
Schedule

First session (90 minutes)

- Overview of architectures for HPCM and CSM (Jeff Hanson, HPE) (30 minutes)
- Ryan Adamson (ORNL) (20 min)
- John Stile (NERSC) (20 min)
- Siqi Deng (NERSC) (20 min)

Second session (90 minutes)

- AIOps (Martin Foltin, HPE) (45 minutes)
- Optional feature enablement (Jeff Hanson, HPE) (45 minutes)
 - Slingshot (15 min)
 - PCIM and PDUs (5 min)
 - SNMP presentation from PCIM (10 min)
 - LDMS (15 min)





**Hewlett Packard
Enterprise**



HPCM Monitoring Architecture Review

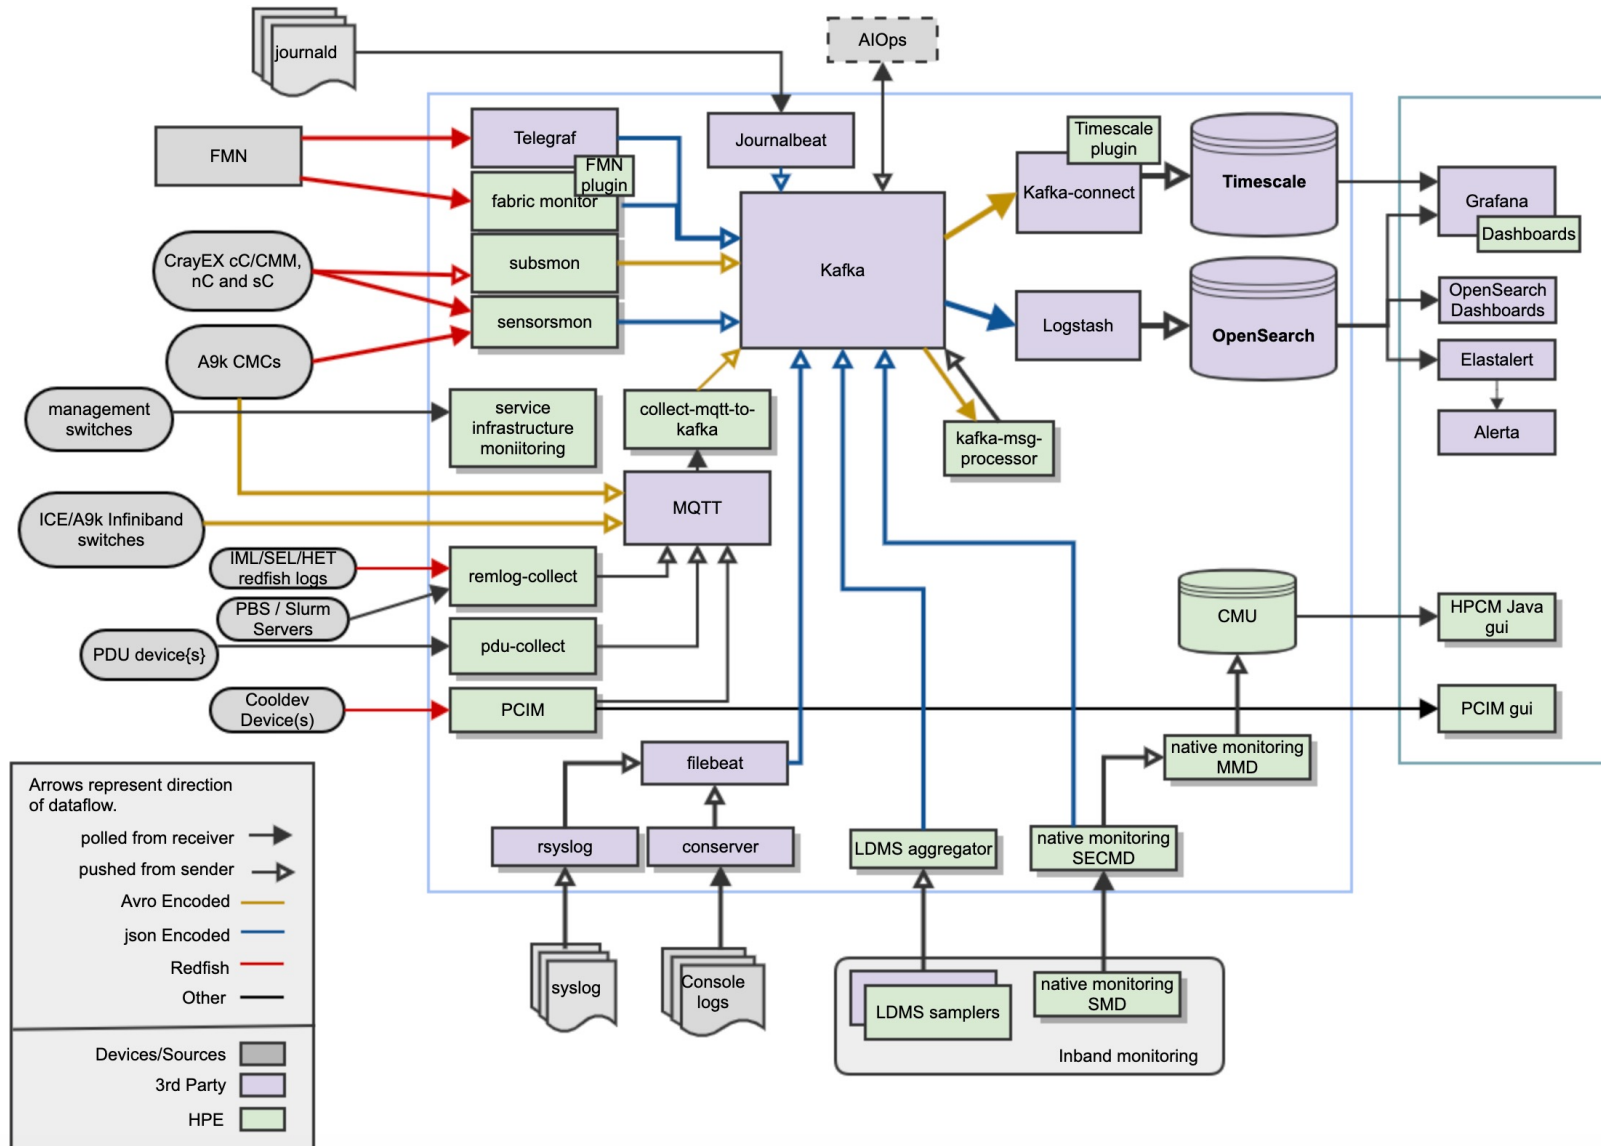
**Jeff Hanson, HPCM/SMA Monitoring
Manager**

