Implementing a Cooperative Multi-Tenant Capable Prometheus

Jonas Große Sundrup

10.08.2018

PromCon 2018

Target users

Users:

- run small-scale infrastructure
- aren’t monitoring experts

Goal: Share monitoring infrastructure

Goal

Drop-dead-simple fire-and-forget low-resource monitoring/alerting solution that just works

Roadmap

1. Prometheus
2. Alertmanager
3. System Architecture
4. Additional Services

Requirements

- One Prometheus per machine
- Multi-tenancy
- Ansible-compatibility
- patch-free operation

Getting data into Prometheus

Deploying scrape targets

```
~/scrptarg
├── node1.yml
├── node2.yml
├── db.yml
└── mail.yml
```

Deploying scrape targets

```
~/scrptarg
├── node1.yml
├── node2.yml
├── db.yml
└── mail.yml
```

```
- job_name: alice

  file_sd_configs:
  - files:
    - /home/alice/scrptarg/*.json
    - /home/alice/scrptarg/*.yml
    - /home/alice/scrptarg/*.yaml
    refresh_interval: 5m

  scheme: https
  basic_auth:
    username: prometheus
    password: <secret>
```

Deploying rules

```
~/rules
├── normal.yml
└── the-apocalypse.yml
```

Getting data out of Prometheus

up

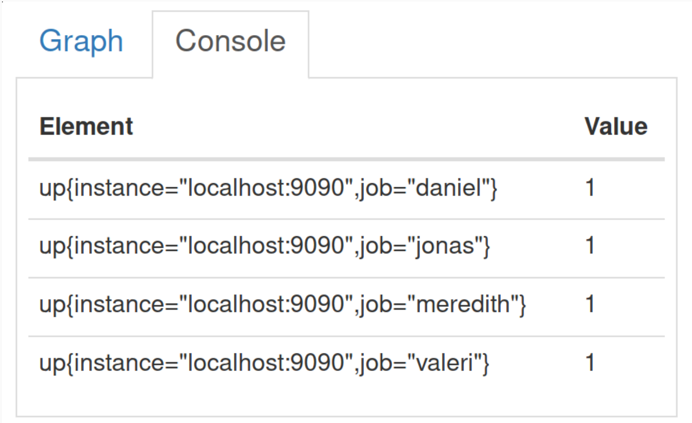

Element	Value
up{instance="localhost:9090",job="daniel"}	1
up{instance="localhost:9090",job="jonas"}	1
up{instance="localhost:9090",job="meredith"}	1
up{instance="localhost:9090",job="valeri"}	1

up{job="jonas"}

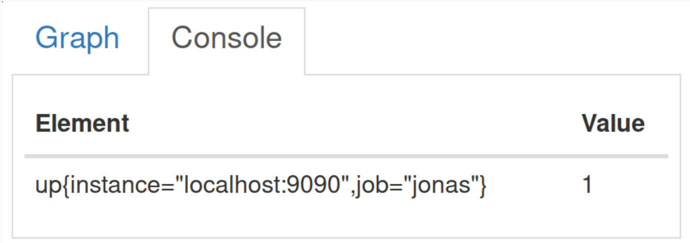

Element	Value
up{instance="localhost:9090",job="jonas"}	1

up{job="jonas"}

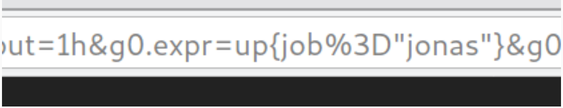

```
out=1h&g0.expr=up{job%3D"jonas"}&g0
```

Injecting a label

- `your_metric{job="hello",instance="mybox"}`

Injecting a label

- `your_metric{job="hello",instance="mybox"}`
- `your_metric{instance="mymachine"}`

Injecting a label

- `your_metric{job="hello",instance="mybox"}`
- `your_metric{instance="mymachine"}`
- `your_metric/my_metric{code="42"}`

Injecting a label

- `your_metric{job="hello",instance="mybox"}`
- `your_metric{instance="mymachine"}`
- `your_metric/my_metric{code="42"}`
- `avg_over_time(your_metric[1h]) offset 1d`

A PromQL-parser

Retrieving Data in PromQL

- `sum(http_requests_total{code="200"})`
- `avg_over_time(your_fancy_metric[1h])`

Rules

```
rules:
- alert: Endpoint down
  expr: up == 0
  for: 10m
  annotations:
    summary: "Must be short, slide isn't wide"
```

Rules

```
rules:
- alert: jonas: Endpoint down
  expr: up{job="jonas"} == 0
  for: 10m
  labels:
    job: jonas
  annotations:
    summary: "Must be short, slide isn't wide"
```

