



# EIC PD Computing and Software Services I: Control and Monitoring Services

De Stefano, *et al*,  
on behalf of the Scientific Computing & Data Facility (SCDF)  
and other collaborators

EIC Computing PD at BNL AHM: 5 Dec 2025



# Disclaimer

This work is the sum of many efforts in and without our organization; only the details of those directly related to this PD, and who have reported their efforts here, are included in detail. Apologies to those whose efforts are not mentioned.

This work is evaluatory in nature and still in development. As such, much of the technology is not yet available outside the working group or for wider testing.

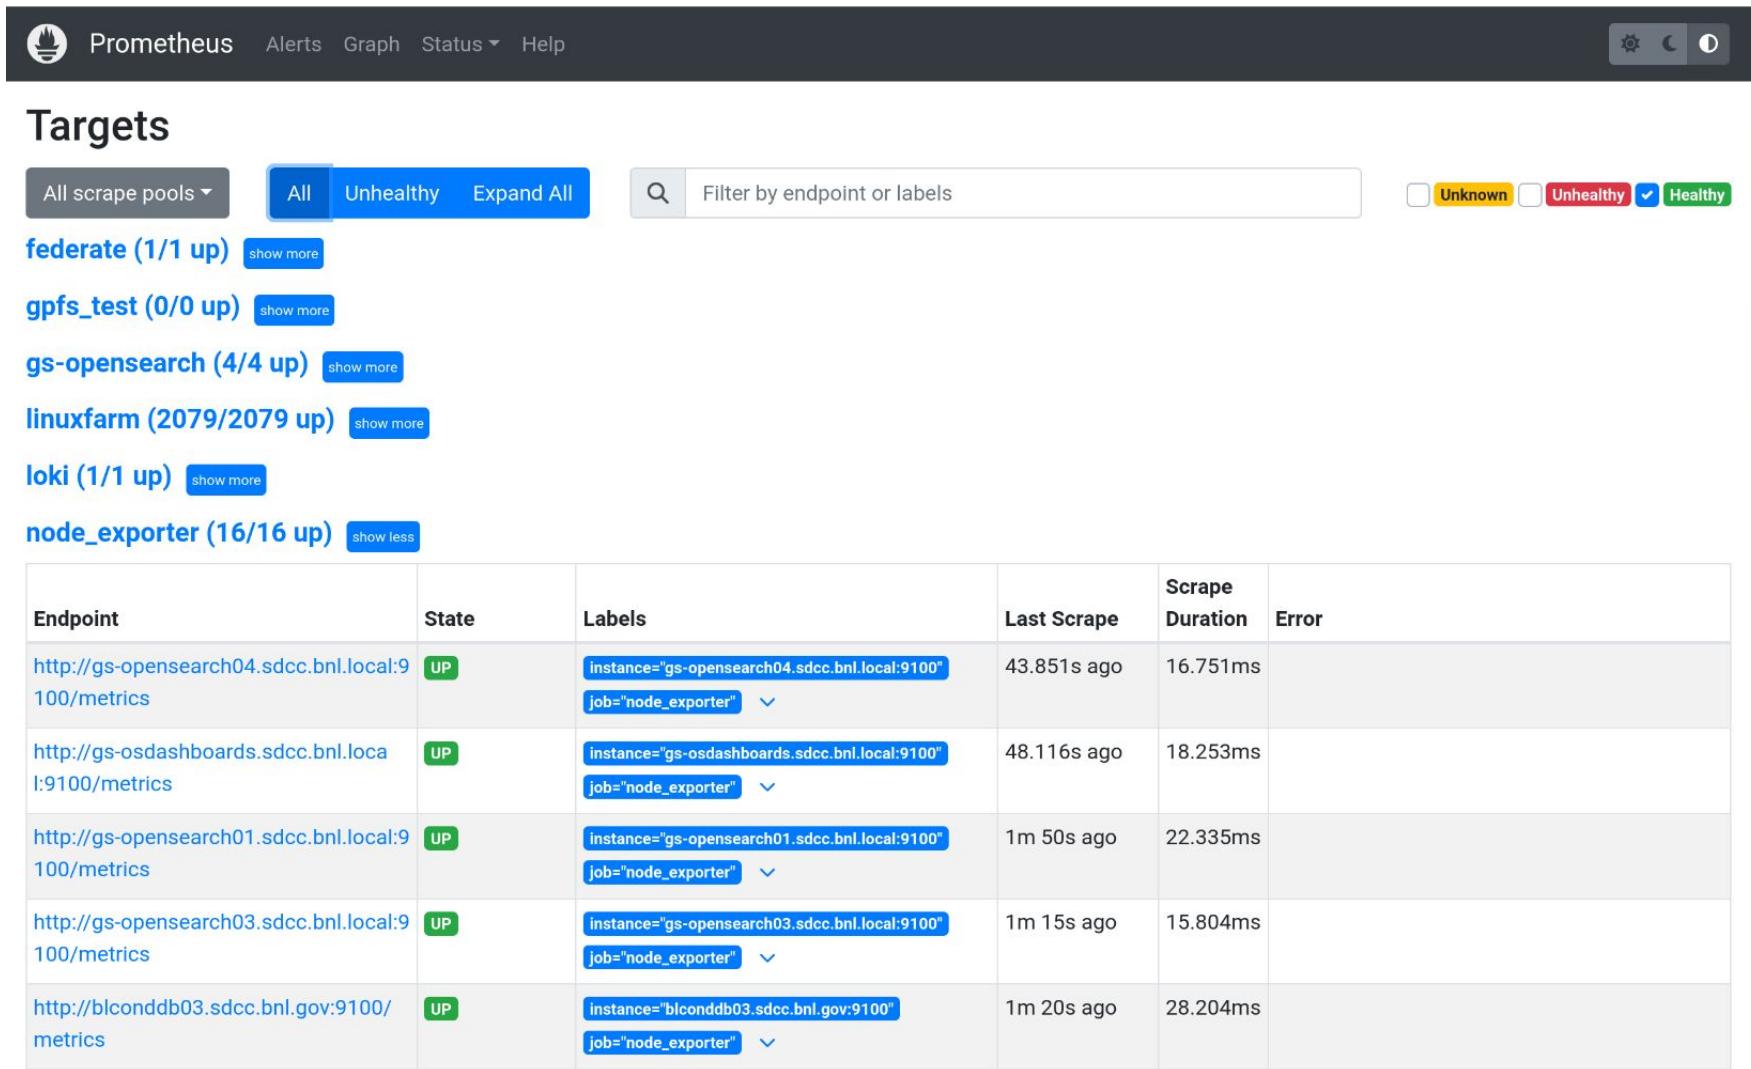
# Monitoring: Status & Achievements

# Monitoring technology: Prometheus

Evaluated and built Prometheus instances for development and production (Pelosi)