May 2023

HPCM Monitoring Architecture Review

- Overview diagram
- Apache Avro and Confluent Schema Registry
- Data sources
 - Redfish
 - IPMI
 - Custom
 - Compute node
- Data persistence
- Visualizations

HPCM Monitoring Overview



Apache Avro and Confluent Schema Registry

- Some data is encoded with [Apache Avro](#) (long term plan is to encode all of it).
- Apache Avro enforces schema and provides serialization of data.
- When coupled with [Confluent Schema Registry](#) the overhead of using Avro is reduced because schema is stored as a reference in Schema Registry rather than in the file.
- HPCM preloads all schemas used on initialization of the Kafka environment. There are roughly 55 schema in use (each one has a key and value schema).
- Other Kafka implementations have Schema Registry ([Strimzi](#)) ([Apicurio](#)). HPCM has no experience with non Confluent.

Data Sources

Redfish - CrayEX

- CrayEX node Controller (nC), chassis Controller (cC), switch Controller (sC)
- HPCM calls them (from cm controller show)
sC example x3001c0r40b0 external_switch_controller and x9000c1r1b0 cmm_switch_controller
nC example x9000c1s0b0 cmm_node_controller
cC x9000cec1 cabinet_environment_controller
- Metrics and Events are gathered via subscription service using subsmon. Controllers push to the subscriber (in the normal Redfish manner).
- Credentials for controllers are stored in HPCM database and used by subsmon to subscribe.
- Use [Redfish Mockup Creator](#) to explore.

