IBM Spectrum Scale

Performance Data analytics with Grafana dashboards

Helene Wassmann

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. Reference Materials

Outline

1. **What is Grafana?**
    a. **Evolution & Key concepts**
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. Reference Materials

What is Grafana?

IBM Spectrum Scale

“The analytics platform for all your metrics
Grafana allows you to query, visualize, alert on **and** understand your metrics no matter where they are stored.
Create, explore, and share dashboards with your team and foster a **data driven culture**.”

https://grafana.com/grafana

Grafana features highlights

IBM Spectrum Scale

Application

- pretty easy to install and configure
- OS: Linux, Mac, Windows, Docker or building from source
- powerful visualization capabilities
- wide array of customization options
- large community of users and active contributors
- hundreds of dashboards and plugins in the official library

Data Source

- built-in integration with Graphite, Prometheus, InfluxDB, openTSDB, Elasticsearch, MySQL, PostgreSQL, and many others...
- for each data source specific query editor
- extend Grafana by writing your own plugins

Dashboard

- dynamic & reusable dashboards
- generic dashboards trough templating
- comprehensive charts
- slice and dice data in any way
- mix different data sources in the same dashboard
- Sharing data & dashboards across the team in many ways (direct link, snapshot, export & import)

Row/Panel

- easy style switch (graph, singlestat, table, heatmap and free text)
- flexible data transformation (time series to rows or time series to columns)
- mix different data sources in the same graph
- query Inspector
- alias patterns

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. Reference Materials

Grafana integration with IBM Spectrum Scale

Grafana

- talks to backends over REST HTTP API

IBM Spectrum Scale **BRIDGE** for Grafana

- standalone Python application
- openTSDB data exchange format
- full set of IBM Spectrum Scale supported metrics
- communicates with Grafana via port 4242 (default by openTSDB)

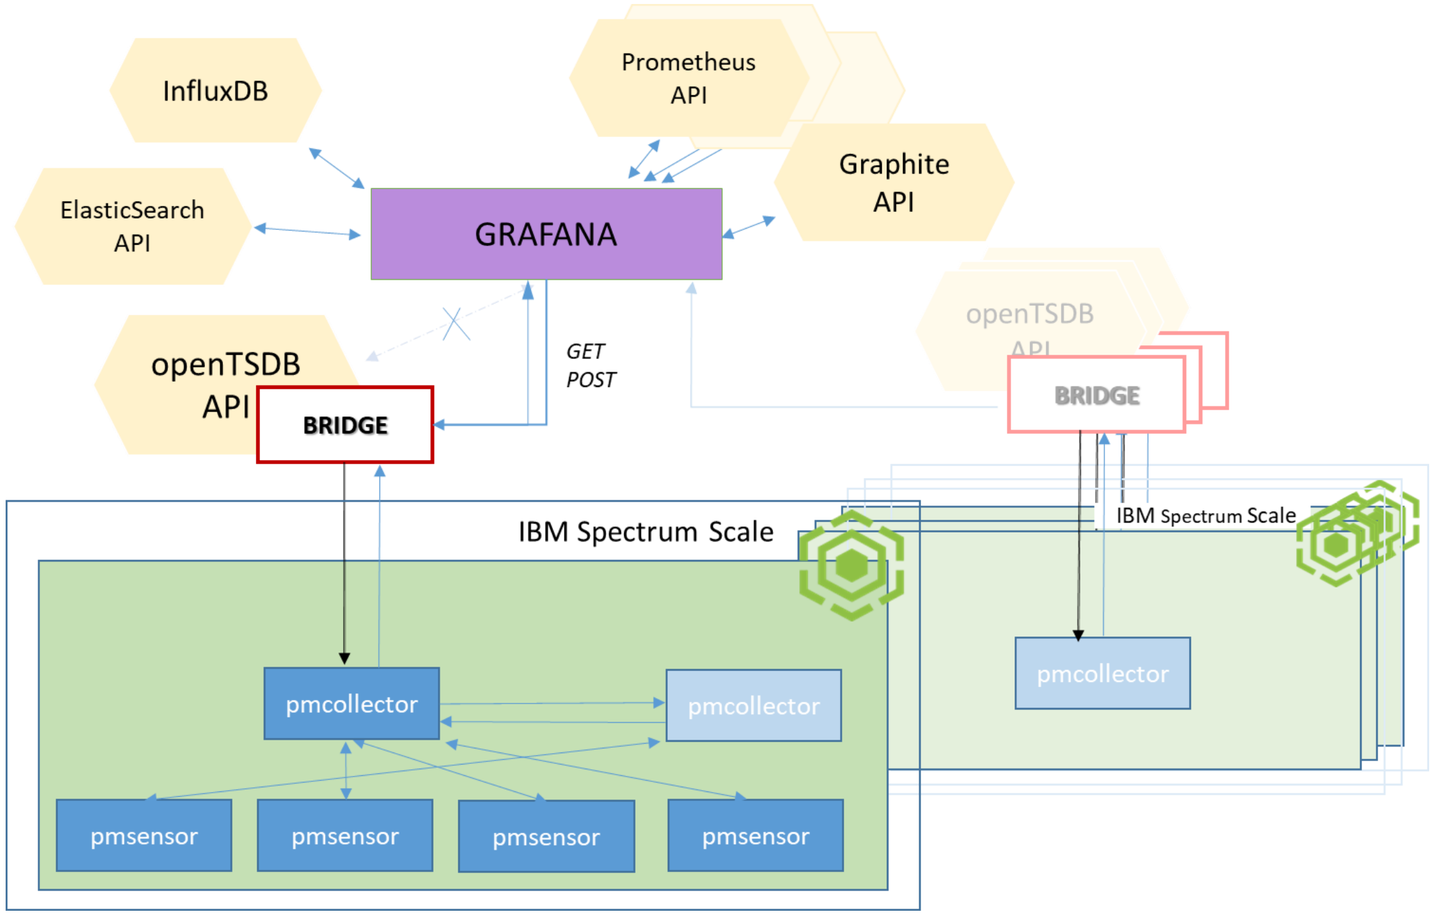

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. **IBM Spectrum Scale Bridge for Grafana updates**
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. Reference Materials

IBM Spectrum Scale **Bridge** updates

IBM Spectrum Scale Bridge version 4.0

NEW !

- will be available at the beginning of April 2019
- tested with Grafana 6.0, Python 2.7, 3.6, CherryPy 18.01
- tested with IBM Spectrum Scale 4.2.3, 5.0.2, 5.0.3

NOTE: CherryPy has dropped support for Python 2 with version 18.0.0

Bridge source code changes

- most improvements done to the selection/setting query interval
  - use metric data polling interval, if the downsampling is disabled explicitly (!)
  - allow MIN, MAX, SUM, AVG calculation for the downsampling

NEW !

New bundle of “Example Dashboards”

- download for free
- will be published on developerWorks , together with the bridge v.4
- organized in sub groups: HOWTO’s, protocols, cloud…

NEW !

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. **New bundle ‘Example Dashboards’**
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. Reference Materials

New bundle “**Example Dashboards**”

IBM Spectrum Scale

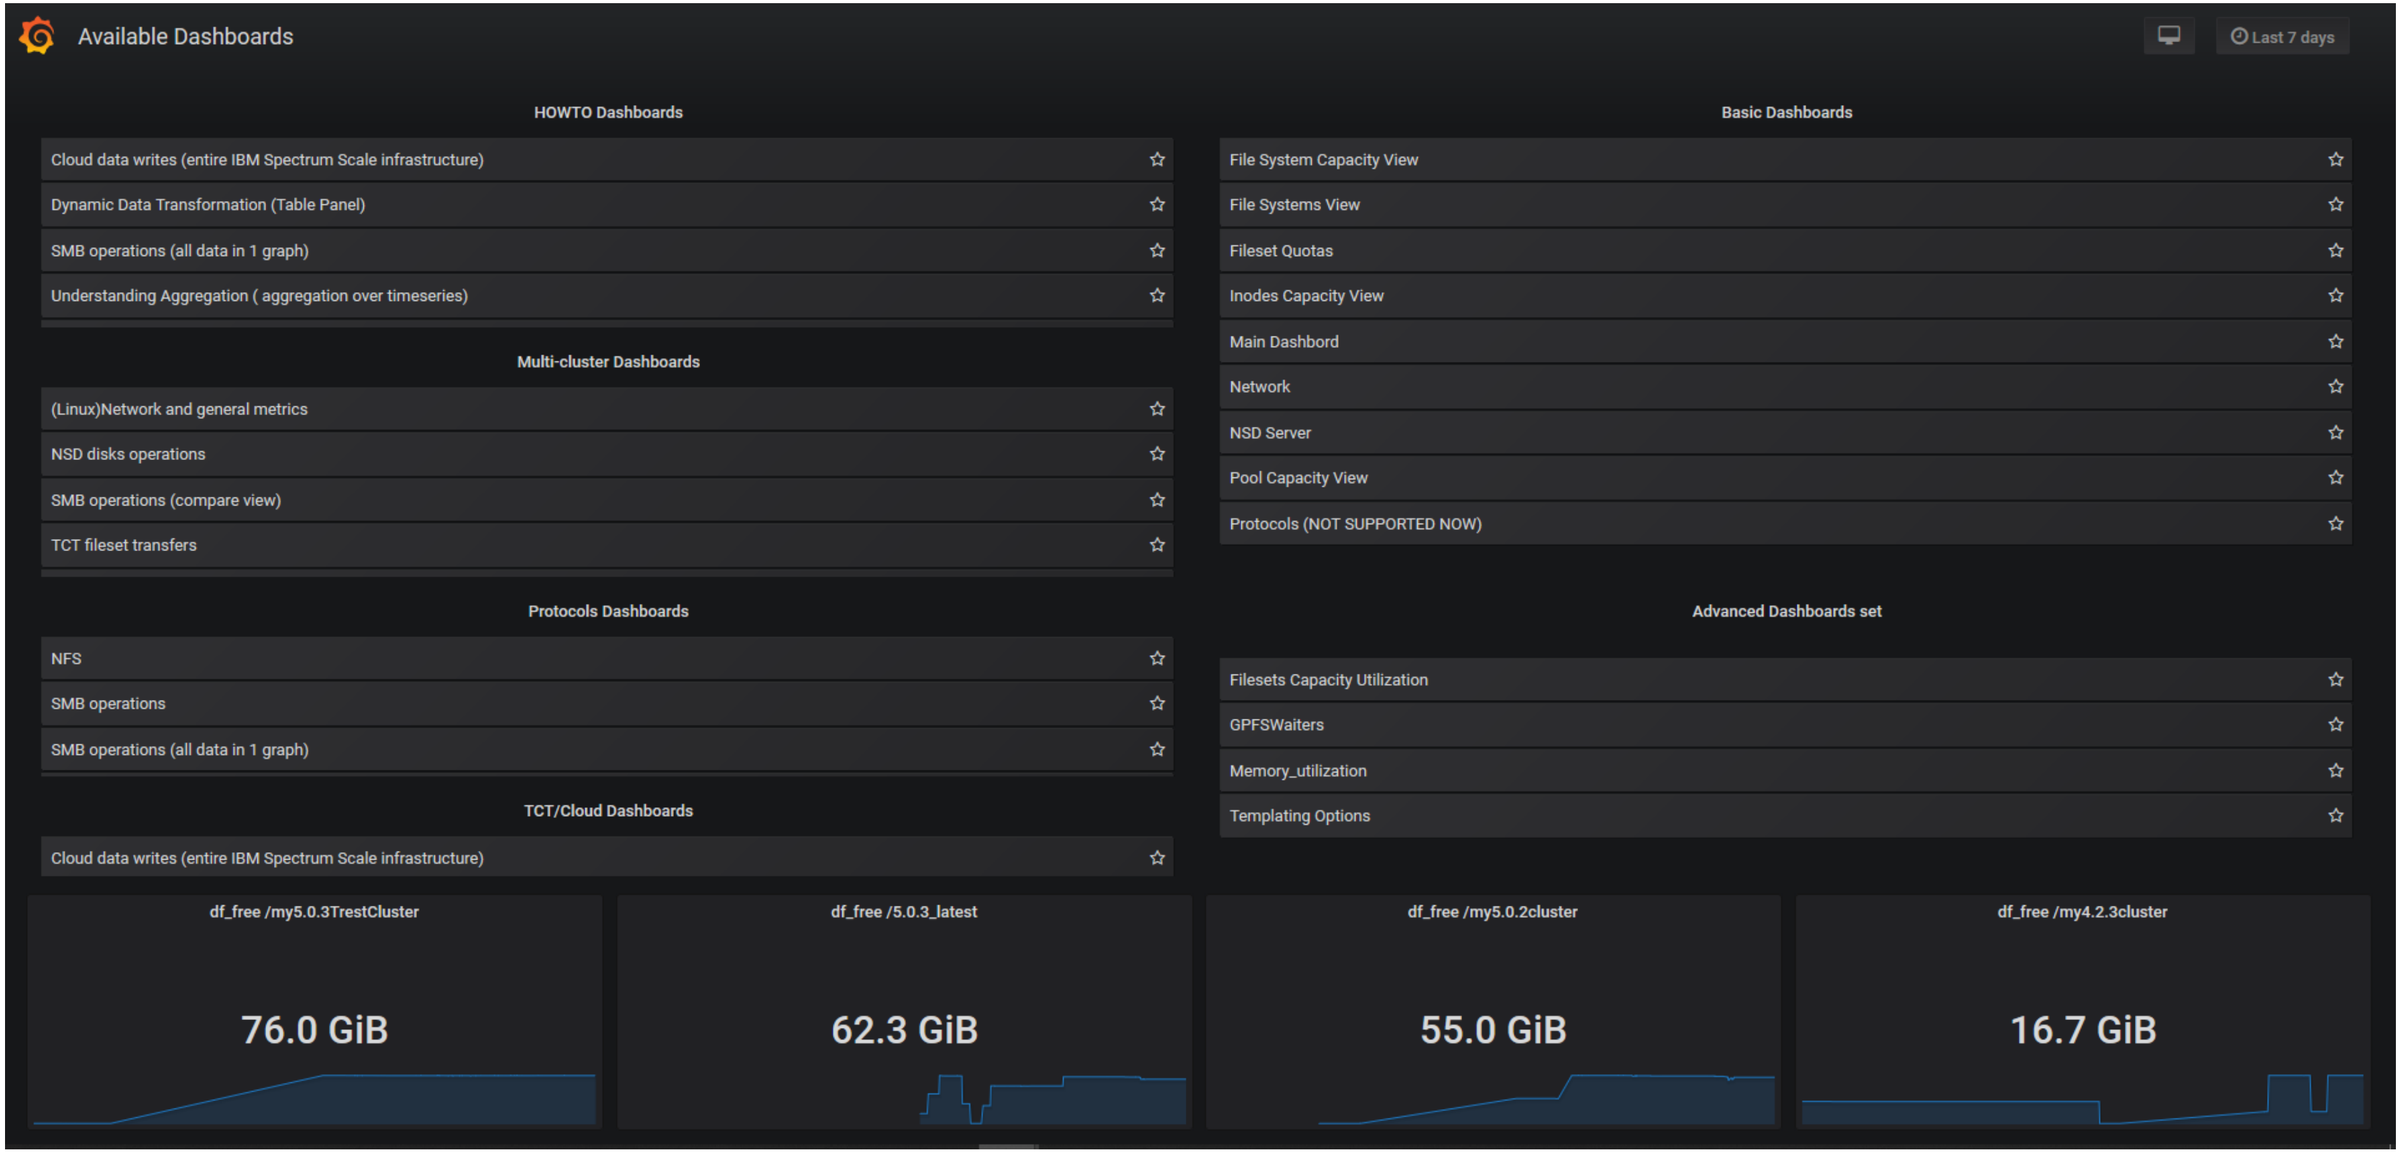

New bundle “**Example Dashboards**”