- Initial deployments of node exporter and other service exporters (e.g., PostgreSQL exporter)

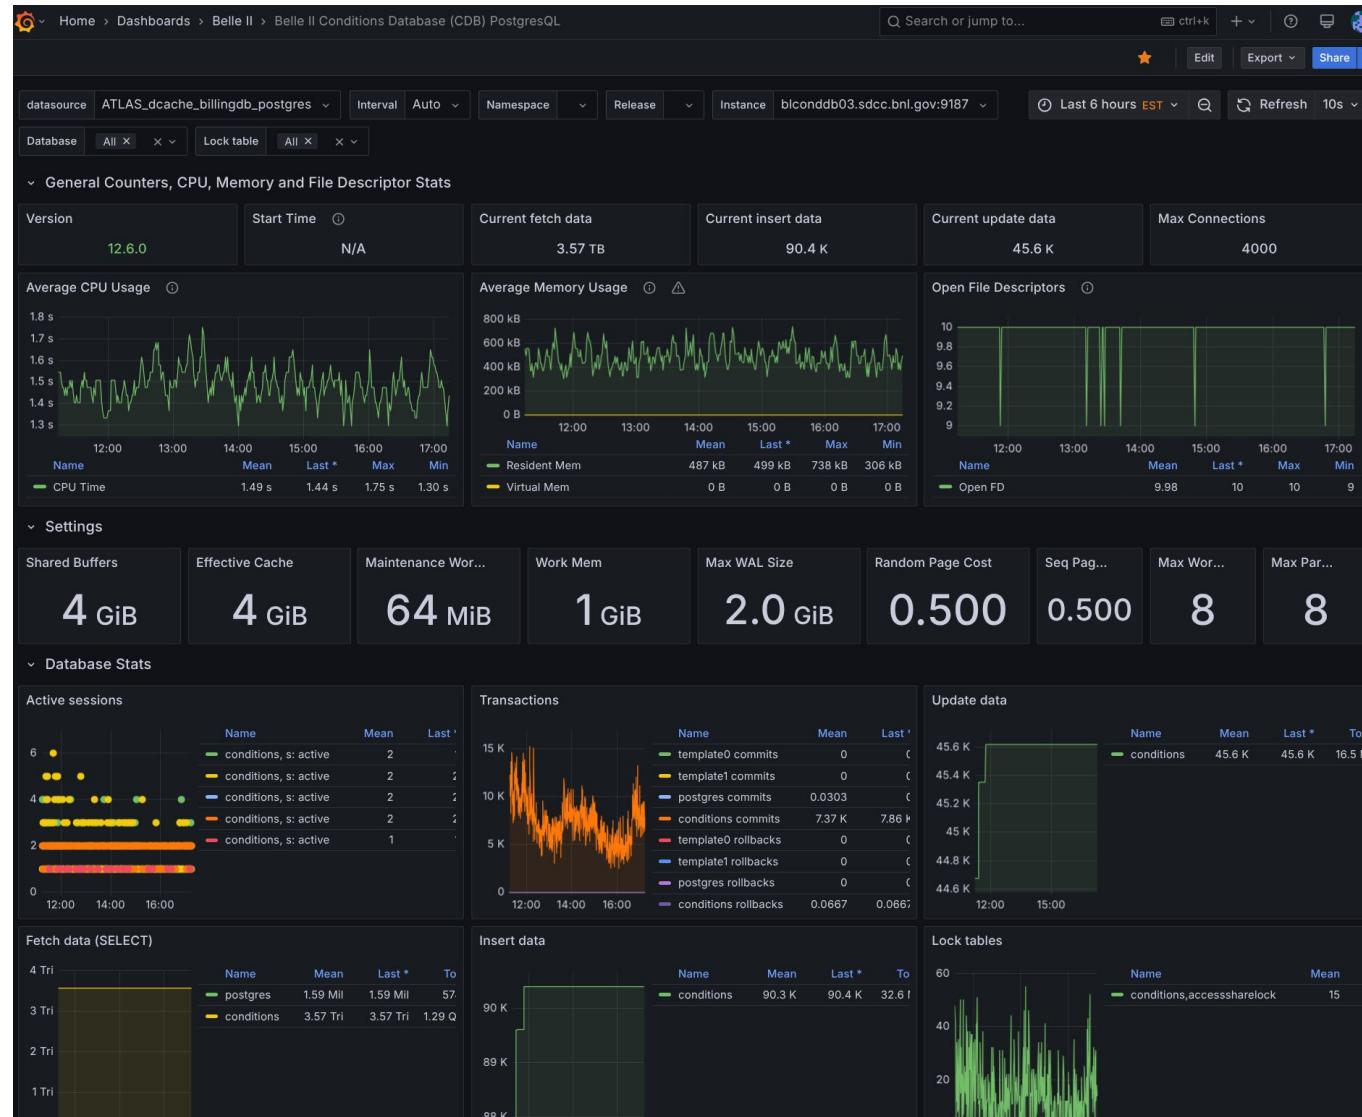
# Prometheus: Targets



The screenshot shows the Prometheus Targets page. At the top, there is a navigation bar with links for Prometheus, Alerts, Graph, Status, and Help. On the right side of the navigation bar are three small icons: a gear, a crescent moon, and a circle. Below the navigation bar, the page title is "Targets". There are several buttons and filters: "All scrape pools" (dropdown), "All" (selected), "Unhealthy", "Expand All", a search bar with placeholder "Filter by endpoint or labels", and a color-coded legend for status: Unknown (yellow), Unhealthy (red), and Healthy (green). The main content area lists monitored endpoints with their status and labels. The table has columns: Endpoint, State, Labels, Last Scrape, Scrape Duration, and Error. The "Labels" column contains dropdown menus for "instance" and "job".

| Endpoint   | State | Labels  | Last Scrape | Scrape Duration | Error |
|--|-------|---|-------------|-----------------|-------|
| <a href="#">http://gs-opensearch04.sdcc.bnl.local:9100/metrics</a> | UP    | instance="gs-opensearch04.sdcc.bnl.local:9100"<br>job="node_exporter" ▾ | 43.851s ago | 16.751ms        |       |
| <a href="#">http://gs-osdashboards.sdcc.bnl.local:9100/metrics</a> | UP    | instance="gs-osdashboards.sdcc.bnl.local:9100"<br>job="node_exporter" ▾ | 48.116s ago | 18.253ms        |       |
| <a href="#">http://gs-opensearch01.sdcc.bnl.local:9100/metrics</a> | UP    | instance="gs-opensearch01.sdcc.bnl.local:9100"<br>job="node_exporter" ▾ | 1m 50s ago  | 22.335ms        |       |
| <a href="#">http://gs-opensearch03.sdcc.bnl.local:9100/metrics</a> | UP    | instance="gs-opensearch03.sdcc.bnl.local:9100"<br>job="node_exporter" ▾ | 1m 15s ago  | 15.804ms        |       |
| <a href="#">http://blconddb03.sdcc.bnl.gov:9100/metrics</a>        | UP    | instance="blconddb03.sdcc.bnl.gov:9100"<br>job="node_exporter" ▾        | 1m 20s ago  | 28.204ms        |       |