Data Sources

Redfish - HPE

- HPE systems using iLO
- HPCM calls them <nodename>-bmc
- Metrics gathered via sensormon which polls the iLOs.
- Credentials for controllers are stored in HPCM database and used by sensormon to poll.
- Events (IML) are polled via remlog-collect

Data Sources

Redfish - Other

- Neither subsmon nor sensormon are general purpose redfish subscribers. They only support what hardware HPE provides in HPCM managed clusters as node hardware.
- Redfish on CDUs is done via PCIM which will be discussed in a later slide.

Data Sources

IPMI

- Sensormon can query IPMI devices limited to what hardware HPE provides in HPCM managed clusters as node hardware.
- Older HPCM has a SEL log scrapper but that has been retired.
- HET (Hardware Event Tracker) provides use PEF for eventing for IPMI devices.

Data Sources

Custom

- PCIM (Power and Cooling Infrastructure Monitoring) is a service that collects metrics and events from CDU and PDU devices.
 - In HPCM cooling devices are automatically managed by PCIM. There is a standalone method in PCIM as well so without HPCM it can be told about CDUs.
 - In CrayEX the CEC are the bridge to CDUs and are queried via Redfish.
 - For PDUs a configuration step is needed to enable SNMP and provide credentials for PCIM to be able to query.
- Fabric manager metrics are sent via telegraf to Kafka

Data Sources

Compute node

- Conservers log to su-leader nodes for compute nodes, other nodes use admin for logging. Filebeat watches the logs and writes to Kafka.
- Rsyslog logs to su-leader nodes for compute nodes, other nodes and controllers use admin for logging. Filebeat watches the logs and writes to Kafka.
- Filebeat writes a limited subset of journald to Kafka.
- Two options (which can be used together or separately) for OS metrics
 - HPCM native monitoring – a proprietary solution that leverages collectl and custom tools. Is extensible.
 - LDMS – supporting all the Cray samplers including DVS and Cassini. Uses currently closed source Kafka store plugin.

Persistence

MQTT

- Several services use mqtt (mosquitto) as a bridge to Kafka
 - PCIM writes out files that it then pushes to mqtt which are read by collect-mqtt-to-kafka to be serialized with Avro.
 - Same for remlog
 - Apollo 9000 CMC and InfiniBand switch controllers have embedded mqtt tools. Mqtt is a better fit for controller devices than Kafka.
- Mqtt brokers are distributed to su-leaders