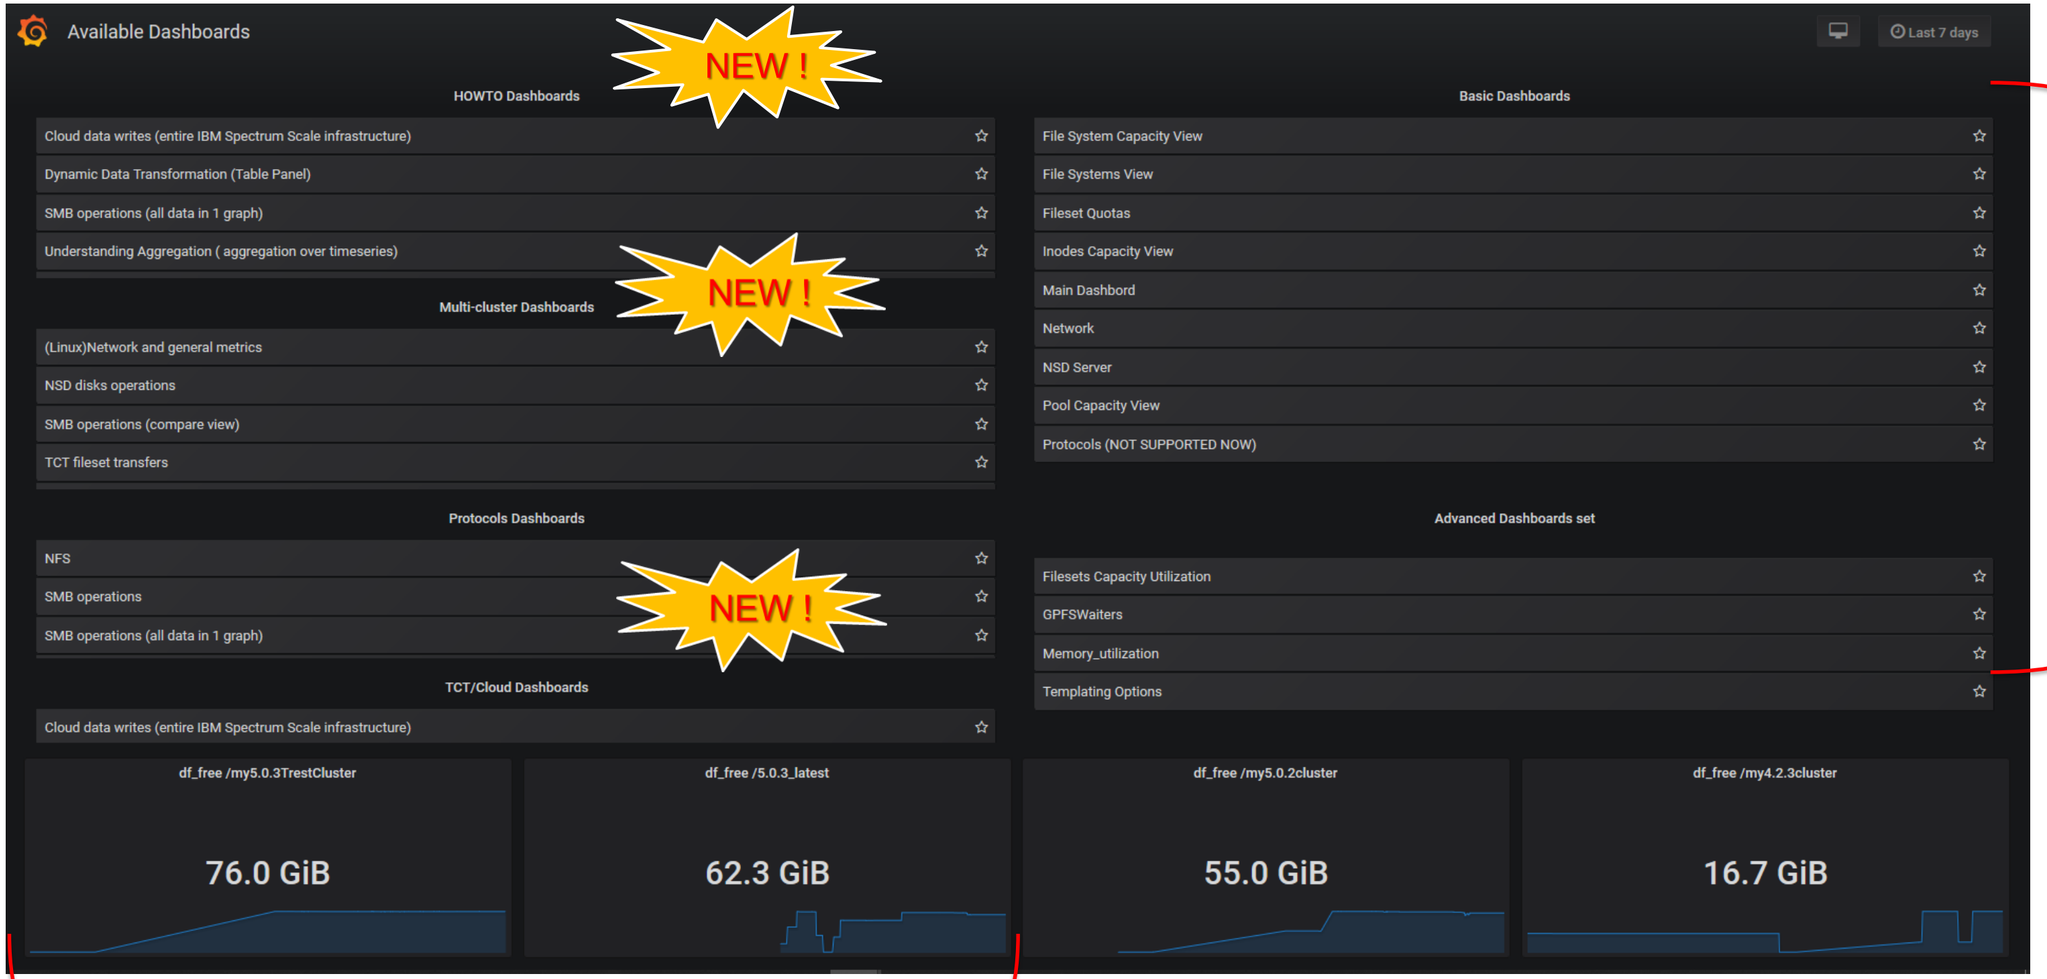

available from the previous Bridge releases

New dashboards download package will be available with the Bridge v.4

New bundle “**Example Dashboards**”

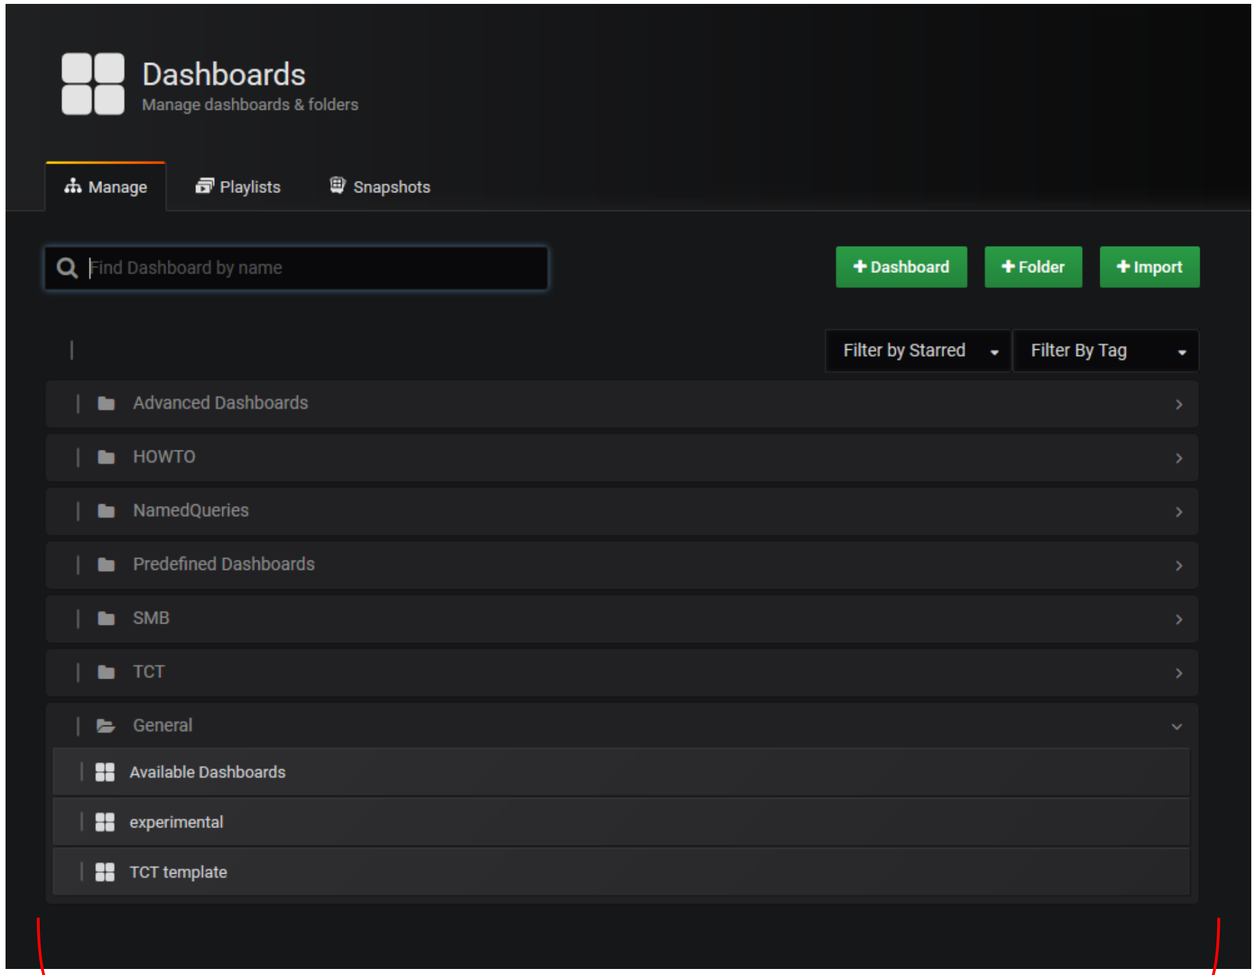

folder structure

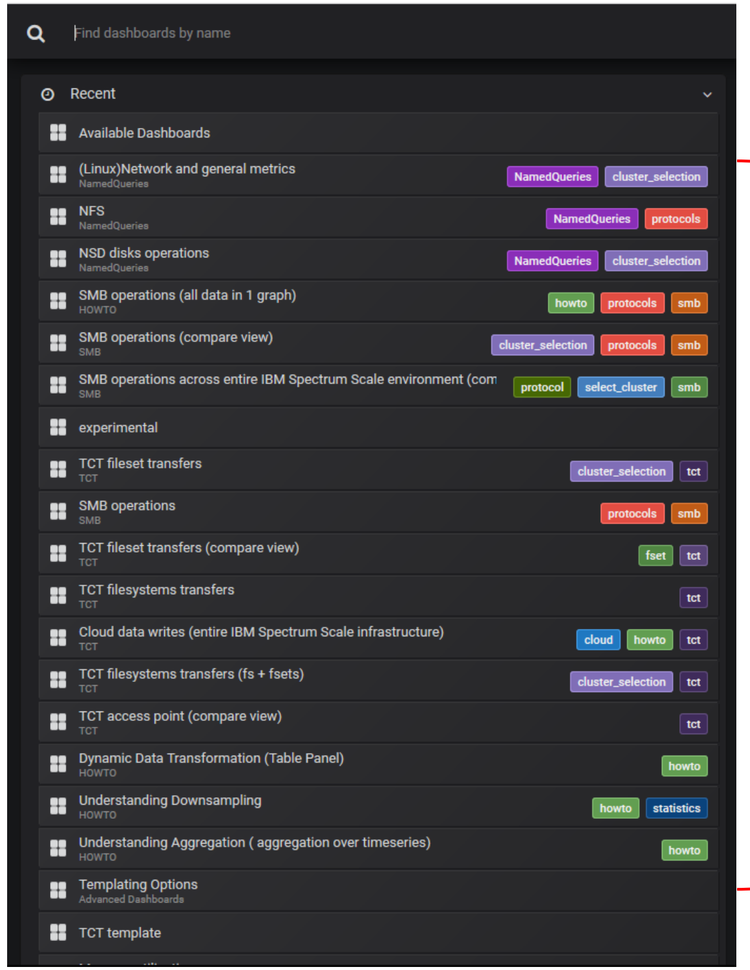

search by *TAG* or dashboard name

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. Reference Materials

Getting Started with “**HOWTO Dashboards**”