# Prometheus: PostgresQL exporter



# Monitoring technology: OpenSearch

Rebuilt legacy ElasticSearch nodes as new Opensearch cluster  
(Pelosi)

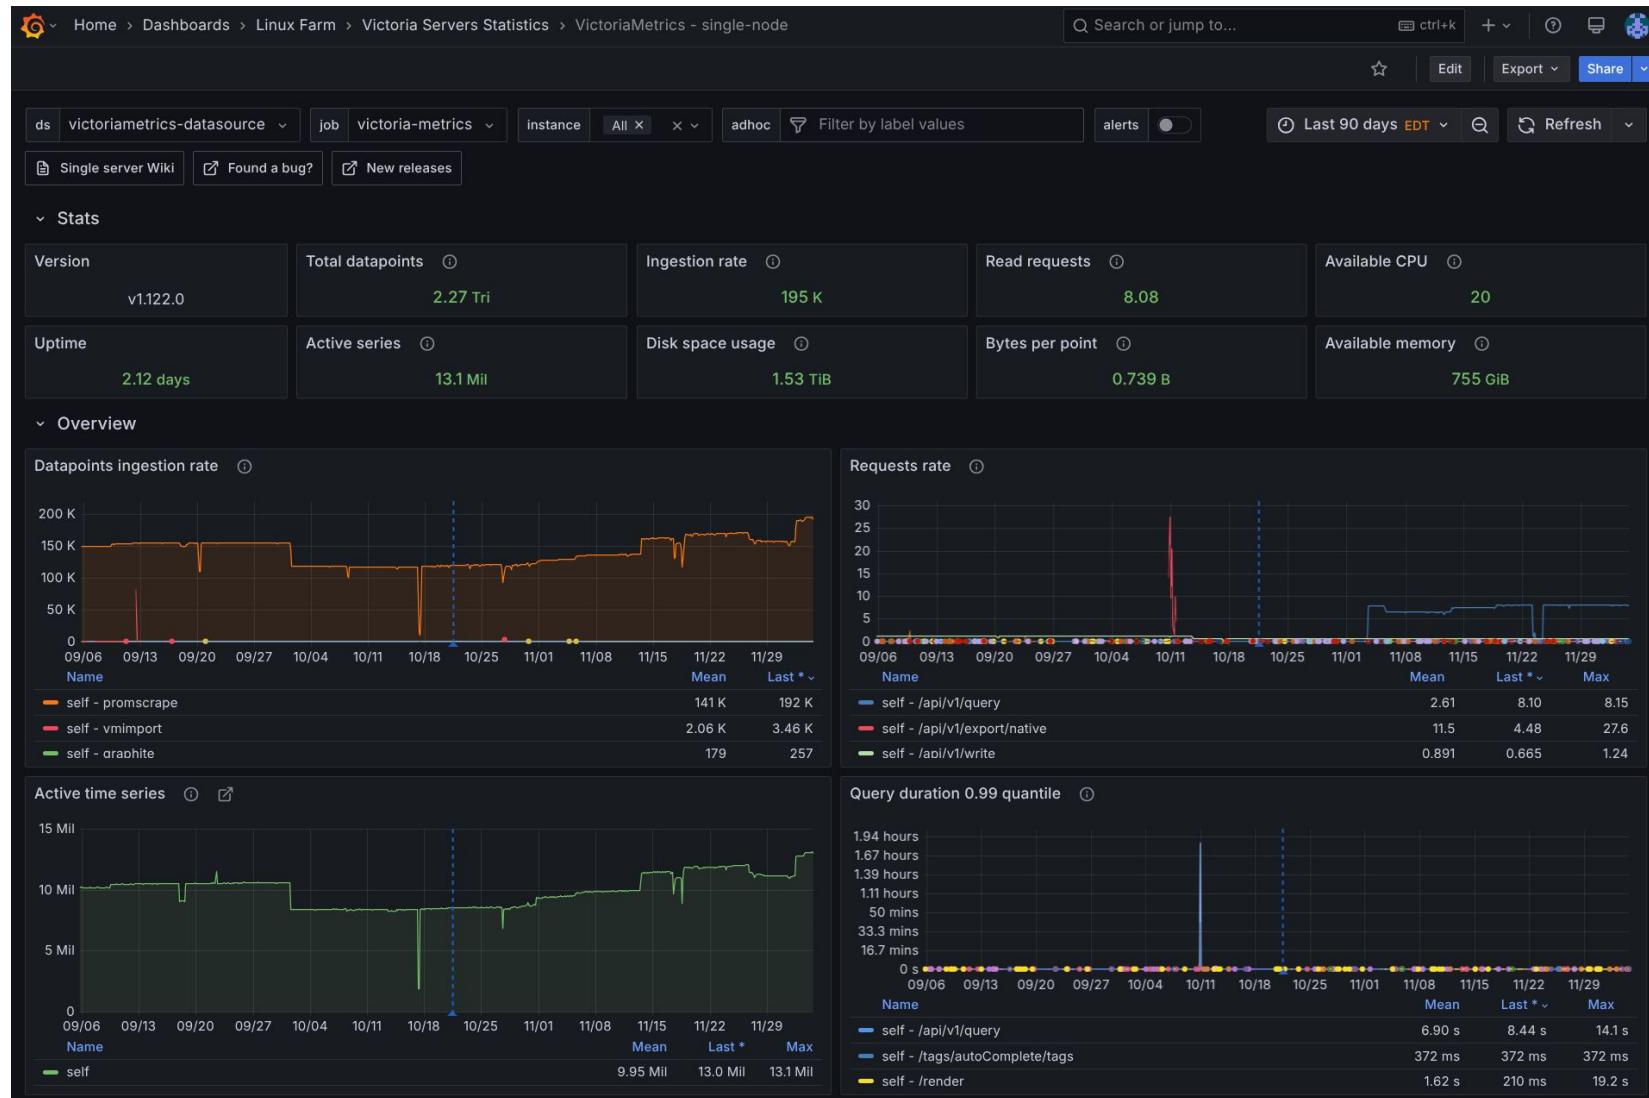
- Rebuilt legacy Logstash as development platform for OpenSearch
- AdStash tested on OpenSearch v1 (T. Smith)
- Kafka, Zookeeper nodes integrated with Logstash, OpenSearch pipelines  
(Garonne)

