- Component *ENTITY* status agregated to a cluster level

  # mmhealth cluster show **COMPONENT entityName** -v -Y | grep **ClusterSummary**

- **`mmhealth event show <eventname> -Y`**

- New Flag: **--**extended **|** -X to return the relevant *UNHEALTHY* **event details**

  # mmhealth node show –h

  Usage:
  mmhealth node show [component [entity]
  [-N {Node[,Node..] | NodeFile | NodeClass}]
  [-X | --extended]
  [-Y] [--verbose] [--unhealthy] [--no-notices] [--color | --nocolor]
  [--resync]

MMHEALTH command improvements

mmhealth node show Filesystem **-X**

Node name: scale-11.vmlocal

Component	Status	Status Change	Reasons & Notices
FILESYSTEM	TIPS	22 hours ago	ill_unbalanced_fs(local_FS2)
cesSharedRoot	HEALTHY	6 days ago	-
local_FS1	HEALTHY	1 day ago	-
local_FS2	TIPS	22 hours ago	**ill_unbalanced_fs(local_FS2)**
remote_fs	HEALTHY	1 day ago	-

Event	Parameter	Severity	Active Since	Event Message
ill_unbalanced_fs	local_FS2	TIP	22 hours ago	The filesystem local_FS2 is not properly balanced.

ill_unbalanced_fs tip event signals local_FS2 to show TIPS

Message: **The filesystem local_FS2 is not properly balanced.**
Cause: The 'mmfsadm eventsExporter get fs <filesystem>' command reports that the filesystem is no longer properly balanced.
Description: A configuration change is causing the filesystem to no longer being properly balanced.
User Action: **Run the 'mmrestripefs' command against the filesystem.**

MMHEALTH command improvements

– 2nd Example

```
[~]# mmhealth node show -X

Node name:      node31.localnet.com
Node status:    DEGRADED
Status Change:  Now

Component       Status        Status Change     Reasons & Notices
------------------------------------------------------------------------
GPFS            HEALTHY       Now               -
...
PERFMON         FAILED        Now               pmsensors_down
THRESHOLD       HEALTHY       11 days ago       -

pmsensors_down error event signals PERFMON to show FAILED
   Message:     The pmsensors service should be started, but is stopped.
   Cause:       The node has the perfmon designation in the 'mmlscluster' command, but the
                pmsensors service is not running.
   Description: The performance monitor sensors are down.
   User Action: Start the service by using the 'service' or 'systemctl' command. Otherwise,
                remove the perfmon designation by using the 'mmchnode' command.
```

– Also an automatic text wrap was added to mmhealth for a nicer CLI experience!

AFM MONITORING IN MMHEALTH

AFM Monitoring with MMHEALTH

Active File Management (AFM) lets Spectrum Scale extend over geographic distances

- Tolerates unreliable, high-latency networks (like a WAN).
- Caches copies of data from a remote file system into the local Spectrum Scale cluster.
- Cached files have the same read and write performance as other local files.
- Default AFM configuration

  # mmlsconfig afmHardMemThreshold

  5368709120
- https://www.ibm.com/support/pages/sites/default/files/inline-files/AFMDeepDive.pdf

New default threshold rule **AFMInQueue_Rule**

- observes the AFM in queue memory usage on Gateway nodes and raise events upon reaching or exceeding the specified limits.
  - **Warning event - 80 % of afmHardMemThreshold**
  - **Error event - 90 % of afmHardMemThreshold**
- Based on the perfmon metric **gpfs_afm_used_q_memory**
- Current limitation: the same memory config is expected for all Gateway nodes.
- Inactive as long as AFM sensors not enabled

