



Prometheus Best Practices and Beastly Pitfalls

Julius Volz, August 17, 2017



Areas

- Instrumentation
- Alerting
- Querying

Instrumentation

What to Instrument

- Every component (including libraries)
- Spread metrics liberally (like log lines)
- "USE Method" (for resources like queues, CPUs, disks...)
Utilization, Saturation, Errors
<http://www.brendangregg.com/usemethod.html>
- "RED Method" (for things like endpoints)
Request rate, Error rate, Duration (distribution)
<https://www.slideshare.net/weaveworks/monitoring-microservices>



Metric and Label Naming

- No enforced server-side typing and units
- **BUT! Conventions:**
 - Unit suffixes
 - Base units (`_seconds` vs. `_milliseconds`)
 - `_total` counter suffixes
 - either `sum()` or `avg()` over metric should make sense
 - See <https://prometheus.io/docs/practices/naming/>

Label Cardinality

- Every unique label set: one series
- Unbounded label values will blow up Prometheus:
 - public IP addresses
 - user IDs
 - SoundCloud track IDs (*ehem*)

Label Cardinality

- Keep label values well-bounded
- Cardinalities are multiplicative
- What ultimately matters:
 - **Ingestion:** total of a couple million series
 - **Queries:** limit to 100s or 1000s of series
- Choose metrics, labels, and #targets accordingly

Errors, Successes, and Totals

Consider two counters:

- failures_total
- successes_total

What do you actually want to do with them?

Often: **error rate ratios!**

Now complicated:

```
rate(failures_total[5m])  
/  
(rate(successes_total[5m]) + rate(failures_total[5m]))
```

Errors, Successes, and Totals

⇒ Track **failures and total requests**, not **failures and successes**.

- failures_total
- requests_total

Ratios are now simpler:

```
rate(failures_total[5m]) / rate(requests_total[5m])
```

Missing Series

Consider a labeled metric:

```
ops_total{optype="<type>"}
```

Series for a given "type" will only appear once something happens for it.



Missing Series

Query trouble:

- `sum(rate(ops_total[5m]))`
⇒ empty result when **no** op has happened yet
- `sum(rate(ops_total{optype="create"}[5m]))`
⇒ empty result when no “create” op has happened yet

Can break alerts and dashboards!

Missing Series

If feasible:

Initialize known label values to 0. In Go:

```
for _, val := range opLabelValues {  
    // Note: No ".Inc()" at the end.  
    ops.WithLabelValues(val)  
}
```

Client libs automatically initialize label-less metrics to 0.

Missing Series

Initializing not always feasible. Consider:

```
http_requests_total{status="<status>"}
```

A `status=~"5.."` filter will break if no `5xx` has occurred.

Either:

- Be aware of this
- Add missing label sets via `or` based on metric that exists (like `up`):

```
<expression> or up{job="myjob"} * 0
```

See <https://www.robustperception.io/existential-issues-with-metrics/>



Metric Normalization

- Avoid non-identifying extra-info labels

Example:

```
cpu_seconds_used_total{role="db-server"}
```

```
disk_usage_bytes{role="db-server"}
```

- Breaks series continuity when role changes
- Instead, join in extra info from separate metric:

<https://www.robustperception.io/how-to-have-labels-for-machine-roles/>



Alerting



General Alerting Guidelines

Rob Ewaschuk's ["My Philosophy on Alerting"](#) (Google it)

Some points:

- Page on user-visible symptoms, not on causes
 - ...and on immediate risks ("disk full in 4h")
- Err on the side of fewer pages
- Use causal metrics to answer **why** something is broken



Unhealthy or Missing Targets

Consider:

ALERT HighErrorRate

IF rate(errors_total{job="myjob"}[5m]) > 10

FOR 5m

Congrats, amazing alert!

But what if **your targets are down or absent in SD?**

⇒ empty expression result, no alert!



Unhealthy or Missing Targets

⇒ Always have an up-ness and presence alert per job:

(Or alert on up ratio or minimum up count).

ALERT MyJobInstanceDown

IF up{job="myjob"} == 0

FOR 5m

ALERT MyJobAbsent

IF absent(up{job="myjob"})

FOR 5m



FOR Duration

Don't make it too short or missing!

ALERT InstanceDown

IF up == 0

Single failed scrape causes alert!



FOR Duration

Don't make it too short or missing!

ALERT InstanceDown

IF up == 0

FOR 5m



FOR Duration

Don't make it too short or missing!

```
ALERT MyJobMissing  
  IF absent(up{job="myjob"})
```

Fresh (or long down) server may immediately alert!

FOR Duration

Don't make it too short or missing!

```
ALERT MyJobMissing  
  IF absent(up{job="myjob"})  
  FOR 5m
```



FOR Duration

⇒ Make this at least 5m (usually)

FOR Duration

Don't make it too long!

```
ALERT InstanceDown
```

```
  IF up == 0
```

```
  FOR 1d
```

No **FOR** persistence across restarts! ([#422](#))



FOR Duration

⇒ Make this at most 1h (usually)

Preserve Common / Useful Labels

Don't:

ALERT HighErrorRate

IF sum(rate(...)) > x

Do (at least):

ALERT HighErrorRate

IF sum by(job) (rate(...)) > x

Useful for later routing/silencing/...



Querying



Scope Selectors to Jobs


- Metric name has single meaning only within one binary (job).
- Guard against metric name collisions between jobs.
- ⇨ Scope metric selectors to jobs (or equivalent):

Don't: `rate(http_request_errors_total[5m])`

Do: `rate(http_request_errors_total{job="api"}[5m])`


Order of rate() and sum()

Counters can reset. `rate()` corrects for this:




Order of rate() and sum()

`sum()` before `rate()` masks resets!



Order of rate() and sum()

`sum()` before `rate()` masks resets!



Order of rate() and sum()

⇒ Take the sum of the rates, not the rate of the sums!

(PromQL makes it hard to get wrong.)



Thanks!