# Monitoring technology: Victoria Metrics

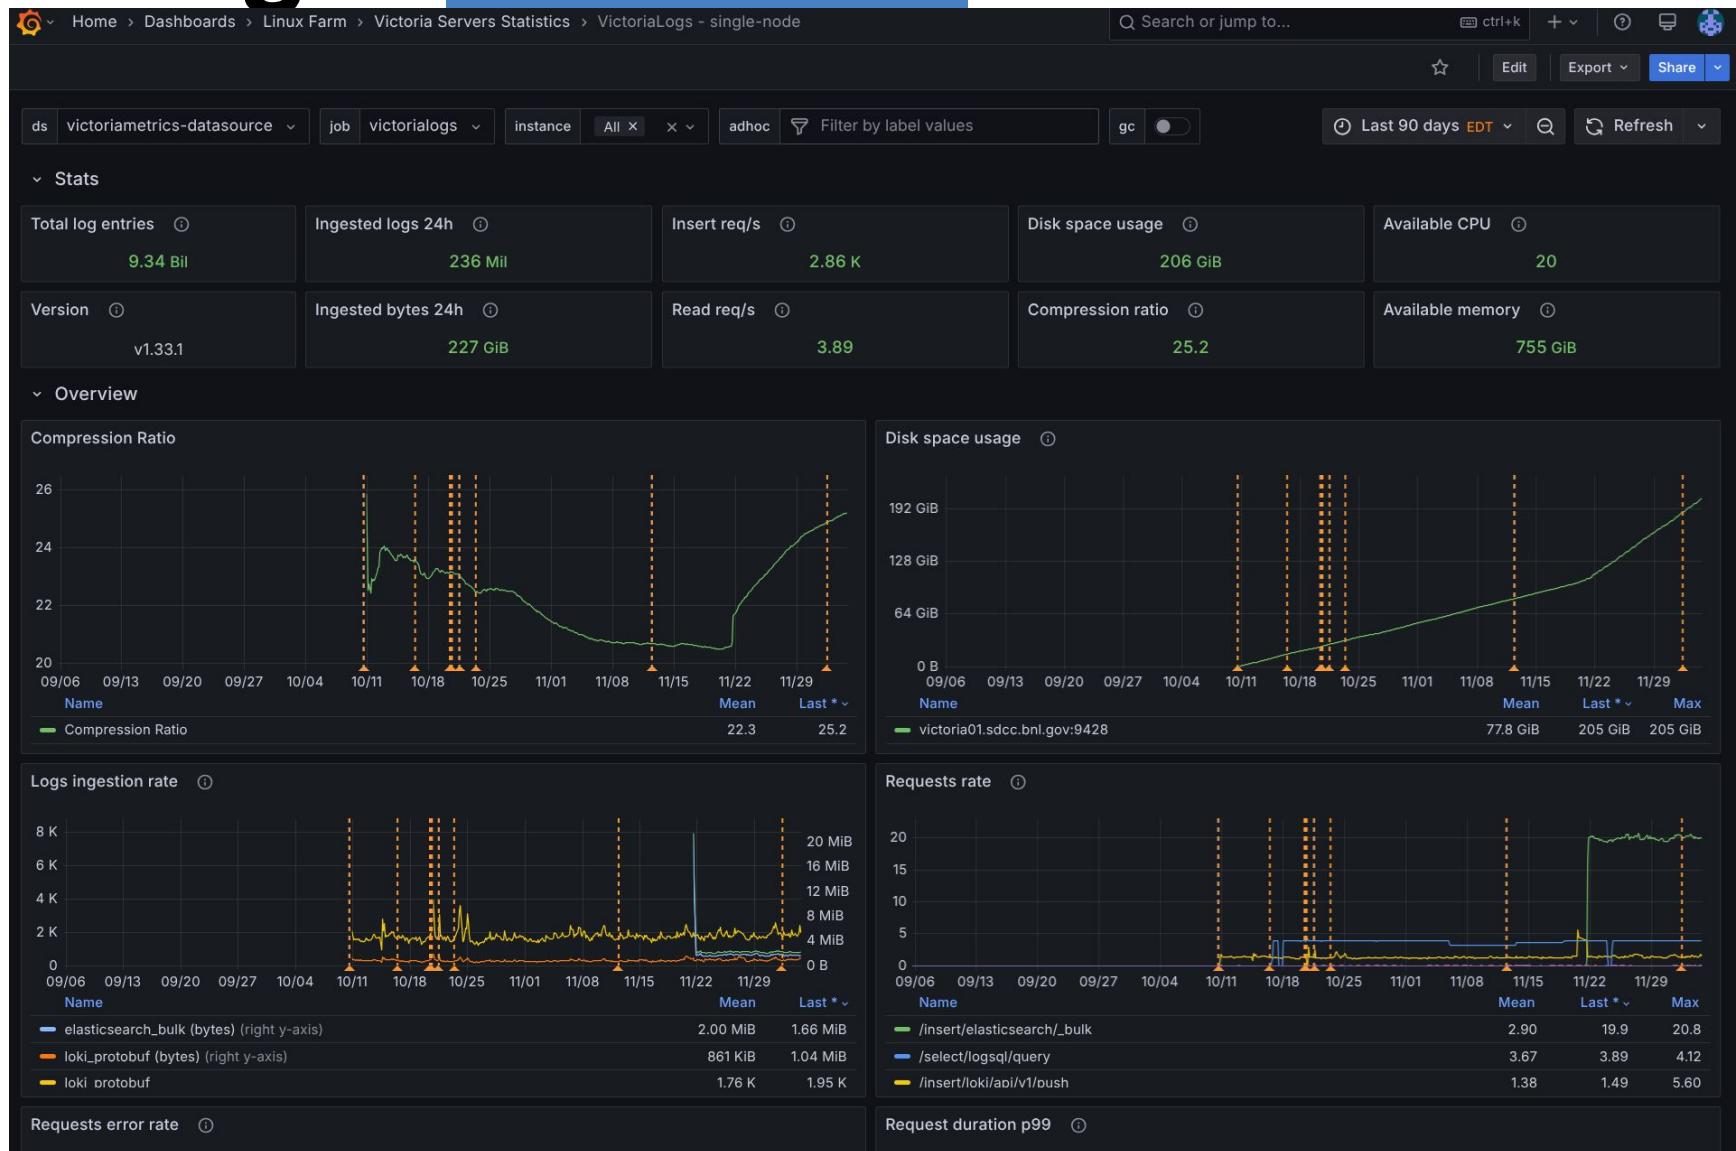
Evaluated Victoria Metrics and Victoria Logs functionality with ELK and Prometheus (Caramarcu)

- Built and evaluated services, moved to production for worker node farm (and beyond)
- Replaced collectd with node exporter and other custom exporters (e.g., CVMFS client exporter)
- Migrating Graphite/Carbon clusters to Victoria Metrics
- Replaced, created Grafana dashboards with Victoria Metrics data source
  - Includes new metrics monitoring, Grafana dashboards for NetApp storage (Frith)

