

Monitoring

Petter A. Urkedal

2022-06-21

News Since Last Meeting

- ▶ Tarpool node exporter.
- ▶ Prometheus TLS and X.509 auth
- ▶ Prometheus PushProx
- ▶ VictoriaMetrics
- ▶ Adjusting some Ganglia scripts to serve Prometheus
- ▶ Production setup running in Ørestaden

Grafana Upgrades are Problematic

After (automatic) upgrade to Grafana 9:

- ▶ Grafana new uses a time format (like “2m1”) for `$_rate_interval` which is yet unimplemented in our Prometheus server.
- ▶ Grafana requires a more recent Elasticsearch.

VictoriaMetrics

- ▶ Provides long-term storage for Prometheus time-series data
- ▶ Compact storage and fast queries. Experience with our setup:
 - ▶ Data from our three Prometheus instances currently consumes 19 GB. The TSDB on the production instance uses 29 GB. TimescaleDB, when we tried it, used about 10 times as much storage.
 - ▶ Query speed only slightly worse than TSDB and much better than TimescaleDB.
 - ▶ Significantly less load on the VM than TimescaleDB.
- ▶ Collects metrics from our three Prometheus instances in one place.

The Tarpool Node Exporter Setup

We are hopefully down two options:

1. We install a node exporter in the dCache account listening on port 19100.
2. Sites which already runs the node exporter can allow us to use it.

We had a third option based on PushProx, which I will be happy to rip out of our Ansible role.

The first option is secured with TLS and certificate authentication. For the second option, at least, the endpoint should be protected by firewall rules.

Prometheus Exporters

Configured exporters:

- ▶ node-exporter with some custom node collectors:
 - ▶ certificate expiration
 - ▶ Ansible inventory groups (useful for joins)
 - ▶ Ganeti DRBD performance metrics from `ganeti-mond`.
 - ▶ tape usage per site and experiment from Chimera
 - ▶ dCache billing log¹ or we can use Elasticsearch.
- ▶ dcache-exporter ([project page][1], [clone][2]).
- ▶ HAProxy build-in exporter
- ▶ postgres
- ▶ (Other: apache, bind, process, prometheus)

¹Extending `dcache-billing-stat.rb` on piggy

Prometheus Metrics for dCache

- ▶ dCache performance and reliability
 - ▶ OS level metrics (CPU, memory, disk, networking) can be used to some extent.
 - ▶ The dcache-exporter provides only three metrics of limited use which are directly performance-related (door load, domain event queue size, and domain thread count).
 - ▶ We have request times and error rates from billing and tapesum.
 - ▶ HAProxy.
- ▶ dCache storage
 - ▶ Provided by the dcache-exporter
 - ▶ Tape storage provided by the Chimera collector
 - ▶ Billing logs