Persistence

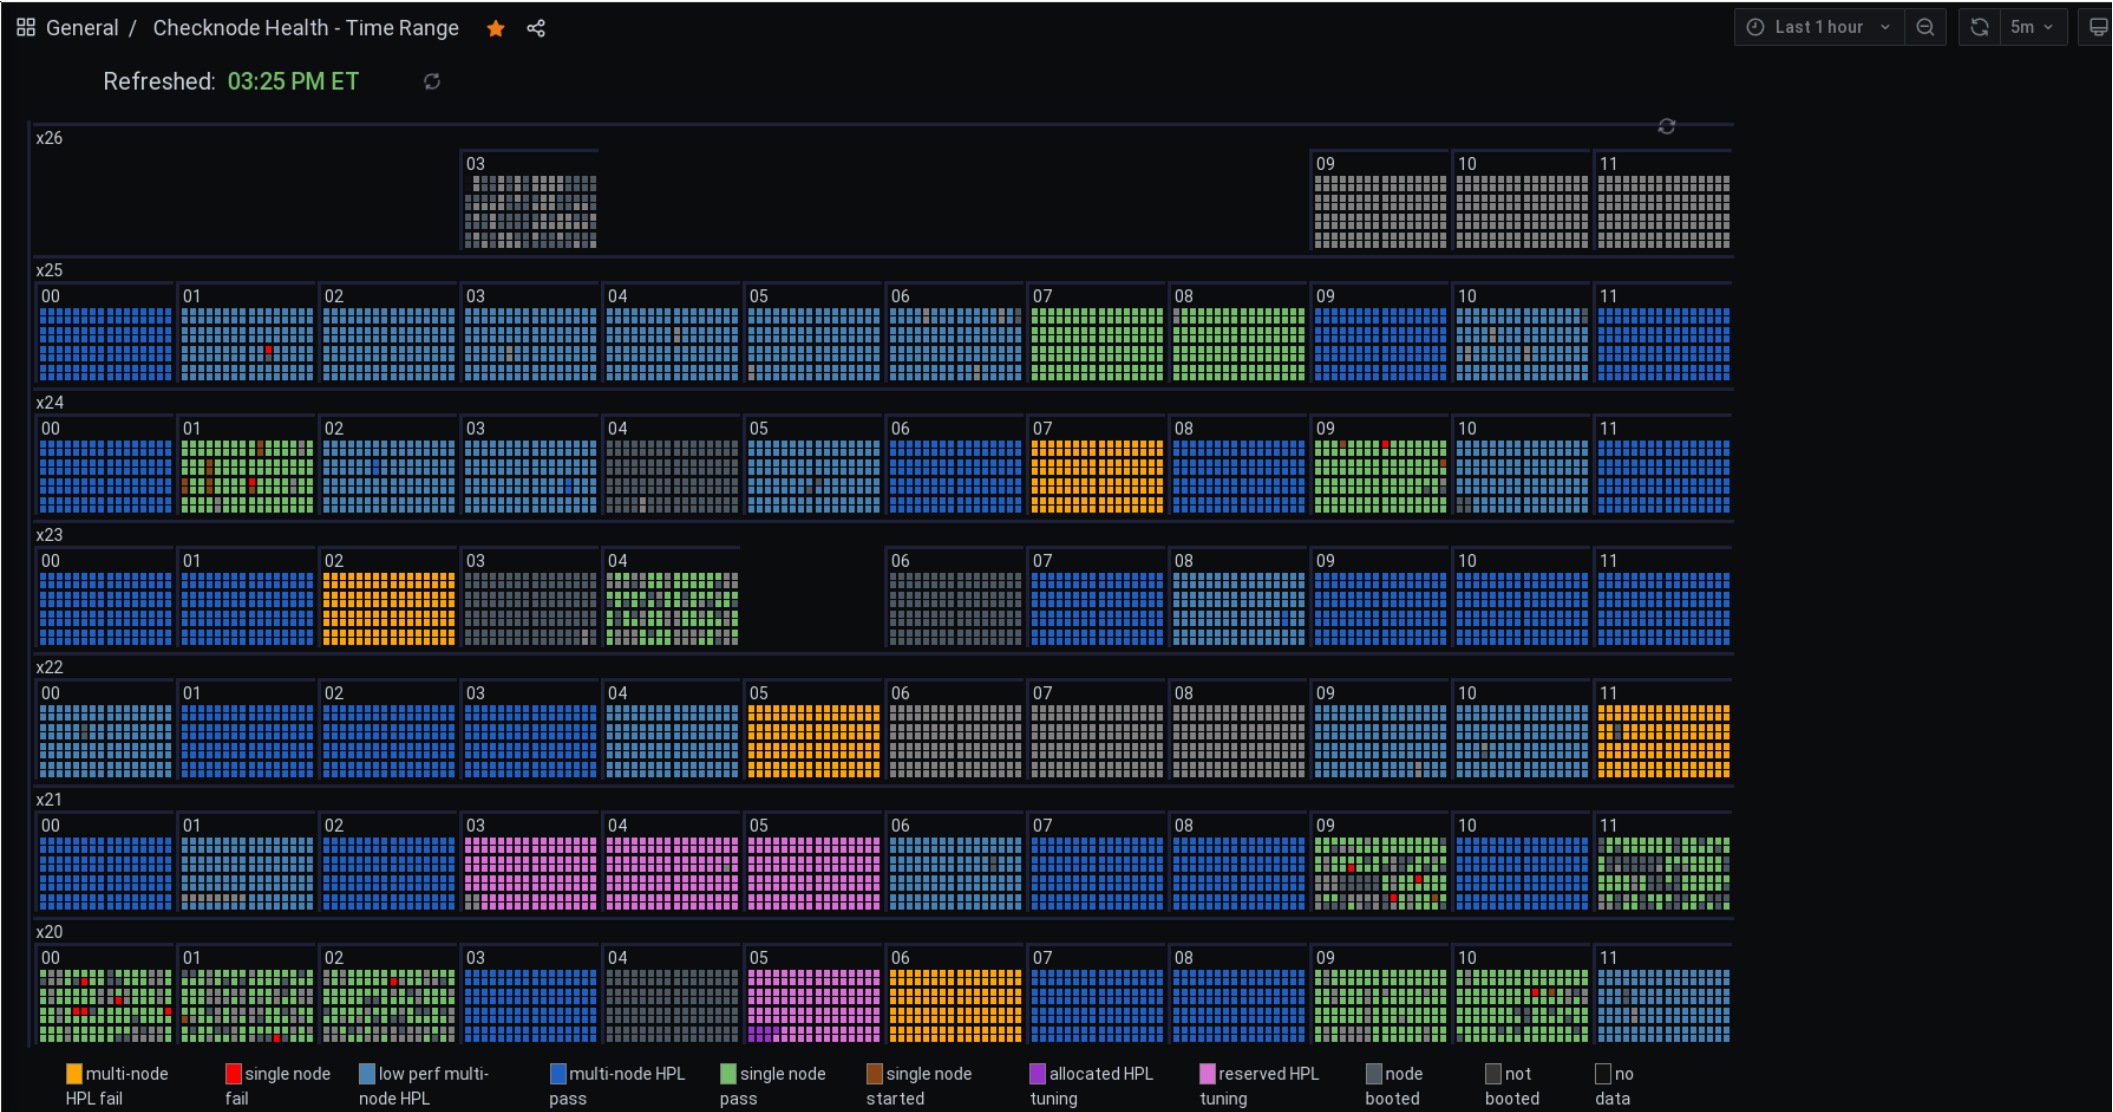
Kafka

- HPCM mantra is Kafka is in the middle always.
- HPCM uses Confluent Kafka with either single node for services or distributed to su-leaders.