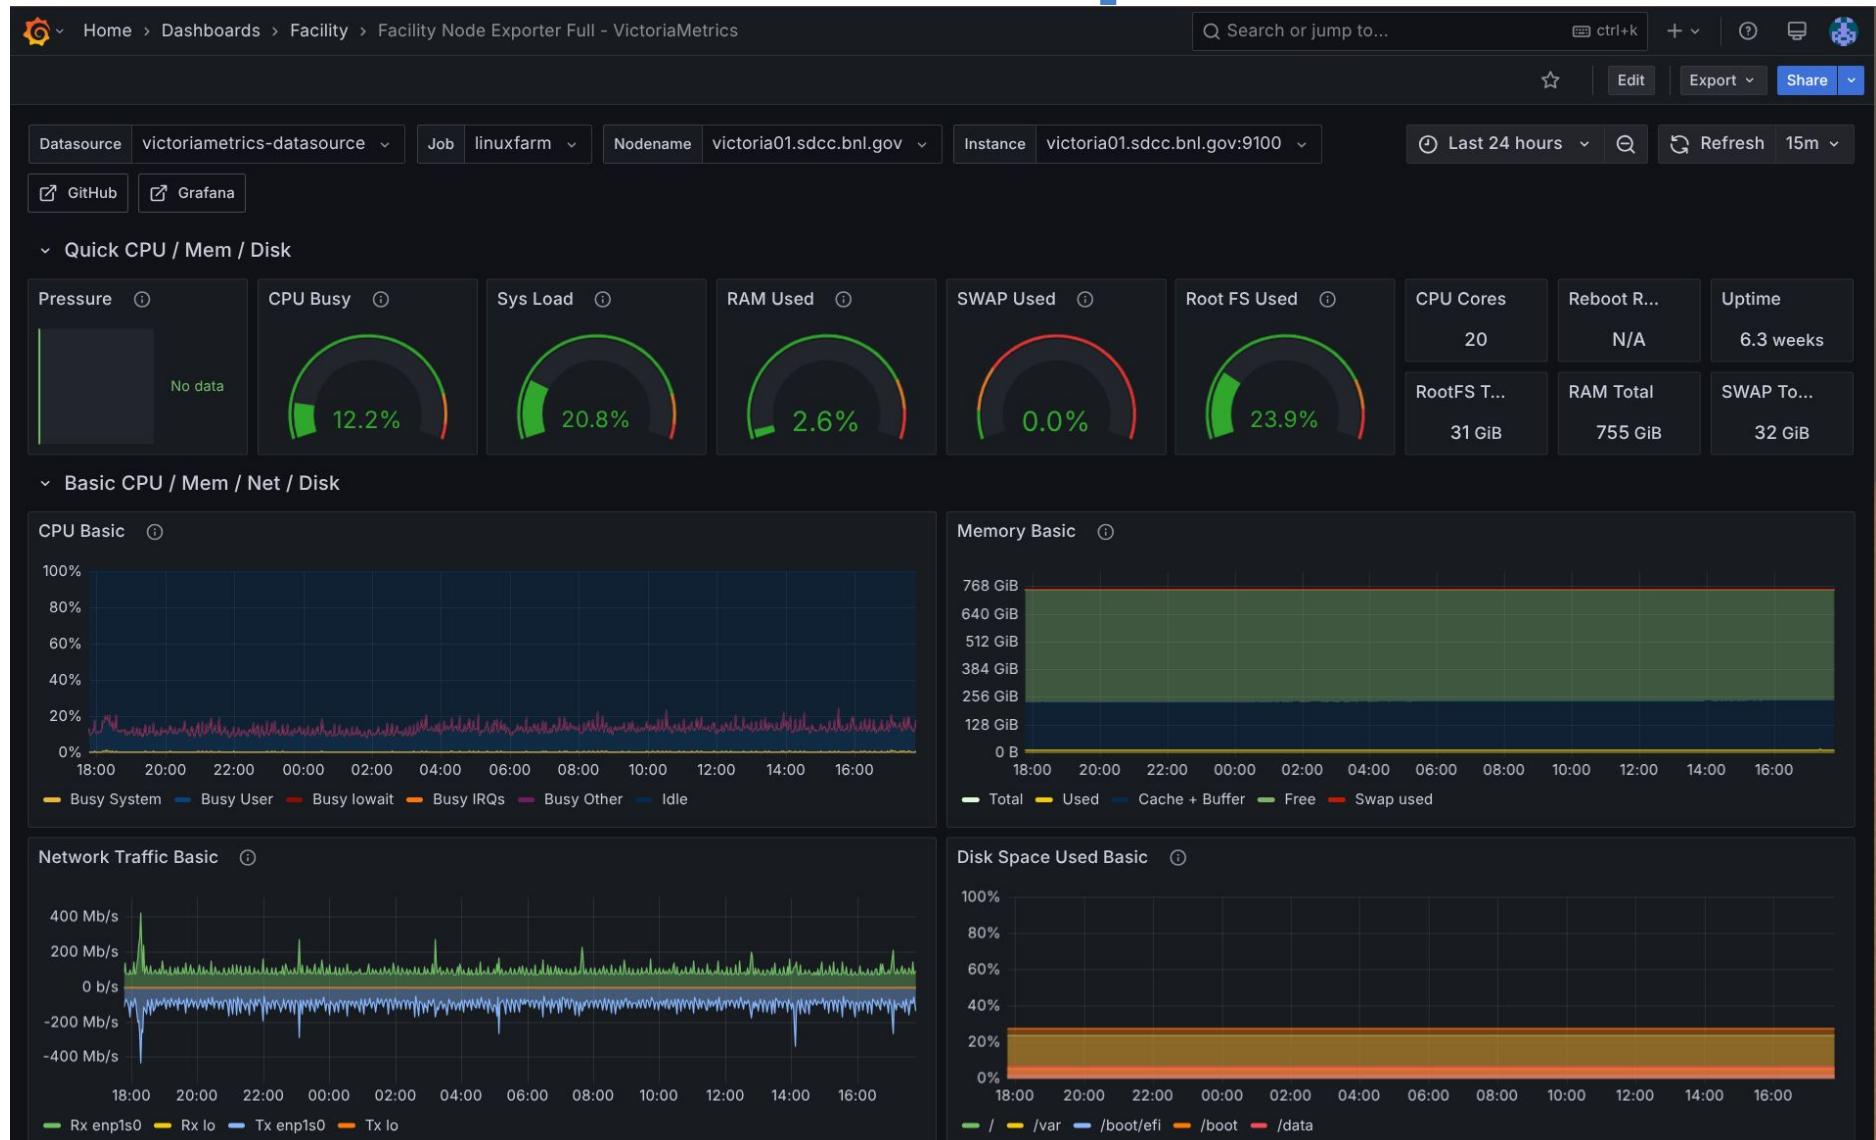
# Victoria Metrics: Statistics



# Victoria Logs: Statistics



# Victoria Metrics: Node exporter



# Victoria Logs: Dashboard

Home > Dashboards > Facility > Protected > Facility VictoriaLogs

VLdatasource: victorialogs-datasource | Host: All | keyword: \* | Filters | Filter by label values | Last 24 hours EST | Refresh