AFM Monitoring with MMHEALTH

mmlscluster

GPFS cluster information
==========

GPFS cluster name: scale-cluster-2.vmlocal

Node	Daemon node name	IP address	Admin node name	Designation
1	scale-21.vmlocal	10.0.100.21	scale-21.vmlocal	quorum-manager-perfmon
2	scale-23.vmlocal	10.0.100.23	scale-23.vmlocal	quorum-manager-perfmon
3	scale-24.vmlocal	10.0.100.24	scale-24.vmlocal	quorum-manager-perfmon
4	scale-22.vmlocal	10.0.100.22	scale-22.vmlocal	perfmon
5	scale-25.vmlocal	10.0.100.25	scale-25.vmlocal	manager-**gateway**-perfmon
6	scale-26.vmlocal	10.0.100.26	scale-26.vmlocal	manager-**gateway**-perfmon
7	scale-27.vmlocal	10.0.100.27	scale-27.vmlocal	manager-**gateway**-perfmon

mmhealth node show -N scale-25

Node name: scale-25.vmlocal
Node status: TIPS
Status Change: 3 days ago

Component	Status	Status Change	Reasons & Notices
GPFS	TIPS	3 days ago	total_memory_small
NETWORK	HEALTHY	3 days ago	-
FILESYSTEM	TIPS	1 day ago	ill_unbalanced_fs(storage_fs)
DISK	HEALTHY	3 days ago	-
AFM	**TIPS**	**1 day ago**	**afm_sensors_inactive(GPFSAFM, GPFSAFMFS, GPFSAFMFSET)**
CES	HEALTHY	3 days ago	-
PERFMON	HEALTHY	3 days ago	-
THRESHOLD	HEALTHY	3 days ago	-

AFM Monitoring with MMHEALTH

mmhealth thresholds list

active_thresholds_monitor: scale-21.vmlocal

Threshold Rules

rule_name	metric	error	warn	direction	filterBy	groupBy	sensitivity
MemFree_Rule	MemoryAvailable_percent	None	5.0	low		node	300-min
DataCapUtil_Rule	DataPool_capUtil	90.0	80.0	high		gpfs_cluster_name,gpfs_fs_name,gpfs_diskpool_name	300
MetaDataCapUtil_Rule	MetaDataPool_capUtil	90.0	80.0	high		gpfs_cluster_name,gpfs_fs_name,gpfs_diskpool_name	300
InodeCapUtil_Rule	Fileset_inode	90.0	80.0	high		gpfs_cluster_name,gpfs_fs_name,gpfs_fset_name	300
SMBConnPerNode_Rule	current_connections	3000	None	high		node	300
SMBConnTotal_Rule	current_connections	20000	None	high			300
AFMInQueue_Rule	**AFMInQueueMemory_percent**	**90.0**	**80.0**	**high**		**node**	**300**

 # mmhealth thresholds list -v -Y | grep AFMInQueue_Rule | grep state

mmhealth_thresholds:THRESHOLD_RULE:0:1::AFMInQueue_Rule:**state**:***inactive***:

AFM Monitoring with MMHEALTH

➢ Enable AFM sensors

```
# mmcrnodeclass afmGateways -N scale-25,scale-26,scale-27
# mmperfmon config update GPFSAFM.period=10 GPFSAFMFS.period=10 GPFSAFMFSET.period=10
# mmperfmon config update GPFSAFM.restrict=afmGateways GPFSAFMFS.restrict=afmGateways GPFSAFMFSET.restrict=afmGateways
```

More info: https://www.ibm.com/docs/en/spectrum-scale/5.1.7?topic=gui-activate-afm-performance-monitoring-sensors

➢ Check the threshold rule AFMInQueue_Rule is active now

```
# mmhealth node eventlog
Node name:      scale-21.vmlocal
```

Timestamp	Event Name	Severity	Details
2023-03-18 11:20:11.453073 CET	ok_unbalanced_fs	INFO	The filesystem cash_fs is properly balanced.
2023-03-18 11:20:11.462619 CET	ok_replicated_fs	INFO	The filesystem cash_fs is properly replicated.
2023-03-18 11:20:11.470317 CET	ok_exposed_fs	INFO	The filesystem cash_fs has no data exposure risk.
2023-03-18 12:26:41.535200 CET	pool-data_normal	INFO	The pool system of file system cash_fs has reached a normal data level.
2023-03-18 12:26:41.543806 CET	inode_normal	INFO	The inode usage of fileset consistGroupB in file system cash_fs reached a normal level.
2023-03-18 12:26:41.552321 CET	inode_normal	INFO	The inode usage of fileset root in file system cash_fs reached a normal level.
2023-03-18 12:26:41.558731 CET	pool-metadata_normal	NFO	The pool system of file system cash_fs has reached a normal metadata level.
2023-03-19 14:18:39.796895 CET	**activate_afm_inqueue_rule**	INFO	**Detected AFM Gateway node scale-25.vmlocal. Enabled AFM In Queue Memory rule for threshold monitoring.**