IBM Spectrum Scale

“Learning by Doing”

- install Grafana & IBM Spectrum Scale bridge for Grafana
  - ➢ follow bridge installation instructions on the developerWorks
- add your cluster to Grafana monitoring sources
- download and import example dashboards (**HOWTO**’s recommended)
- try out Grafana & bridge features interactively
- modify and design your own dashboards

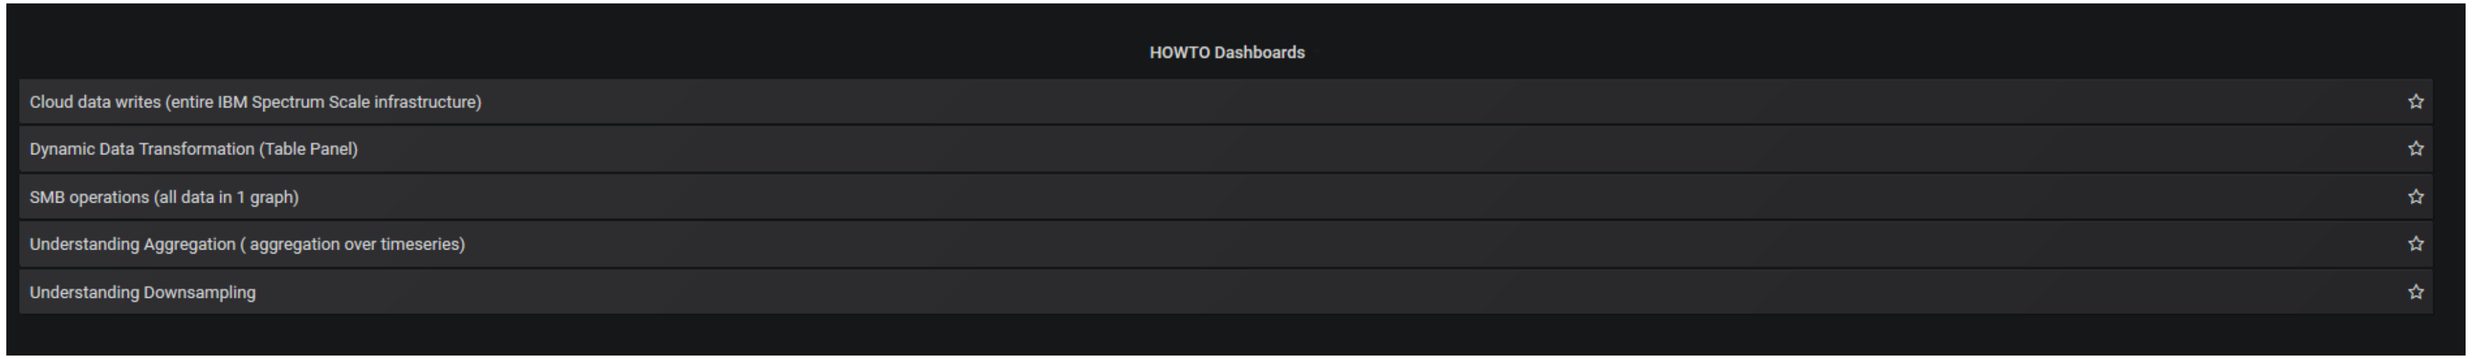

Getting Started with “**HOWTO Dashboards**”

IBM Spectrum Scale

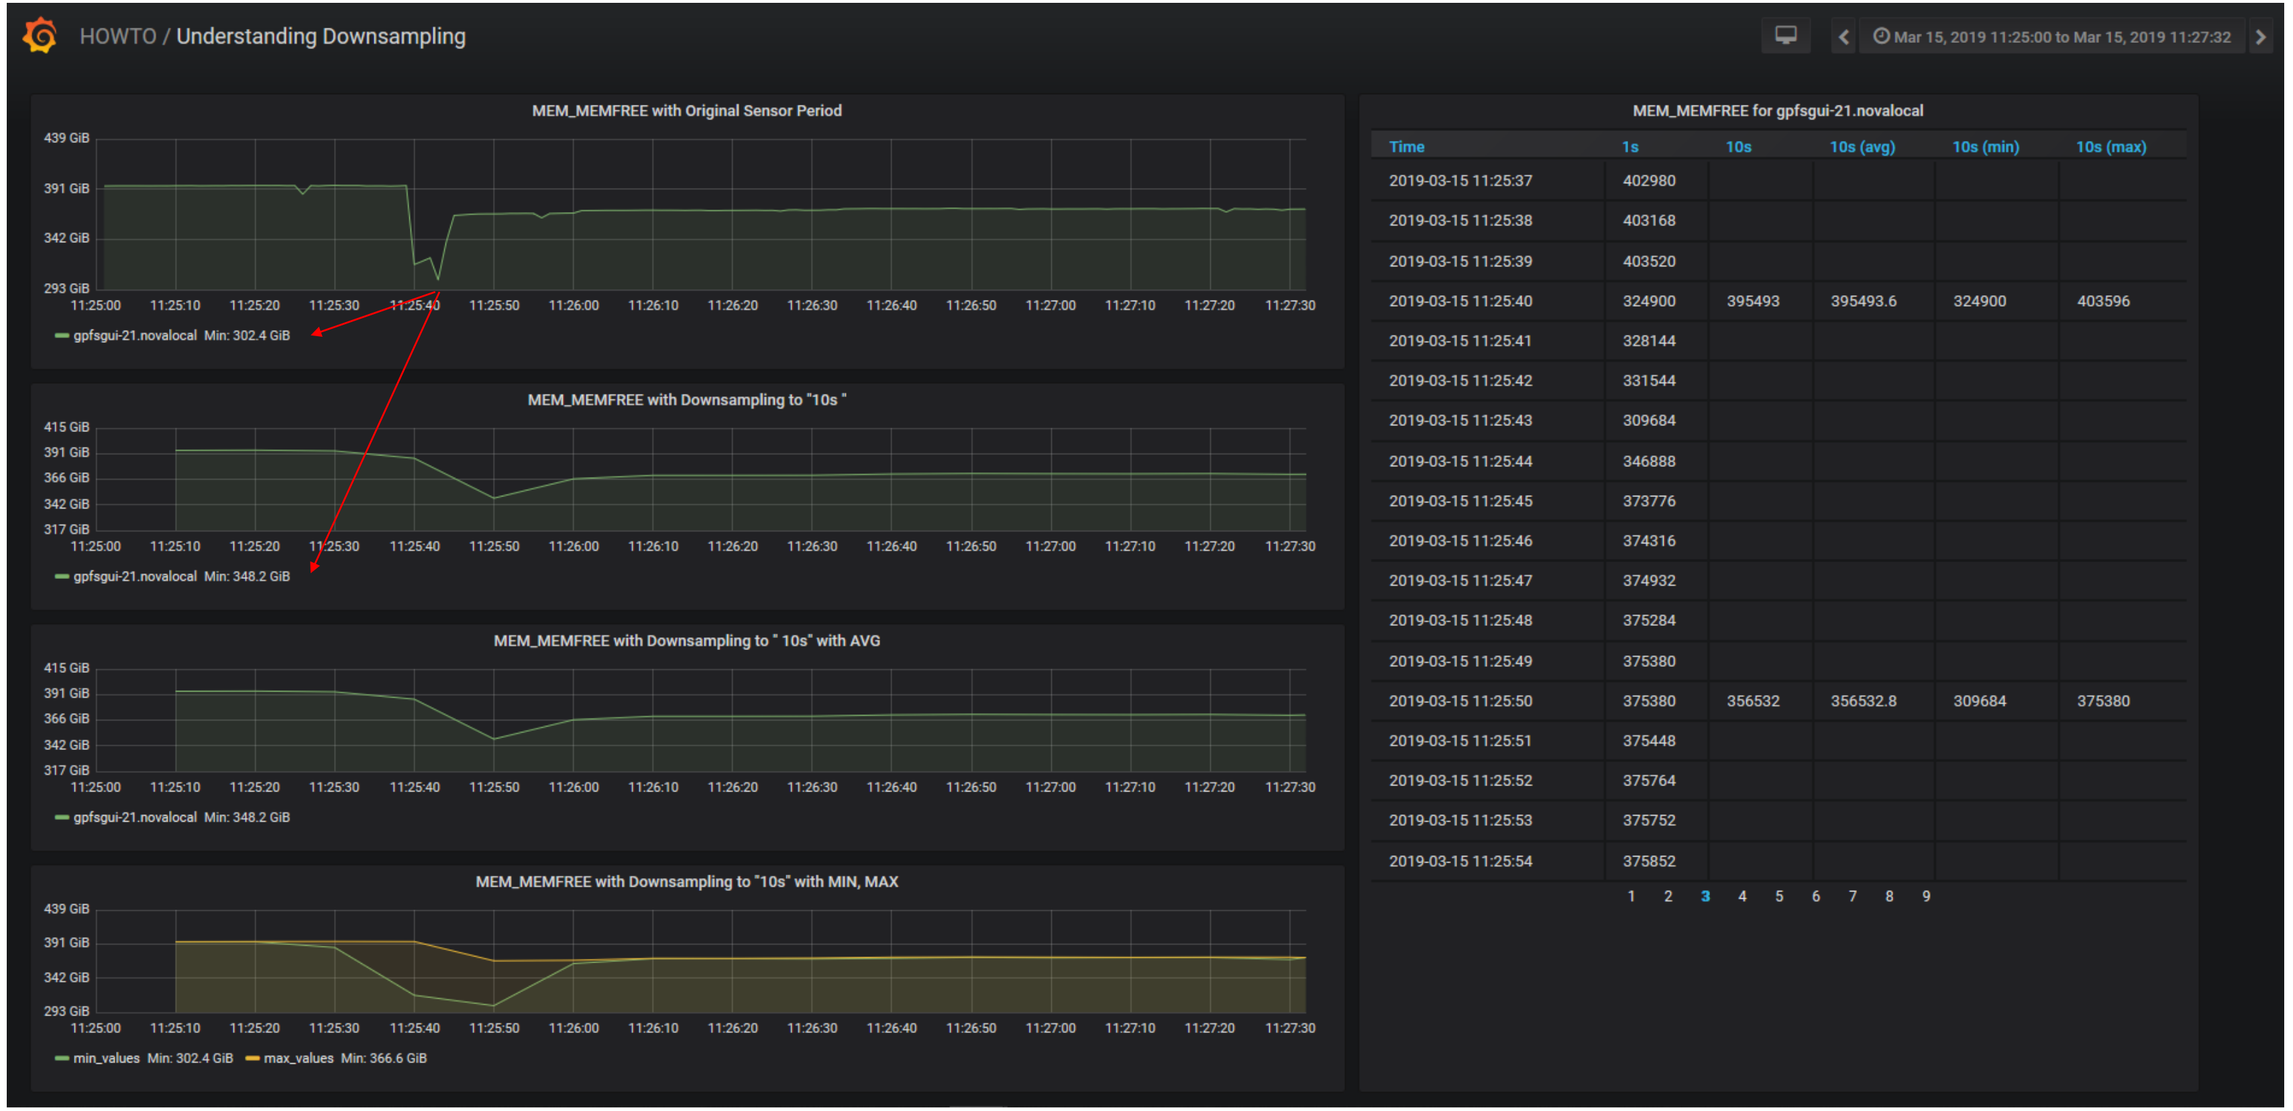

Getting Started with “**HOWTO Dashboards**”