Top 20 hosts generating messages

| Host               | Count    |
|--------------------|----------|
| os-master00.sdc... | 9198477  |
| cvmfs-cache4       | 16801309 |
| cvmfs-s1c          | 3545982  |
| cvmfs-cache3       | 2287543  |
| gs-elk03           | 7853305  |
| os-master02.sdc... | 7164642  |
| idm05              | 6216273  |
| cvmfs-reverse3     | 6183365  |
| cvmfs-reverse4     | 5629683  |
| os-master01.sdc... | 5510775  |
| master1.opensh...  | 5483079  |
| ns0                | 4874129  |
| os-worker04.sdc... | 4827500  |
| varnish01.sdcc.... | 4405216  |
| webproxy02         | 4130351  |
| os-worker01.sdc... | 3802018  |
| master0.opensh...  | 3668828  |

Top 20 programs generating messages

| Program             | Count    |
|---------------------|----------|
| _time:[1764802...   | 8548979  |
| squid               | 29075821 |
| ns-slapd            | 18582791 |
| (squid-1)           | 17752535 |
| http                | 14211873 |
| named               | 11886175 |
| systemd             | 9124102  |
| metricbeat          | 7898584  |
| varnishncsa         | 4403750  |
| bash                | 4171165  |
| rucio-reaper        | 4041934  |
| baloo_file_extra... | 3698720  |
| logstash            | 3055195  |
| kernel              | 2799105  |
| rsyslogd            | 2786132  |
| hpss_ptpd           | 2336184  |
| lgetidentity        | 2326274  |

Top 20 hosts with keyword: \* in messages

| Host               | Count    |
|--------------------|----------|
| os-master00.sdc... | 9198477  |
| cvmfs-cache4       | 16801309 |
| cvmfs-s1c          | 3545982  |
| cvmfs-cache3       | 2287543  |
| gs-elk03           | 7853305  |
| os-master02.sdc... | 7164642  |
| idm05              | 6216273  |
| cvmfs-reverse3     | 6183365  |
| cvmfs-reverse4     | 5629683  |
| os-master01.sdc... | 5510775  |
| master1.opensh...  | 5483089  |
| ns0                | 4874129  |
| os-worker04.sdc... | 4827500  |
| varnish01.sdcc.... | 4405216  |
| webproxy02         | 4130351  |
| os-worker01.sdc... | 3802018  |
| master0.opensh...  | 3668828  |

Messages for All (6 hours max)

```
> 2025-12-04 17:47:31.998 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.981 {"level": "info", "ts": "2025-12-04T22:47:31Z", "l
> 2025-12-04 17:47:31.981 {"level": "info", "ts": "2025-12-04T22:47:31Z", "l
> 2025-12-04 17:47:31.981 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.975 ::ffff:202.13.202.140:55756 [04/Dec/2025:22:47
> 2025-12-04 17:47:31.974 time="2025-12-04T22:47:31Z" level=debug msg=<
> 2025-12-04 17:47:31.974 time="2025-12-04T22:47:31Z" level=debug msg=<
> 2025-12-04 17:47:31.972 time="2025-12-04T22:47:31Z" level=debug msg=>
> 2025-12-04 17:47:31.972 time="2025-12-04T22:47:31Z" level=debug msg=>
> 2025-12-04 17:47:31.972 time="2025-12-04T22:47:31Z" level=debug msg=<
> 2025-12-04 17:47:31.970 time="2025-12-04T22:47:31Z" level=debug msg=>
> 2025-12-04 17:47:31.970 time="2025-12-04T22:47:31Z" level=debug msg=>
> 2025-12-04 17:47:31.969 time="2025-12-04T22:47:31Z" level=debug msg=>
> 2025-12-04 17:47:31.963 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.963 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.960 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.960 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.956 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.956 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.956 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.952 missing _msg field; see https://docs.victoriam
> 2025-12-04 17:47:31.949 time="2025-12-04T22:47:31Z" level=debug msg=<
> 2025-12-04 17:47:31.949 time="2025-12-04T22:47:31Z" level=debug msg=<
> 2025-12-04 17:47:31.948 time="2025-12-04T22:47:31Z" level=debug msg=<
> 2025-12-04 17:47:31.945 time="2025-12-04T22:47:31Z" level=debug msg=>
> 2025-12-04 17:47:31.944 time="2025-12-04T22:47:31Z" level=debug msg=<
> 2025-12-04 17:47:31.944 time="2025-12-04T22:47:31Z" level=debug msg=<
> 2025-12-04 17:47:31.942 missing _msg field; see https://docs.victoriam
```

Top 20 hosts generating error messages

| Host               | Count   |
|--------------------|---------|
| pandaharvester...  | 1711118 |
| www                | 145049  |
| os-worker04.sdc... | 173720  |
| dqos002            | 172462  |
| os-worker06.sdc... | 134369  |
| lustremondb        | 130979  |
| acas0924.usatli... | 115594  |
| os-worker02.sdc... | 115409  |