Alertmanager

Alertmanager

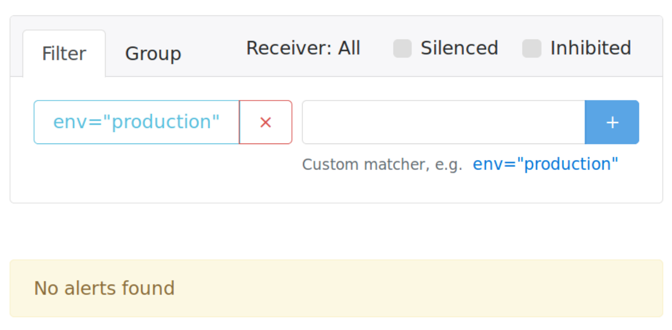

env="production" ×

Custom matcher, e.g. env="production"

No alerts found

Alertmanager

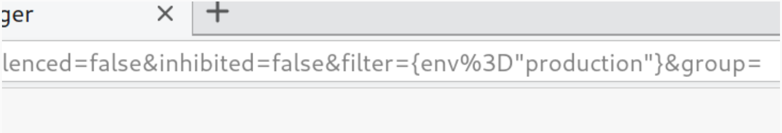

lenced=false&inhibited=false&filter={env%3D"production"}&group=

Silences

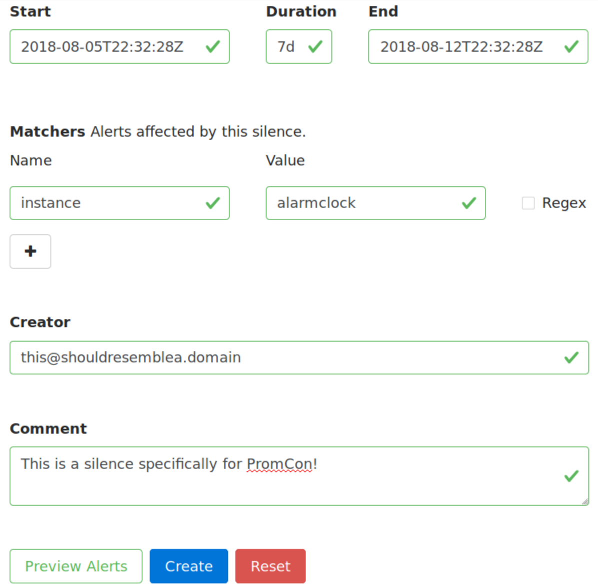

Start 2018-08-05T22:32:28Z ✓

Duration 7d ✓

End 2018-08-12T22:32:28Z ✓

Matchers Alerts affected by this silence.

Name	Value	
instance ✓	alarmclock ✓	□ Regex

Creator this@shouldresemblea.domain ✓

Comment This is a silence specifically for PromCon! ✓

Architecture overview

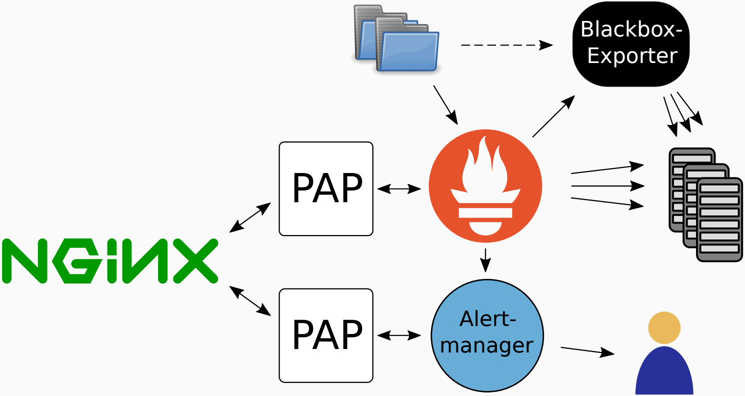

Blackbox-Exporter

Blackbox-Targets

```
~/blackboxtargets/
├── http_2xx_ipv4
│   └── websites.yml
└── tcp_connect_ipv6
    └── tcp.yml
```

Blackbox-Exporter: Configuration

```
- job_name: jonas-blackbox-http_2xx-ipv4
  params:
    module:
    - http_2xx_ipv4
  metrics_path: /probe
  file_sd_configs:
  - files:
    - /home/jonas/blackbox/http_2xx_ipv4/*.json
    - /home/jonas/blackbox/http_2xx_ipv4/*.yml
    - /home/jonas/blackbox/http_2xx_ipv4/*.yaml
    refresh_interval: 5m
  relabel_configs:
    ...
```

Blackbox-Exporter: Relabelling

```
relabel_configs:
  ...
- target_label: job
  replacement: jonas
  action: replace
  ...
```

Blackbox-Exporter: Relabelling

```
relabel_configs:
  ...
- target_label: blackbox_module
  replacement: http_2xx
  action: replace
- target_label: ip_version
  replacement: ipv4
  action: replace
```

Conclusion

Features:

- User separation: yes
- low memory profile: yes
- Ansible-compatibility: yes
- Alerting: yes
- Ease of use: yes

Limitations:

- No resource isolation
- Only one set of target credentials
- Preselected Featureset

`https://github.com/cherti/promauthproxy`

jonas@grosse-sundrup.com