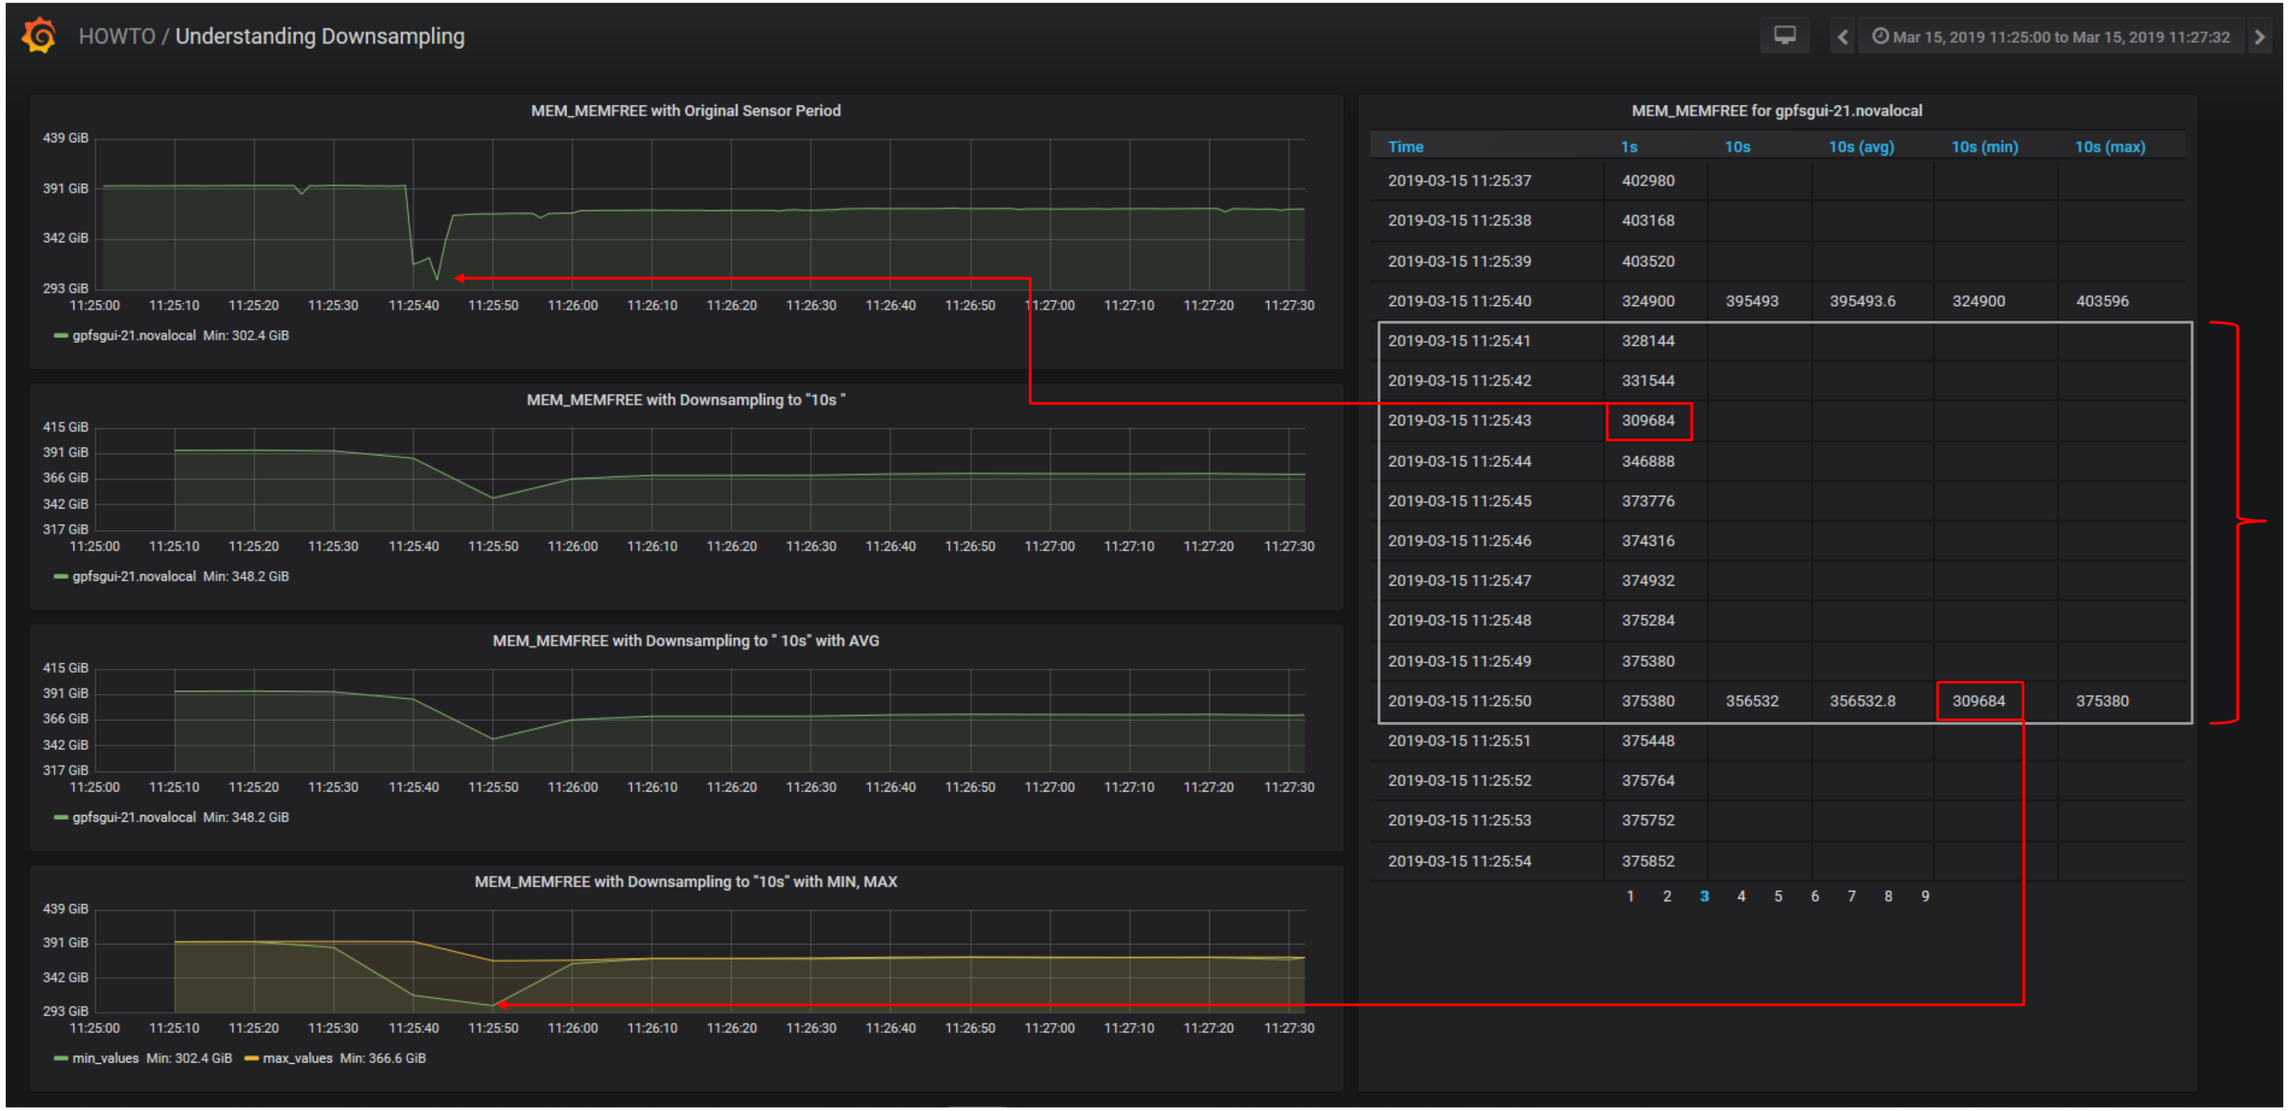

interval (10s)

Getting Started with “**HOWTO Dashboards**”

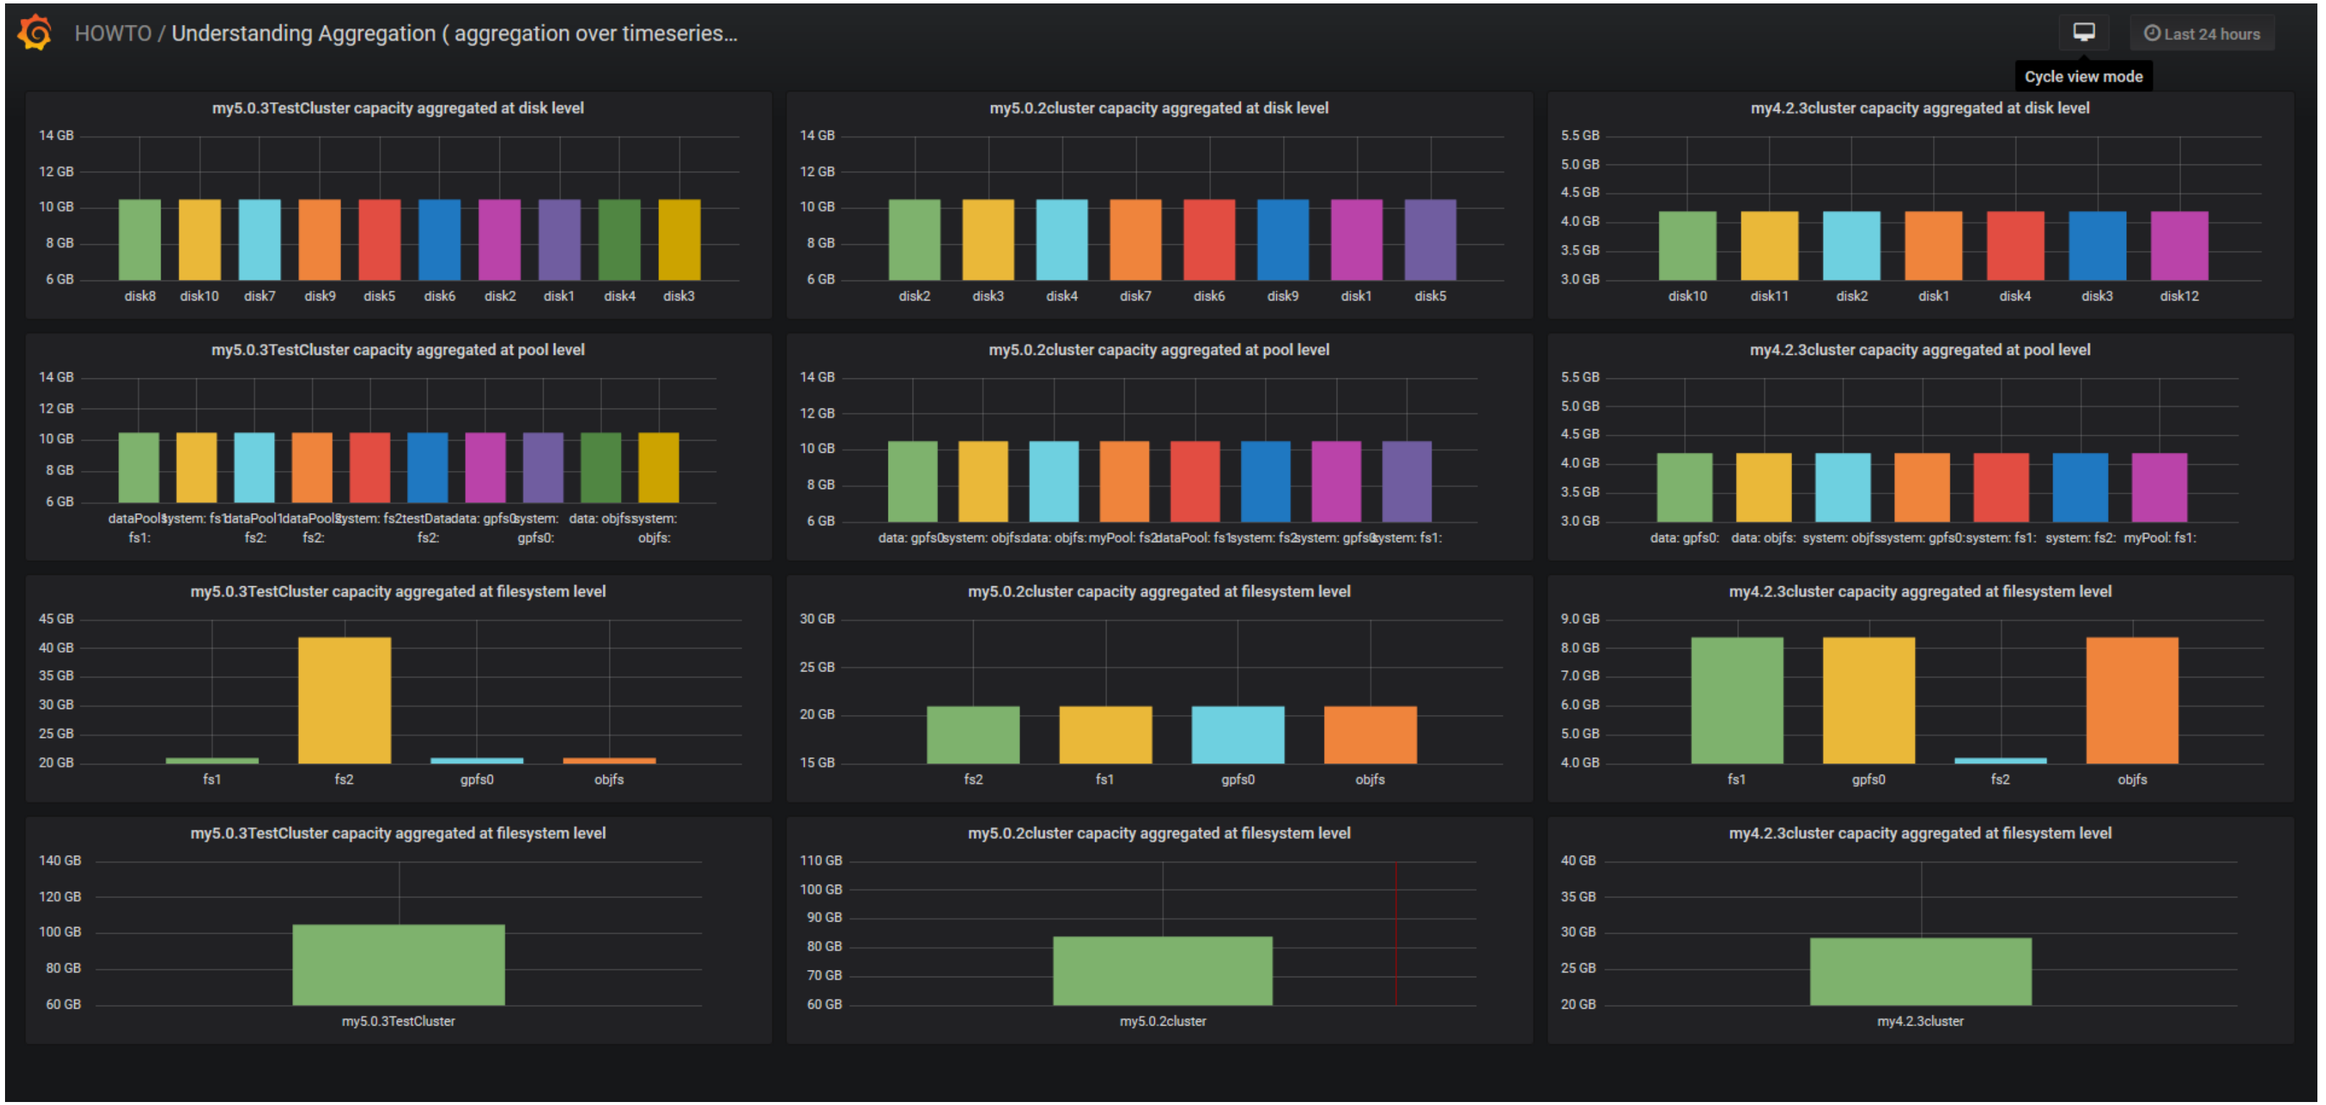

Getting Started with “**HOWTO Dashboards**”

the number of time series reduced from 25 to 3

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. Reference Materials

Use case “**SMB operations**”