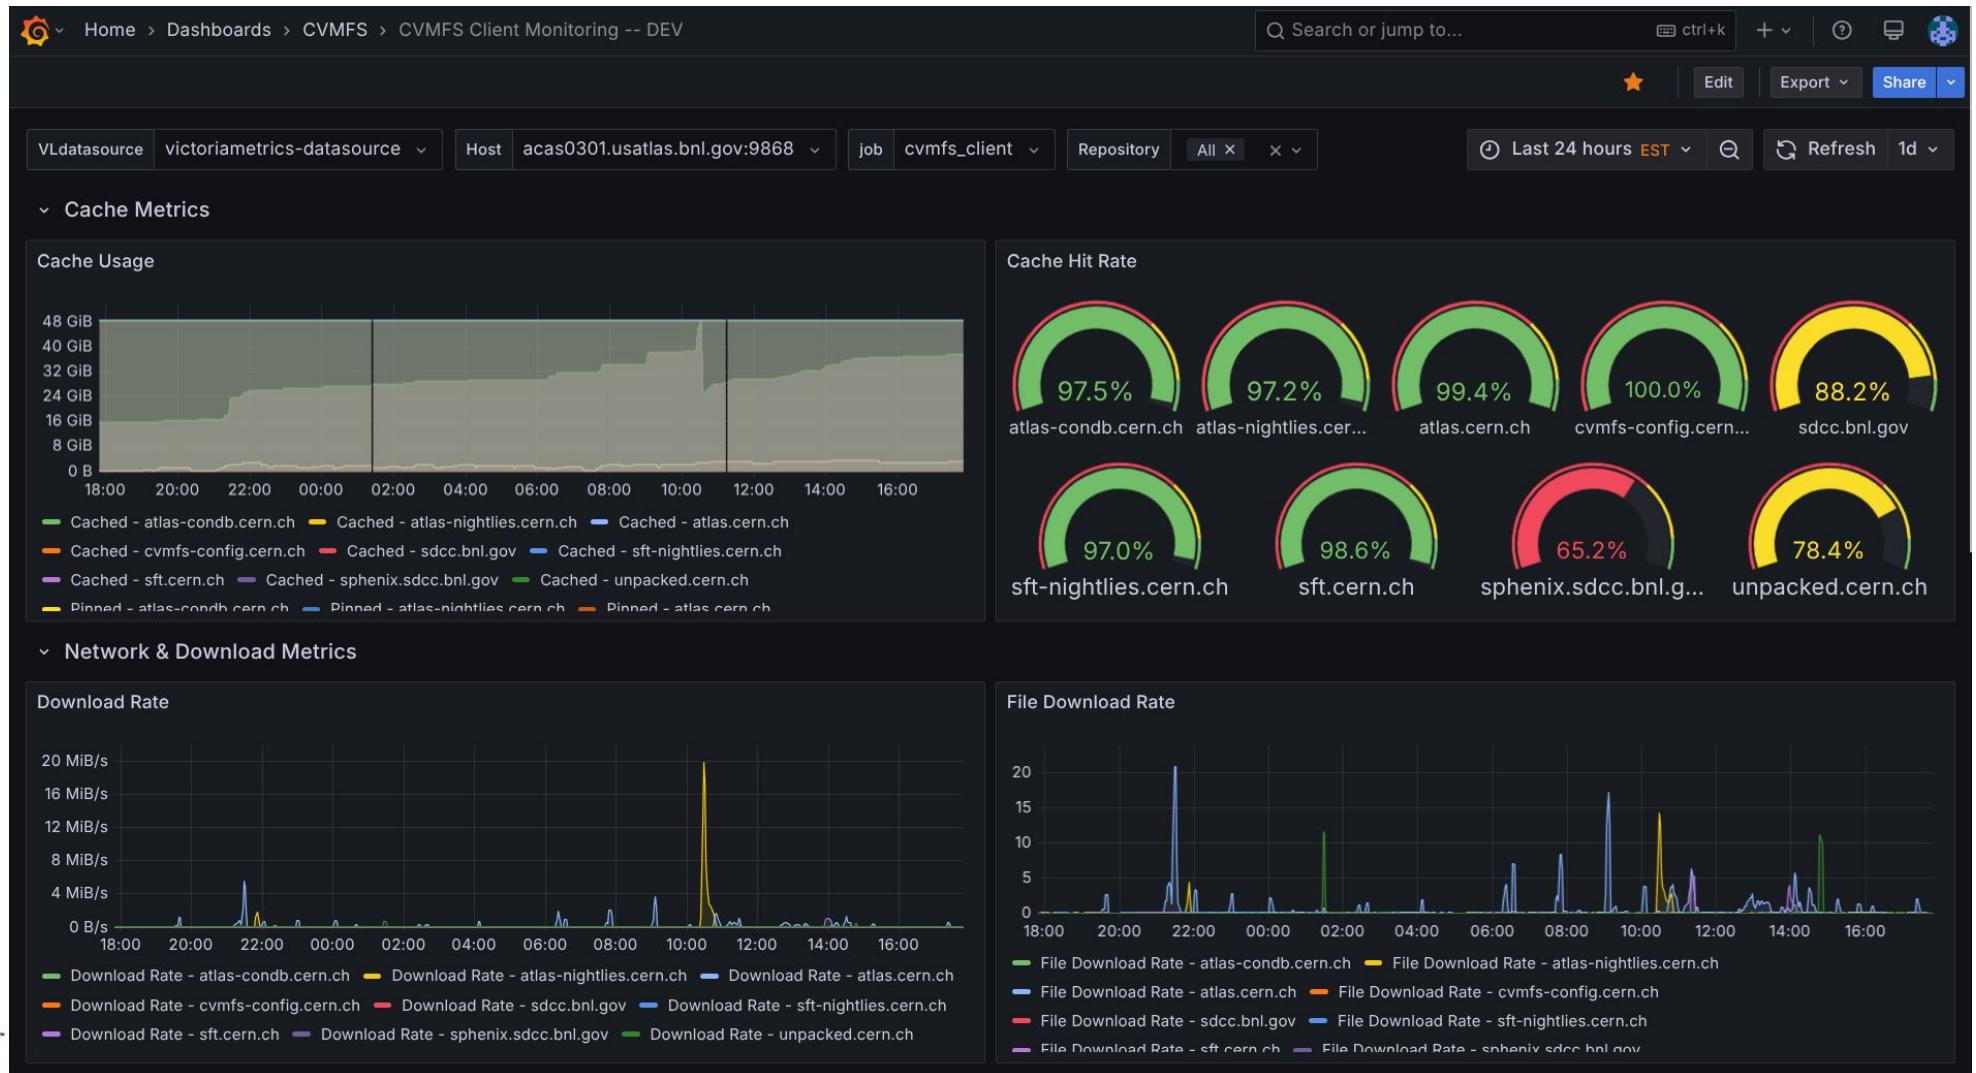
Top 20 programs generating error messages

| Program           | Count   |
|-------------------|---------|
| rsyslogd          | 1709619 |
| _time:[1764802... | 1719182 |
| kernel            | 450602  |
| httpdlog          | 346768  |
| loki-linux-amd64  | 130979  |
| (squid-1)         | 113208  |
| ssh               | 74040   |
| opensearch-da...  | 72249   |

Top 20 programs generating \* messages

| Program           | Count    |
|-------------------|----------|
| _time:[1764802... | 8551670  |
| squid             | 29075821 |
| ns-slapd          | 18582809 |
| (squid-1)         | 17752533 |
| http              | 14211874 |
| named             | 11886185 |
| systemd           | 9124103  |
| metricbeat        | 7898584  |

