

OpenMetrics

Standing on the shoulders of Titans

Richard Hartmann,
RichiH@{freenode,OFTC,IRCnet},
richih@{fosdem,debian,richih}.org,
@TwitchiH

2018-12-11

‘whoami‘

- Richard "RichiH" Hartmann
- Swiss army chainsaw at SpaceNet
 - Currently responsible for building one of the most modern datacenters in Europe
 - ...and always looking for nice co-workers in the Munich area
- FOSDEM, DebConf, DENOGx, PromCon staff
- Author of `https://github.com/RichiH/vcsh`
- Debian Developer
- Prometheus team member



What's Prometheus?

You can't talk about OpenMetrics without mentioning Prometheus



What's Prometheus?

Show of hands: Who has heard of Prometheus?

Before Prometheus

- Historically, the monitoring landscape has been highly fragmented
- Many solutions based on ancient technology
- Most data formats are proprietary, hard to implement, or both
- ..and most with hierarchical data models
- Almost none with a focus on metrics

Before Prometheus

- Only thing approaching a standard: SNMP
- Who has heard of SNMP?
- Who likes SNMP?
- XKCD 927 applies

After Prometheus

- Prometheus has become a de-facto standard in cloud-native metric monitoring
- Ease of exposing data has lead to an explosion in compatible metrics endpoints
- Prometheus' exposition format is based on a lot of operational experience, but has been designed between few people
- Some other projects & vendors are torn about adopting something from a “competing” product
- Especially traditional vendor prefer to support official standards

Goldrush

- So about fragmentation...
- Cloud native is in the buzz phase
- Explosive growth always equals goldrush and stake claiming, aka fragmentation
- I don't have a profit motive, so I can easily decide to try and consolidate this space without profit loss
 - ..so that's what I decided to do

Acknowledgements

- Main work has been done by
 - Prometheus team
 - Ben Kochie
 - Brian Brazil
 - myself
 - Google
 - Sumeer Bhola
 - Uber
 - Jerome Froelich
 - Rob Skillington

First commitments, too many for full list

- Cloudflare
- CNCF at large
- GitLab
- Google
- Grafana
- InfluxData
- Prometheus ;)
- RobustPerception
- SpaceNet
- Uber

Approach

- Spin out Prometheus exposition format into its own standard
 - Neutral brand
 - Umbrella for wider cooperation in FLOSS and industry
- Broadening designer base ensures covering of many/all use cases
- Official standard (work in progress)
 - IETF RFC
 - Registered content-type/MIME-type
 - Maybe IANA port assignment

Current status

- We finished the discussion phase and moved into implementation
 - Still finding minor bugs during implementation
- RFC currently blocked on me finding time
- Prometheus: experimental support since 2.5.0
- Python client library: experimental support since 0.4.0
- Test your own OM output:
`robustperception.io/checking-openmetrics-output-is-valid`
 - Google and Uber want to create another reference parser to weed out bugs

Next steps

- Full OpenMetrics support in Prometheus, InfluxDB, OpenCensus, M3DB, etc
- Spreading the word
- CNCF sandbox to incubating

Beyond cloud

- OpenMetrics is intended to go beyond "just" cloud-native fields
- Want to get more traditional projects and vendors on board
 - Arista already on board
 - Vertiv very interested

Beyond metrics

- OpenMetrics supports more than just metrics
- Every single data point in a time series can point to one single event
- Especially useful if you emit one trace id per histogram bucket, i.e. exemplars
- Some integrations already support this concept, e.g. OpenCensus
- Ingestors are free to discard this optional data, e.g. Prometheus

Bringing observability back together

- IT architectures breathe
- Out/Insourcing, Mainframe/client-server/containers, etc
- You solve the current problem, and then face another
- By focussing on Metrics, Prometheus breathed in
- OpenMetrics intends to breathe out by creating one consistent user story encompassing all events in the next version
 - OpenStats - OpenMetrics for push
 - Events - Logs, Traces, Coredumps, whatever

Prometheus

```
http_requests_total{env="prod",method="post",code="200"} 1027
http_requests_total{env="prod",method="post",code="400"} 3
http_requests_total{env="prod",method="post",code="500"} 12
http_requests_total{env="prod",method="get",code="200"} 20
http_requests_total{env="test",method="post",code="200"} 372
http_requests_total{env="test",method="post",code="400"} 75
```

OpenMetrics

```
http_requests_total{env="prod",method="post",code="200"} 1027
http_requests_total{env="prod",method="post",code="400"} 3
http_requests_total{env="prod",method="post",code="500"} 12
http_requests_total{env="prod",method="get",code="200"} 20
http_requests_total{env="test",method="post",code="200"} 372
http_requests_total{env="test",method="post",code="400"} 75
```

OpenMetrics

Prometheus

```
http_requests_total{env="prod",method="post",code="200"} 1027 1544554800000
```

OpenMetrics

```
http_requests_total{env="prod",method="post",code="200"} 1027 1544554800
```

```
histogram_bucket{le="1"} 123 # {foo="bar"} 42 1544554800
```

```
histogram_bucket{le="2"} 234 # {foo="bar"} 23 1544554799.123
```

```
histogram_bucket{le="3"} 345 1544554800 # {foo="bar"} 11 1544554799
```

Further reading

- Prometheus 2017 Dev Summit: <https://docs.google.com/document/d/1DaHFao0saZ3MDt9yuuxLaCQg8WGad08s44i3cxSARcM/edit>
- OpenMetrics: <https://github.com/OpenObservability/OpenMetrics>
- OpenMetrics mailing list: <https://groups.google.com/forum/#\protect\kern-.1667em\relaxforum/openmetrics>
- This and other talks: <https://github.com/RichiH/talks/>

Thanks!

Thanks for listening!

Questions?

Email me if you want a job in Munich.

See slide footer for contact info.