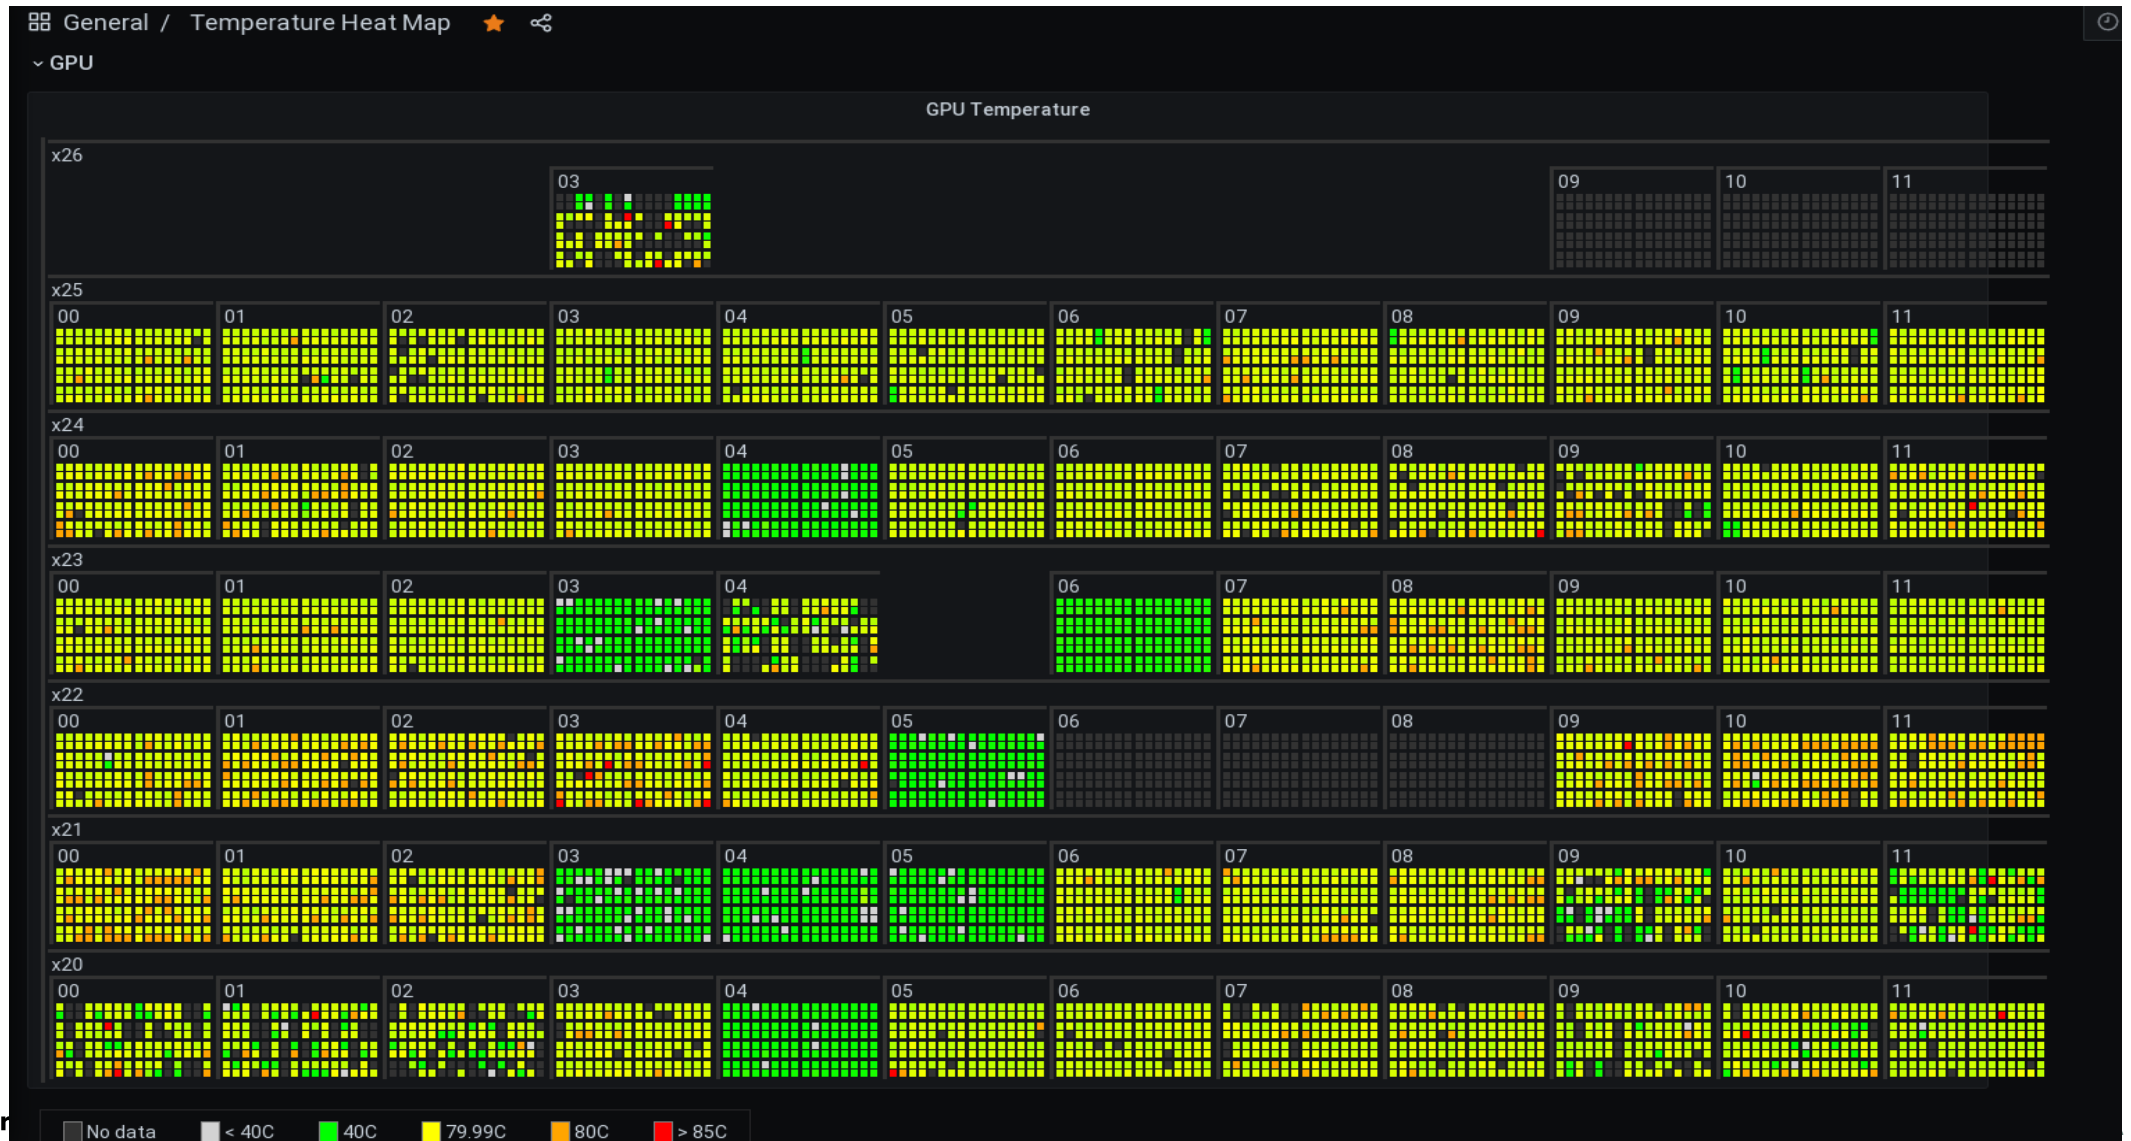
Persistence Databases

- OpenSearch is database for event and log data. Some metrics are still there to enable existing ElastAlert rules to continue to work. Prior to HPCM 1.8 all data was in ElasticSearch
- TimeScaleDB is database for metric data.
- Both can be distributed to su-leader nodes.