AFM Monitoring with MMHEALTH

```
# mmhealth node show threshold -v -N scale-25

Node name:      scale-25.vmlocal
Component          Status        Status Change           Reasons & Notices
------------------------------------------------------------------------------
THRESHOLD          HEALTHY       2023-03-19 14:18:30     -
  AFMInQueue_Rule  HEALTHY       2023-03-19 14:21:33     -
Event              Parameter     Severity    Active Since        Event Message
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------
thresholds_normal  AFMInQueue_Rule  INFO        2023-03-19 14:21:33  The value of AFMInQueueMemory_percent defined in
                                                                      AFMInQueue_Rule for component
                                                                      AFMInQueue_Rule/scale-25 reached a normal level.
```

- **In case of exceeding the rule boundaries**

  `"thresholds_warn"` or `"thresholds_error„` event will be raised

  suggestion for an **user action**:

  The AFM Used Queue memory has reached a critical threshold value. Check and fix any AFM gateway network or load balancing issues. Increase the afmHardMemThreshold value to be 40-50% of available memory. Add more gateway nodes to handle the AFM workload.

AFM Monitoring with Grafana

https://github.com/IBM/ibm-spectrum-scale-bridge-for-grafana/wiki/Monitoring-AFM

– AFM relationship by using the NFS protocol

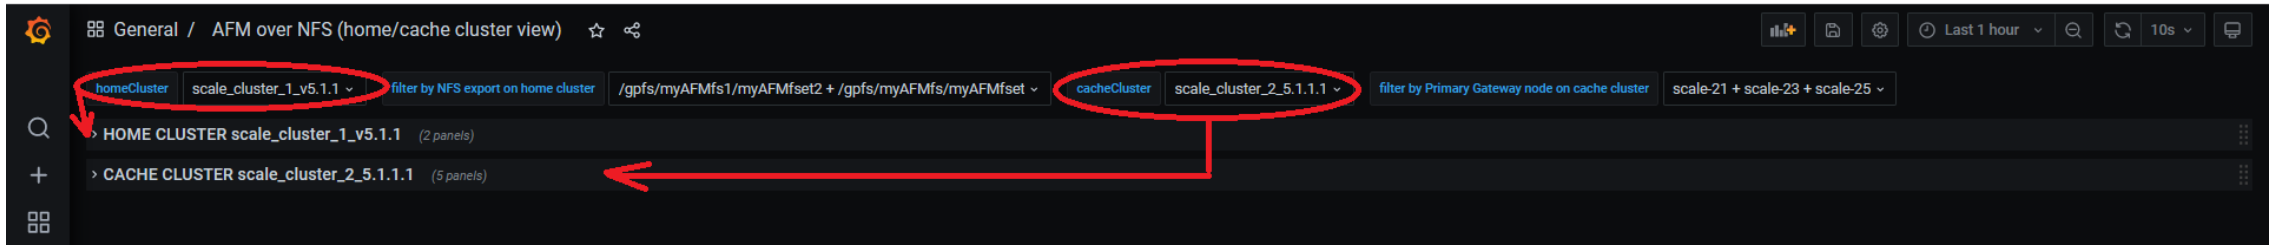

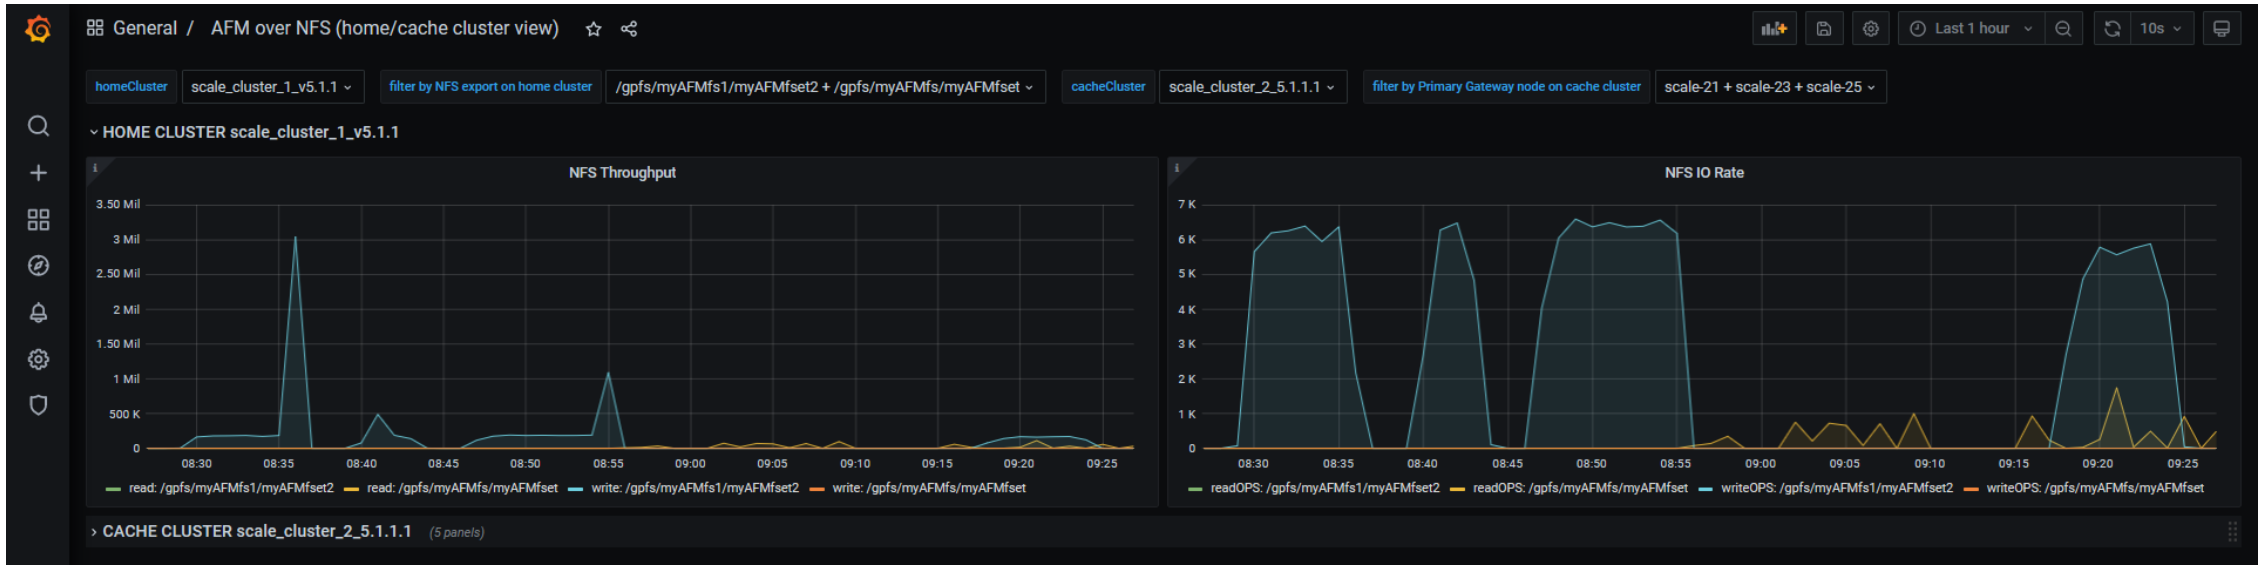

AFM Monitoring with Grafana

https://github.com/IBM/ibm-spectrum-scale-bridge-for-grafana/wiki/Monitoring-AFM

– AFM relationship by using the NFS protocol

Thank you for using
IBM Spectrum Scale!