IBM Spectrum Scale

Use case “**SMB operations**”

Use case “**SMB operations**”

IBM Spectrum Scale

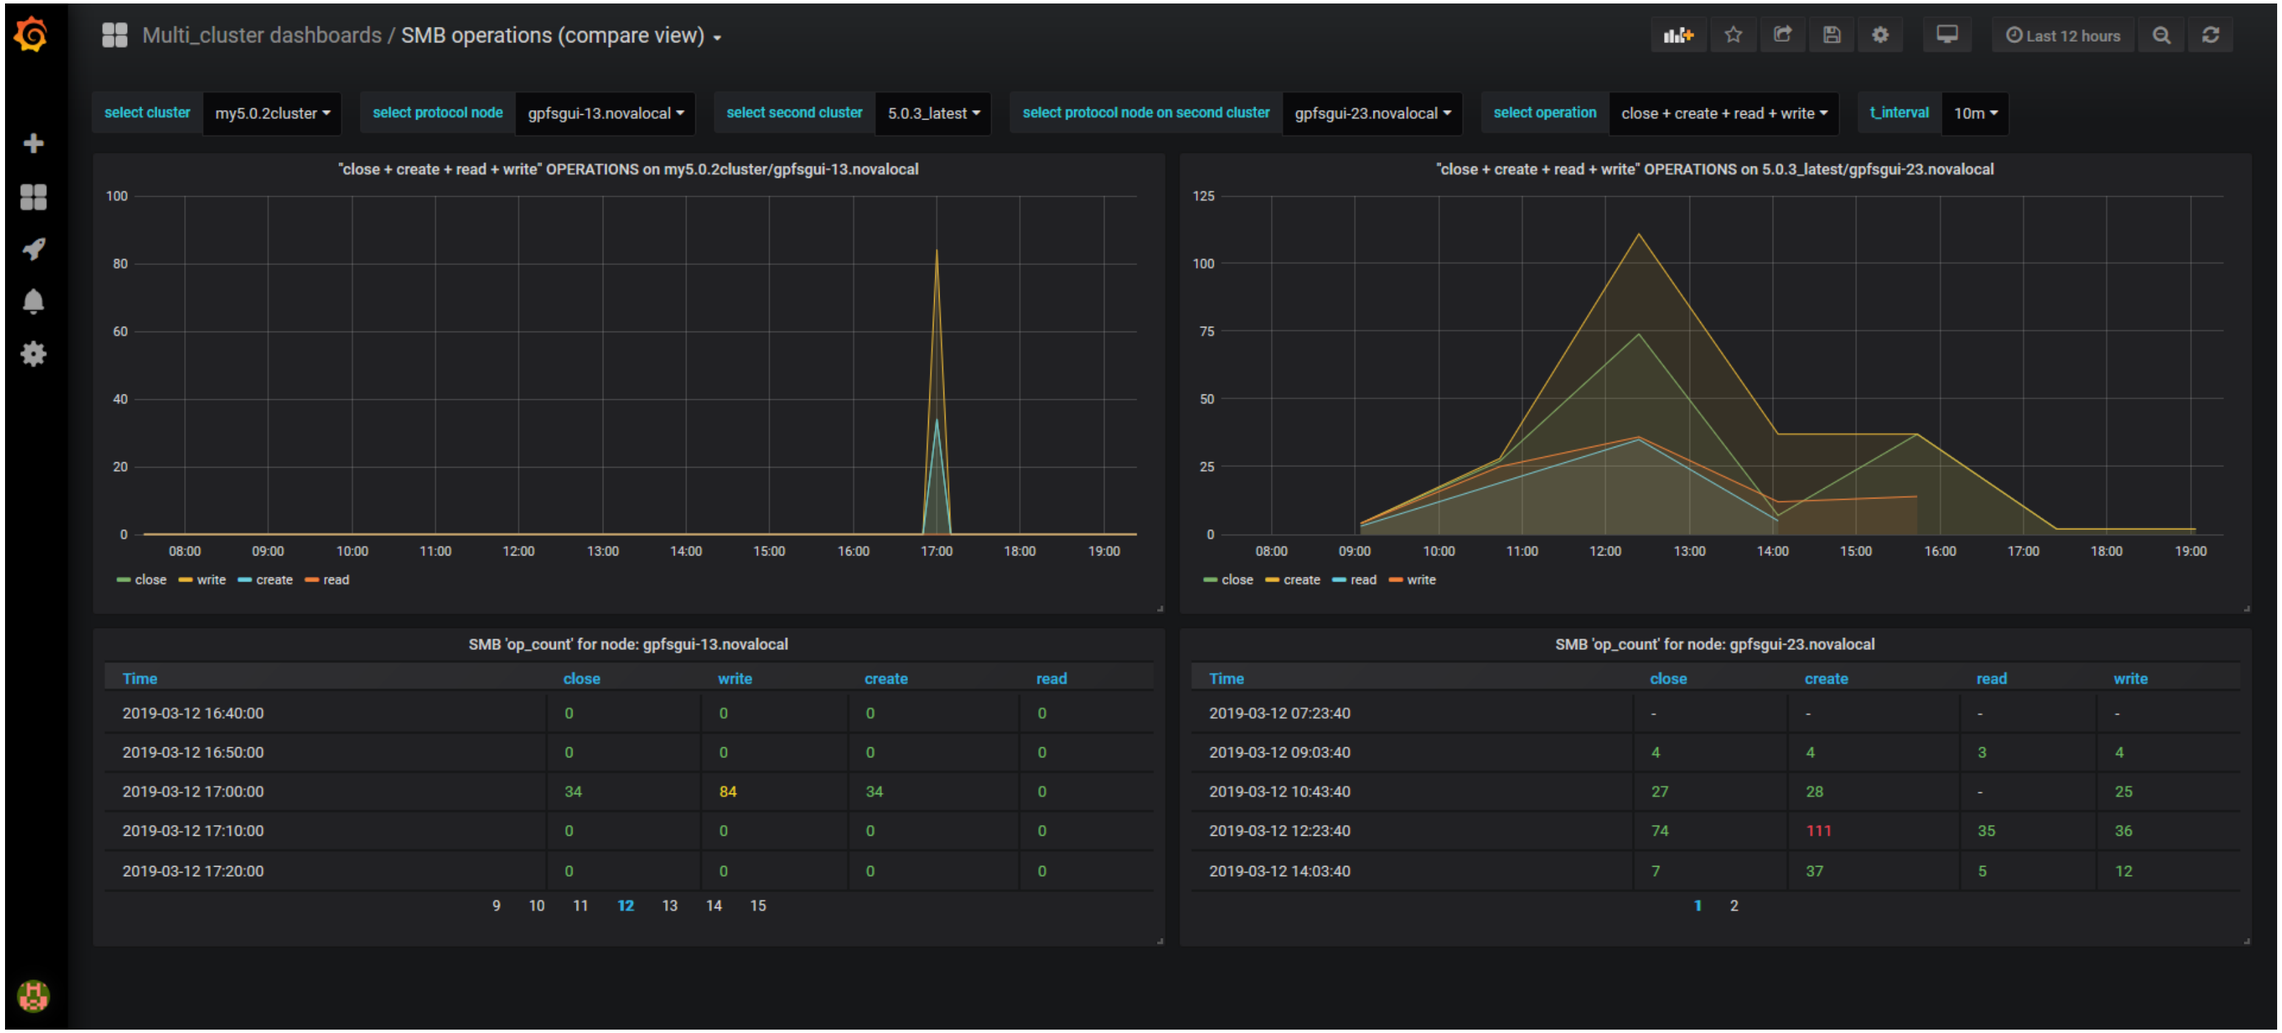

Use case “**SMB operations**”

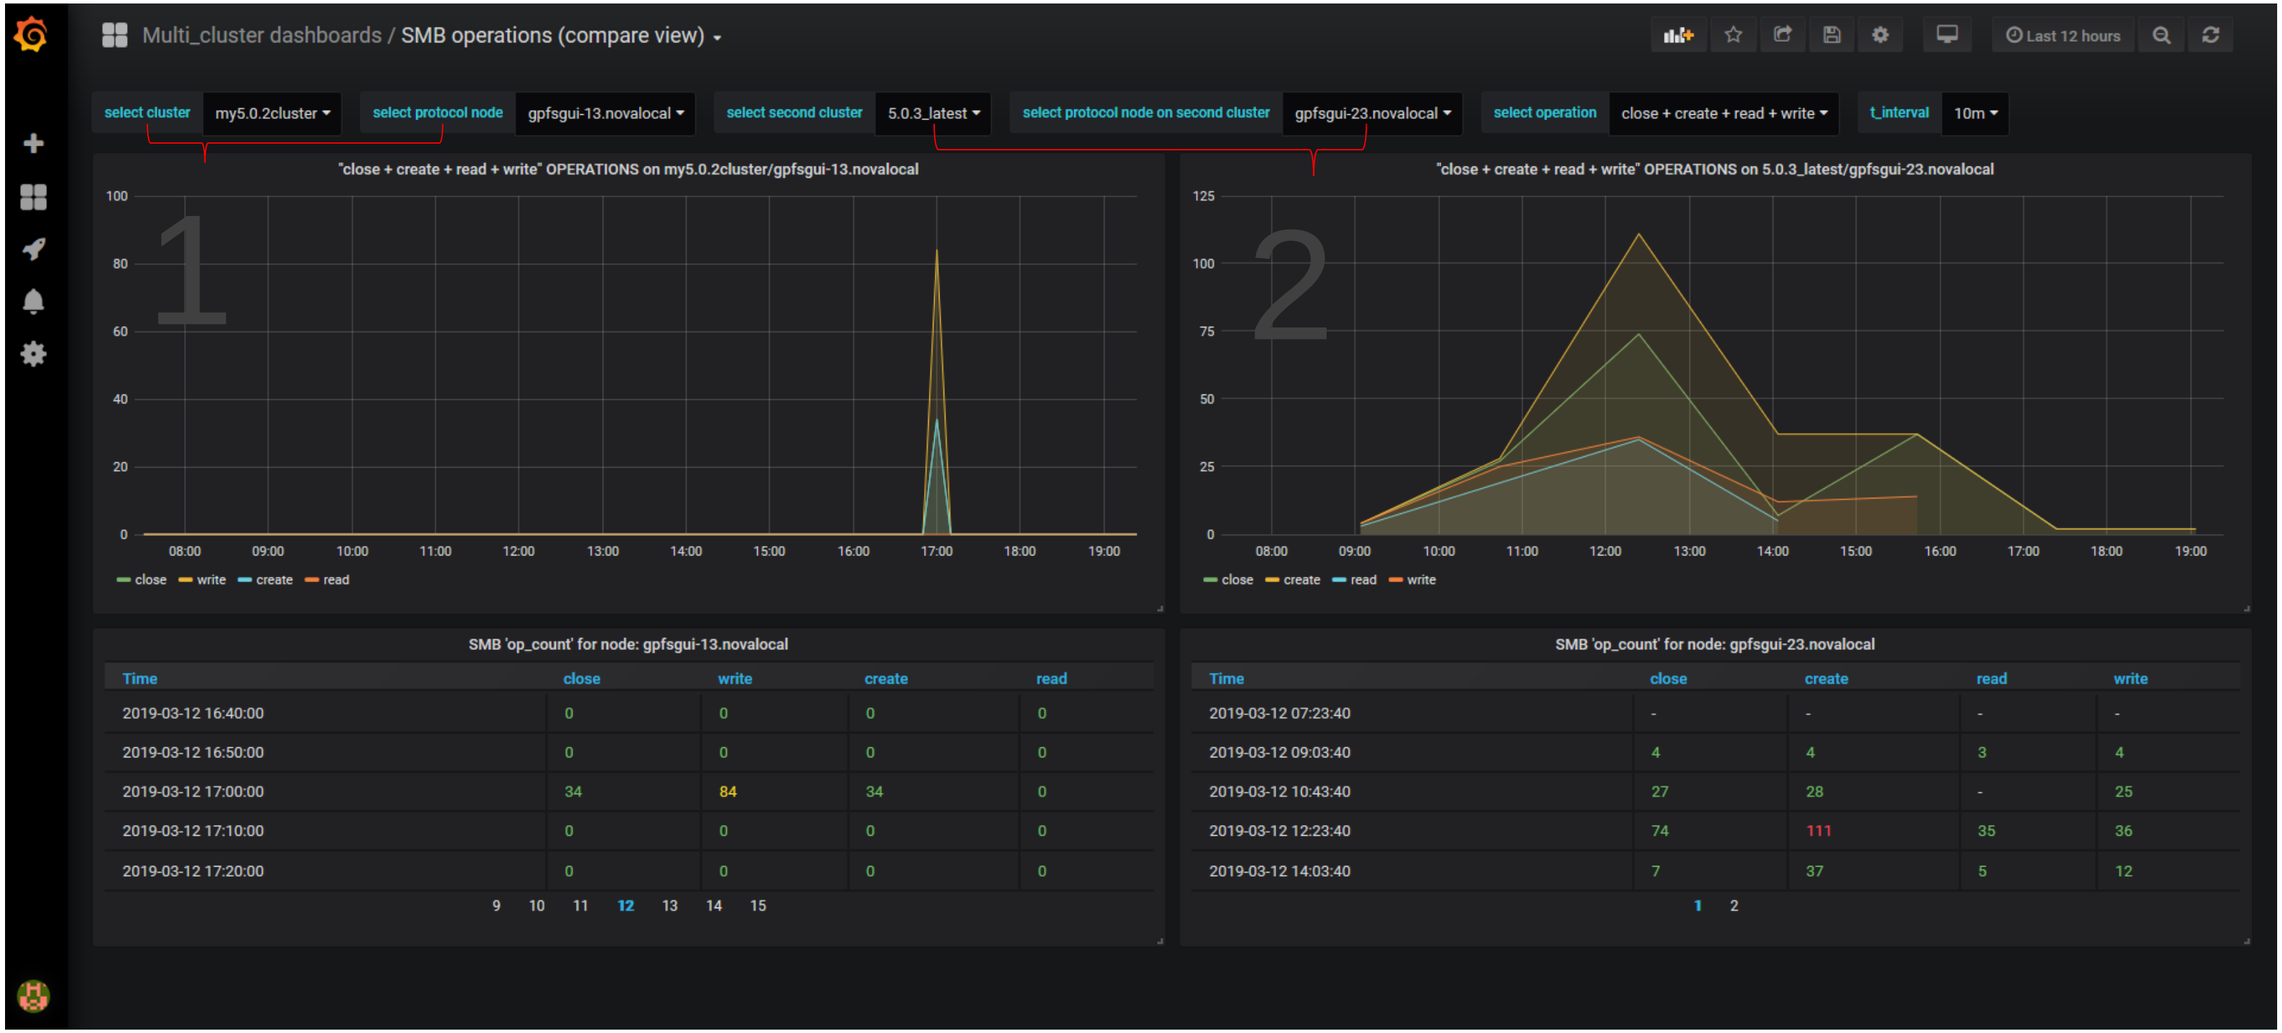

Use case “**SMB operations**”