Main dashboard showing realistic view of cluster



Main dashboard showing simple heatmap of a metric



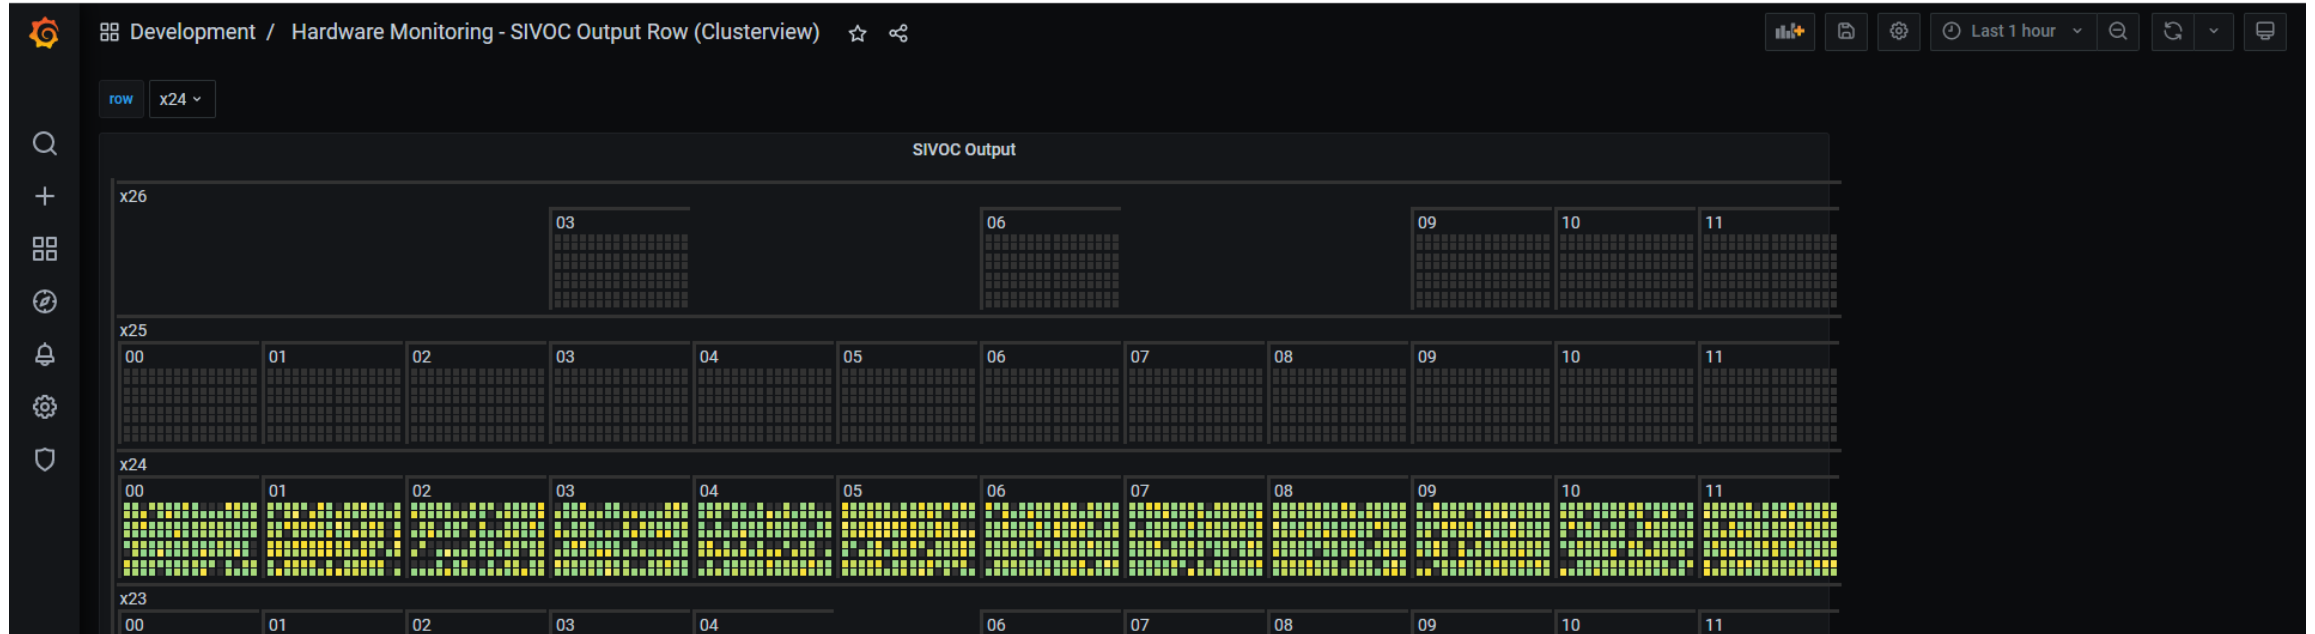
SIVOC OUTPUT MINIMUMS CLUSTERVIEW

Are lower SIVOC Output values more common in certain racks? (x2504, x2508, x2708)



SIVOC MINIMUM VOLTAGE (CLUSTER VIEW)

Identify which nodes under a PDU have lower density



[0-41]-Red
[41-42]-Light Orange
[42-43]-Dark Orange
[43-44]-Light Yellow
[44-45]-Dark Yellow
[45-46]-Light Green
[46-48]-Dark Green

Node level dashboard showing hardware metrics



Fabric monitoring, data from Frontier