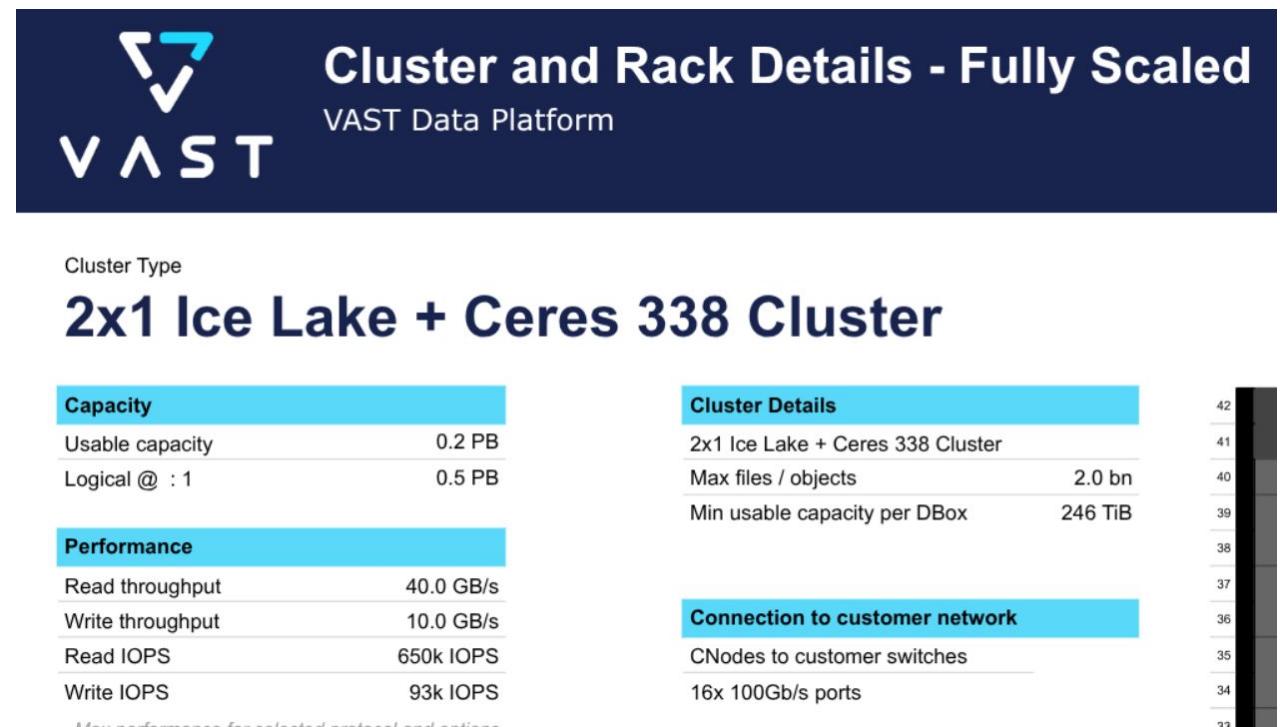
# Victoria Metrics: CVMFS client exporter



# Evaluation: Storage

## VAST storage (Hancock):

- DF-3015
- 338 TB (200 TiB licensed, 260 usable)
- Added new workflows for network block storage (NVMe over TCP)



# Evaluation: CPU

## Dell servers (Hancock):

- PowerEdge R7525 R1 (x2)
- Integrated into OpenShift cluster as worker nodes dedicated to monitoring hosts and workflows

### **PowerEdge R7525 Server**

SAS/SATA Backplane

PowerEdge R7525 Server

No Trusted Platform Module

16X 2.5 SAS/SATA with XGMI supports 6X T4 cards

AMD EPYC 7543 2.8GHz, 32C/64T, 256M Cache (225W) DDR4-3200

AMD EPYC 7543 2.8GHz, 32C/64T, 256M Cache (225W) DDR4-3200

Heatsink for 2 CPU configuration (CPU greater than or equal to 180W)

Performance Optimized

3200MT/s RDIMMs

24x 64GB RDIMM, 3200MT/s, Dual Rank, 16Gb

C7, Unconfigured RAID for HDDs or SSDs (Mixed Drive Types Allowed)

PERC H755 Adapter FH

2x 960GB SSD SATA Read Intensive 6Gbps 512e 2.5in Hot-plug AG Drive, 1 DWPD

3x 3.84TB SSD SATA Read Intensive 6Gbps 512e 2.5in Hot-plug AG Drive, 1 DWPD

# Monitoring technology: HPSS

## Modernization of HPSS web infrastructure (Undrus)

- New RHEL 8 host baseline
- Migration from PHP 5 to PHP 8
- Code refactoring & cleanup
  - Current code size: 211,000 PHP and 344,000 HTML
  - Remove deprecated APIs, modernize language constructs, and introduce stricter typing, error handling
- Security tuning to align configuration with best practices

# Monitoring technology: HPSS

## Preparing the next-gen HPSS Web Platform

This testbed is dedicated to the preparation and deployment of the **Upgrade & Modernize HPSS Web Infrastructure** project, including migration to **RHEL 8**, migration from **PHP 5** to **PHP 8**, and associated code refactoring and hardening.

● RHEL 8 migration   ● PHP 8 upgrade   ● Legacy PHP 5 decommissioning

● Code refactoring & testing

 Why PHP 8 vs PHP 5?

 Open New HPSS Web

### ENVIRONMENT

● Isolated testbed – safe for experimentation

### SCOPE

RHEL 8 adoption, PHP 8 runtime, dependency refresh, and systematic refactoring of the HPSS web codebase to remove PHP 5 patterns and legacy assumptions.

### UPGRADE WORKSTREAMS

RHEL 8 · PHP 8 · Refactor

#### RHEL 8 baseline

New OS image, updated libraries, modern TLS/OpenSSL defaults, and hardened defaults for HPSS web workloads.

#### PHP 8 migration

Replace PHP 5 with PHP 8, align extensions, and validate compatibility with existing HPSS web modules.

#### Code refactoring & cleanup

Remove deprecated APIs, modernize language constructs, and introduce stricter typing, error handling, and test coverage.

#### Performance & security tuning

Benchmark critical HPSS paths, enable PHP 8 Just-In-Time (JIT) compiler where appropriate, and align configuration with security best practices.

 Safe space for breaking changes & rollbacks

 Security-driven modernization

 Measured performance impact

### TARGET STACK

Operating system

PHP runtime

Web tier

Application code

Modernized HPSS Web Platform

RHEL 8 ← RHEL 6 (legacy)

PHP 8.x ← PHP 5.x (current HPSS Web)

Apache (tuned for PHP-FPM)

Refactored HPSS web modules with reduced legacy & improved testability

All experiments here are isolated from production.

# Monitoring: Plans & Goals

# To do:

- Continue deployment and configuration of key services (Victoria Metrics, Prometheus, OpenSearch)
  - Rebuild OpenSearch, Logstash instances on VAST storage
  - Prometheus: complement Victoria Metrics, or merge?
- Evaluate VAST storage:
  - compare performance with other NAS solutions (NetApp)
  - compare storage types (e.g., block vs. image storage)
- Evaluate potential AI use cases and solutions:
  - Anomaly detection (e.g., [vmanomaly](#)), predictive analysis, observability, natural language
- Continue modernization of HPSS monitoring
  - Update web services, consolidate and update background processes, upgrade backend database
- Convert Loki log output to Victoria Logs (Liu)
- (Potential) Dashboard reorganization and clean-up, data source consolidation (Yang)