Use case “**SMB operations**”

Use case “**SMB operations**”

IBM Spectrum Scale

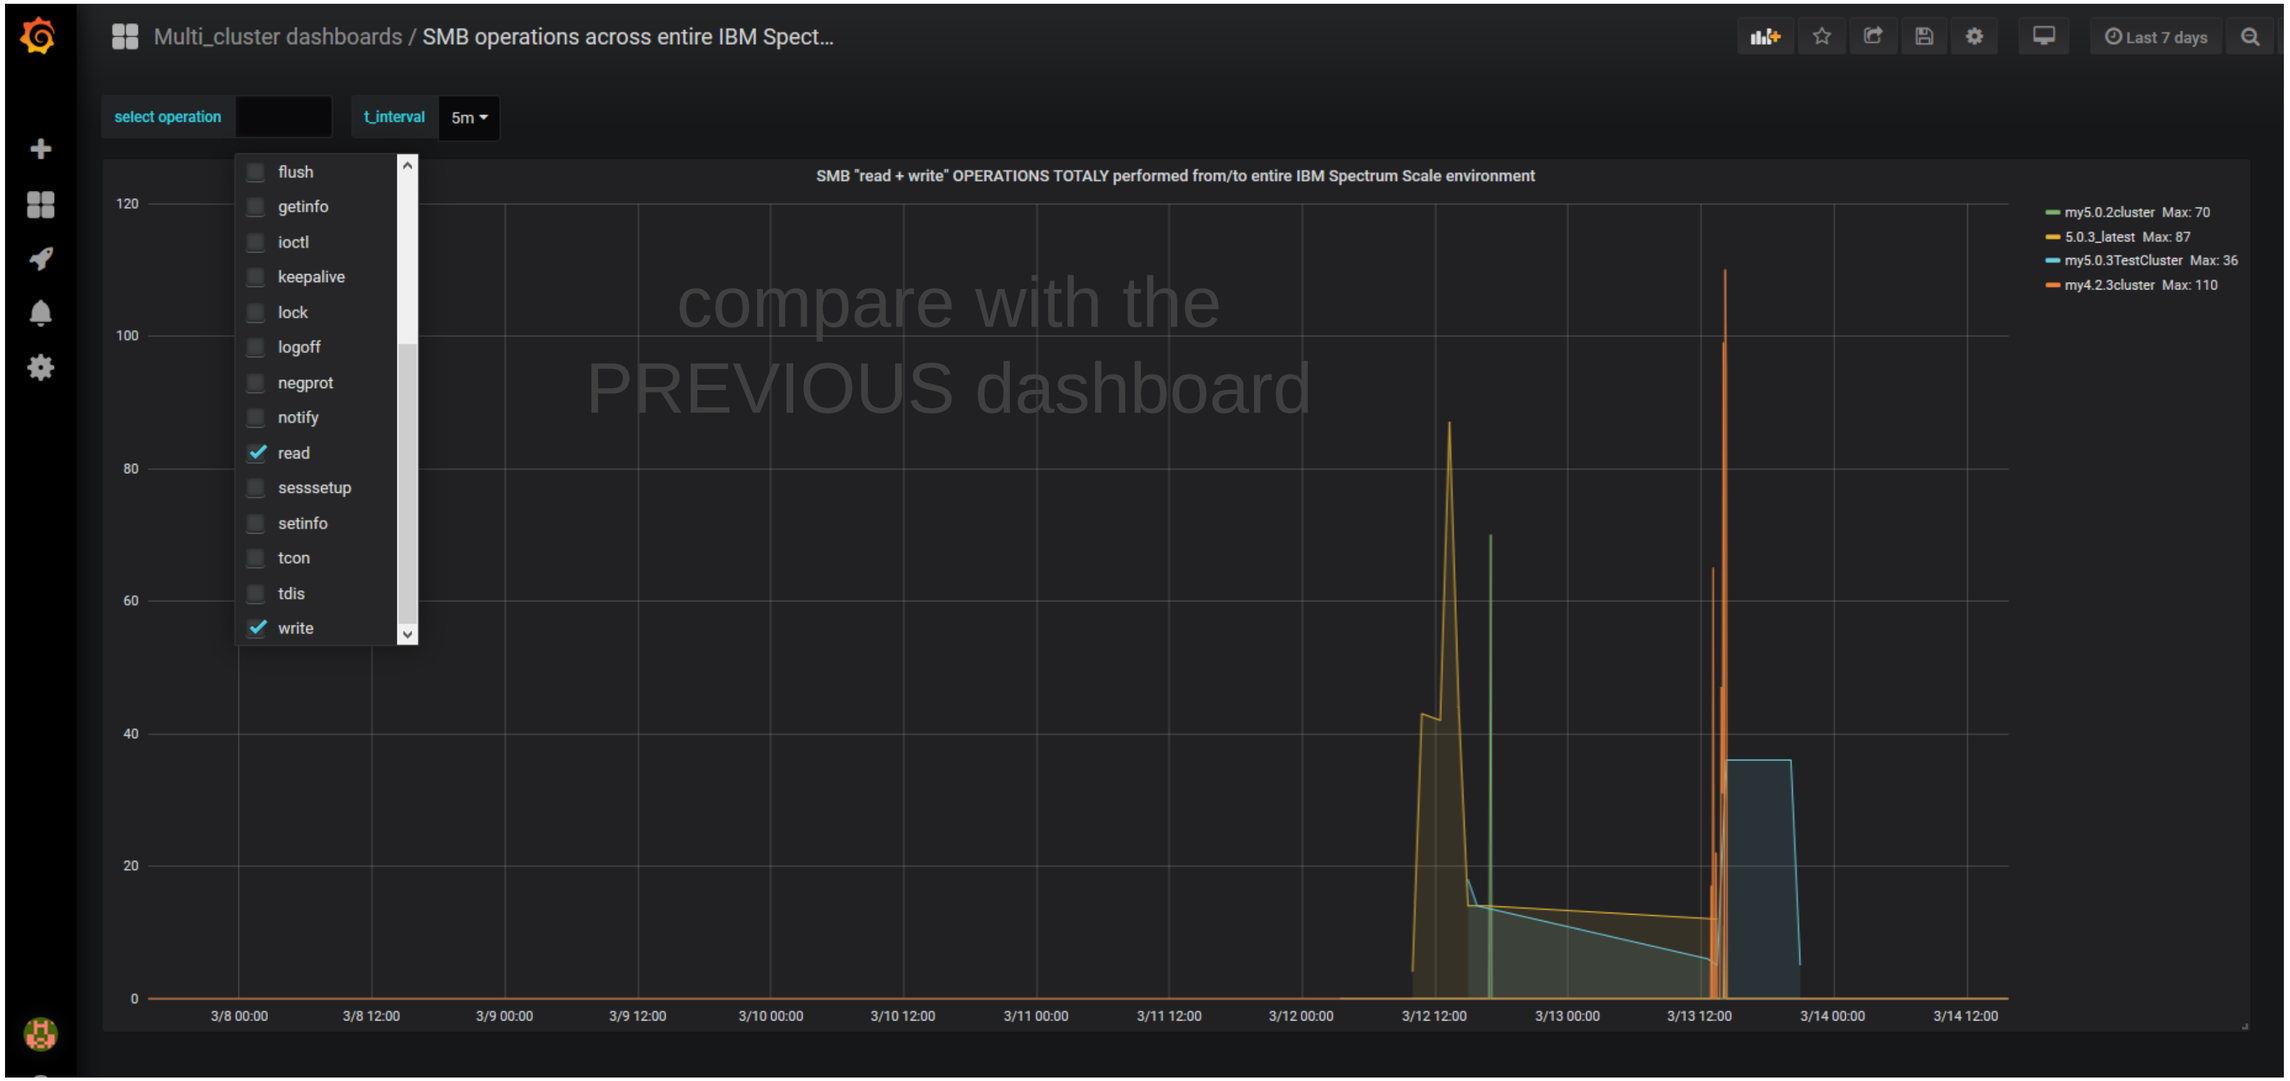

Use case “**SMB operations**”

IBM Spectrum Scale

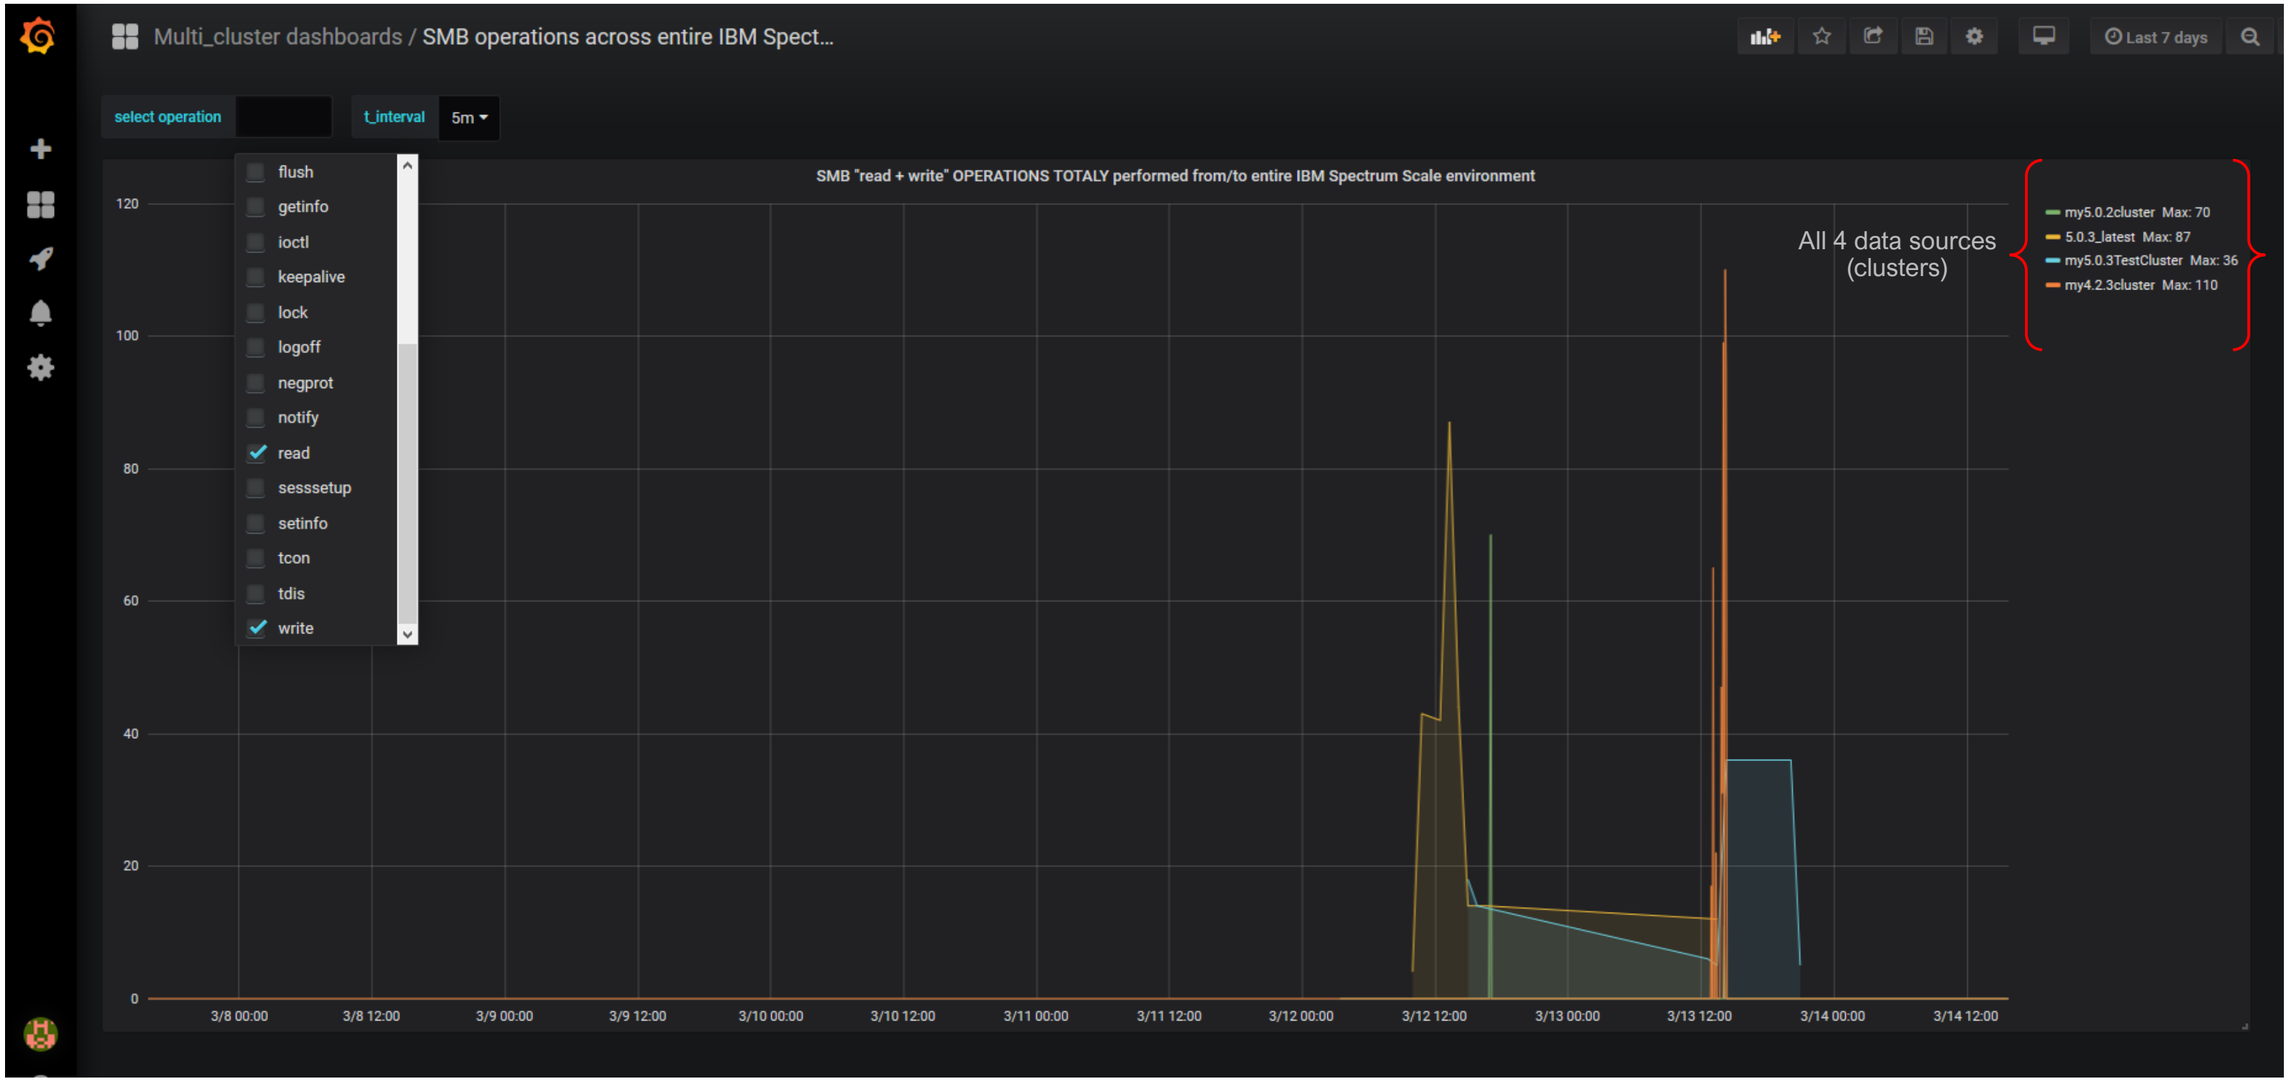

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. Reference Materials

Use Case “**data transfers to a cloud**”

IBM Spectrum Scale

Use Case “**data transfers to a cloud**”

IBM Spectrum Scale

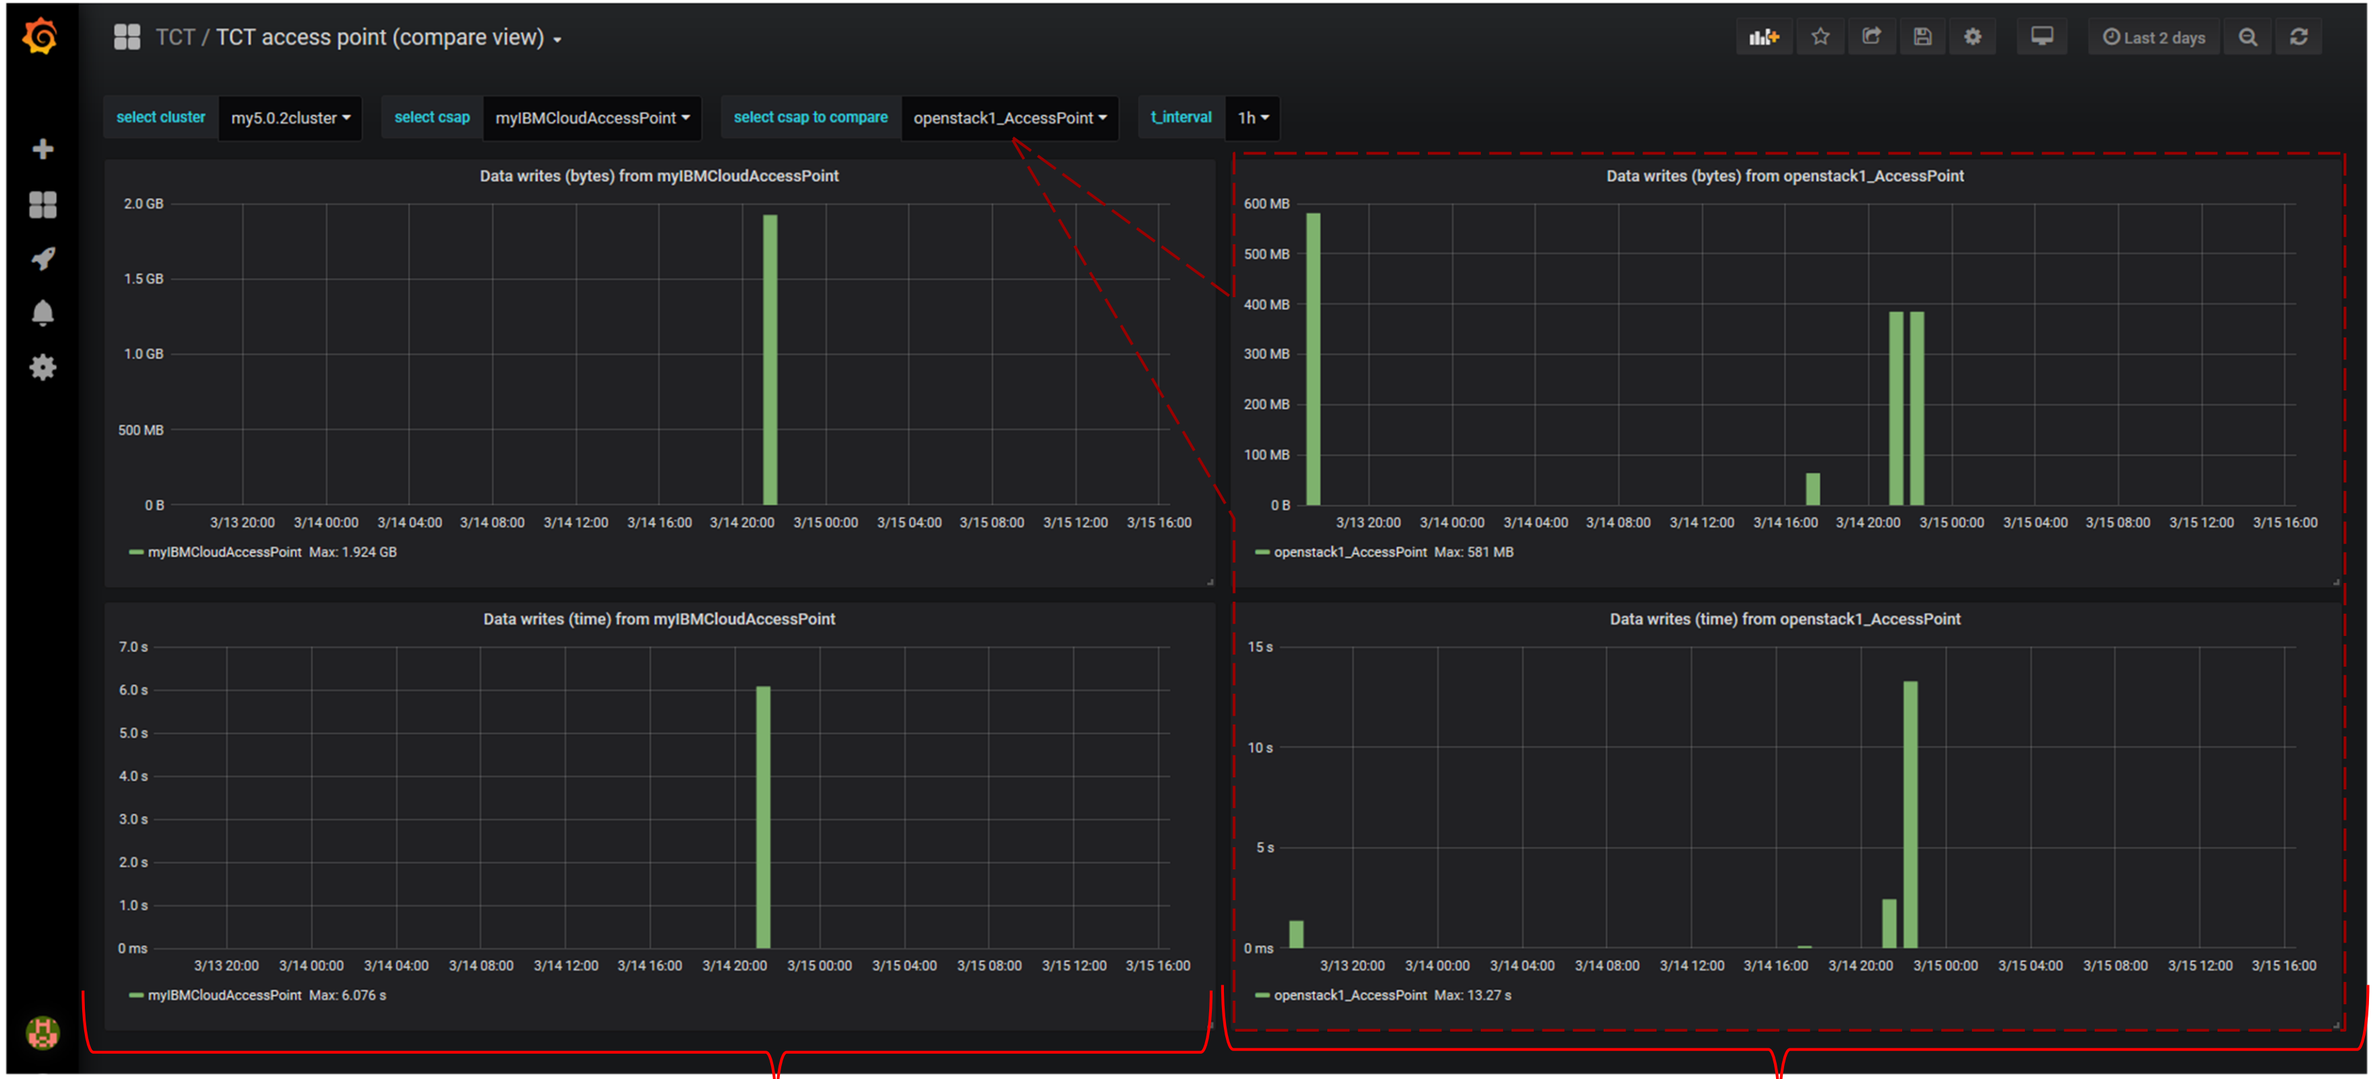

cloud-type = CLEVERSAFE-NEW

cloud-type = OPENSTACK-SWIFT

Use Case “**data transfers to a cloud**”

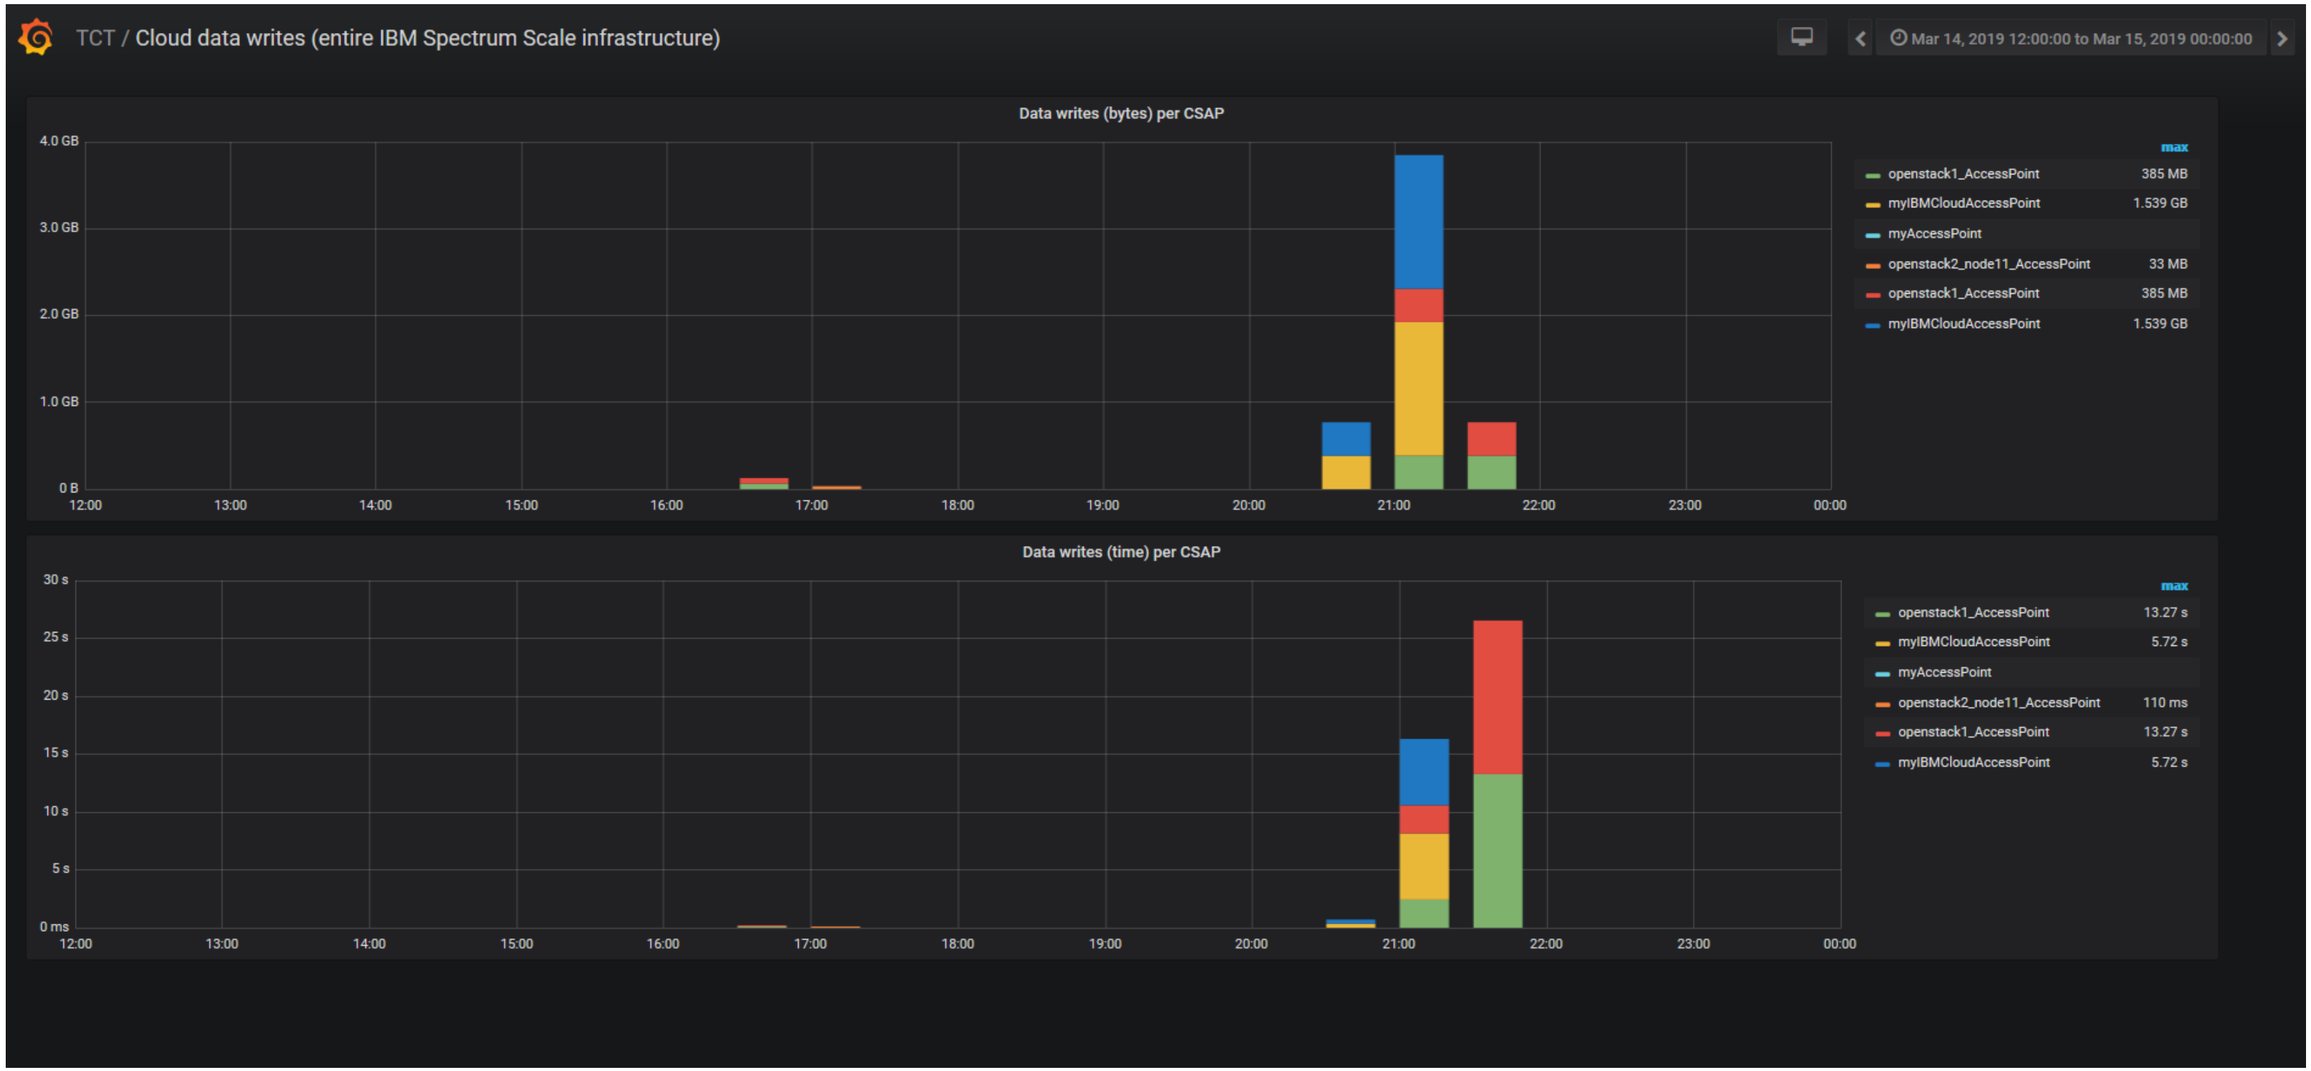

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. **Performance Data analysis strategies**
5. Summary
6. Reference Materials

Performance Data analysis strategies

IBM Spectrum Scale

Time series analysis goals:

- identifying the nature of the phenomenon represented by the sequence of observations
- forecasting (predicting future values of the time series variable)

Performance is based on measurements and forecasts in a controlled environment. The performance of IBM Spectrum Scale environment might vary depending on:

- the amount of multiprogramming in the users job stream
- the I/O configuration
- the storage configuration
- the workload processed

Performance Data analysis strategies

IBM Spectrum Scale

Short time span queries

- a lot of ‘single time series’ graphs in one dashboard ?
- can benefit from filtering, selecting groups, combining corresponding metrics in a table
- useful for random error observations, examination of anomalies (Monitoring)

Wide time span queries

- number of data points more than number of pixels in a display screen ?
- thousands of data points per graph can affect response time of the browser
- can benefit from downsampling, aggregating
- useful for creating high level charts, “trend charts” (Trending)

Performance Data analysis strategies

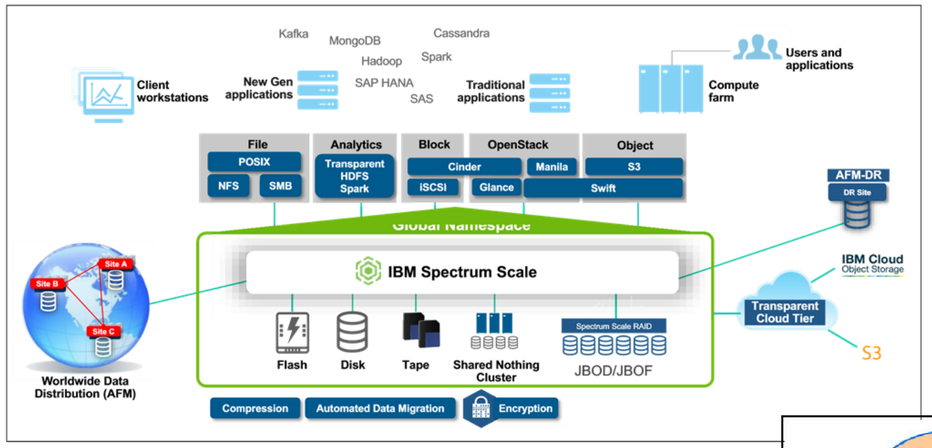

focus on a specific/individual use case

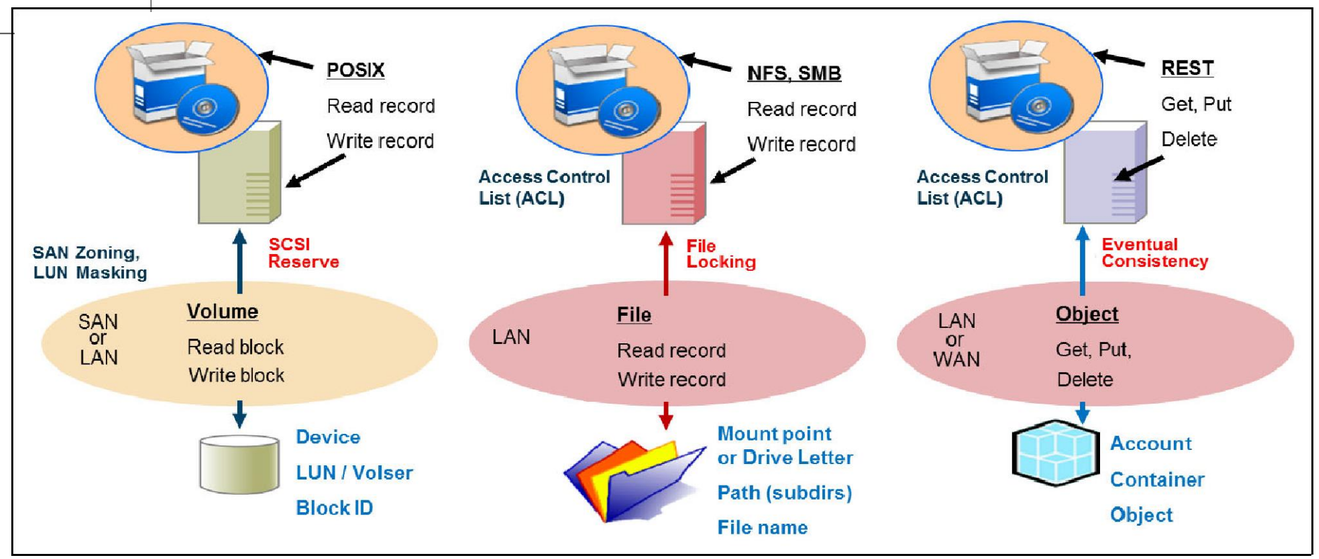

keep in mind the complexity of a whole environment

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. **Summary**
6. Reference Materials

Summary

IBM Spectrum Scale

- analysis methods vary depending on use case
- visualizing data provides visibility required for understanding what transpired at a given point in time
- Grafana is one of the most popular dashboard composer software (open source)
- the usage of Grafana is beneficial for:
  - data monitoring of multiple IBM Spectrum Scale clusters
  - data visualizing for mixed environments
  - large-scale displays throughout
  - comfortable usability during data research & product development experimentation phase
  - advanced performance monitoring

Outline

1. What is Grafana?
    a. Evolution & Key concepts
    b. Grafana features highlights
    c. Integration with IBM Spectrum Scale
2. IBM Spectrum Scale Bridge for Grafana updates
3. New bundle ‘Example Dashboards’
    a. Getting Started with ‘HOWTO Dashboards’
    b. Use Case ‘SMB operations’
    c. Use Case ‘data transfers to a cloud’
4. Performance Data analysis strategies
5. Summary
6. **Reference Materials**

Reference Materials

IBM Spectrum Scale bridge for Grafana

- IBM Spectrum Scale Bridge for Grafana on the developerWorks/IBM Spectrum Scale Wiki
- IBM Spectrum Scale bridge for Grafana in the IBM Knowledge Center

IBM Spectrum Scale useful links

- IBM Spectrum Scale - Performance Monitoring
- IBM Spectrum Scale - Viewing and analyzing the performance data
- IBM Spectrum Scale - list of performance metrics
- Enabling Hybrid Cloud Storage for IBM Spectrum Scale Using Transparent Cloud Tiering
- A Deployment Guide for IBM Spectrum Scale Unified File and Object Storage
- Cloud Object Storage as a Service
- IBM Software-Defined Storage Guide

Reference Materials

IBM Spectrum Scale

Grafana

- http://docs.grafana.org/
- https://community.grafana.com/c/howto/faq
- Checkout the CHANGELOG.md file for a complete list of new features, changes, and bug fixes

External materials

- http://opentsdb.net/docs/build/html/user_guide/query/aggregators.html
- http://opentsdb.net/docs/build/html/user_guide/query/performance.html
- http://docs.cherrypy.org/en/latest/
- https://guides.github.com/features/mastering-markdown/
- https://logz.io/blog/grafana-vs-kibana/
- https://www.betterevaluation.org/en/evaluation-options/timeseriesanalysis

Experts talks (articles in German)

- https://www.embarc.de/wp-content/uploads/2016/08/Traenen_luegen_nicht-Dashboards_schon_deploy.pdf
- https://www.admin-magazin.de/Das-Heft/2017/06/Skalierbares-Monitoring-mit-Prometheus
- http://www.linux-magazin.de/downloads/der-datenvisualisierer-grafana-und-seine-backends/attachment/034-039_grafana/
- https://www.informatik-aktuell.de/entwicklung/methoden/monitoring-das-maechtigste-werkzeug-fuer-cloud-microservices-und-business.html