Using Loki for Simplifying the Usage of Shasta Logs

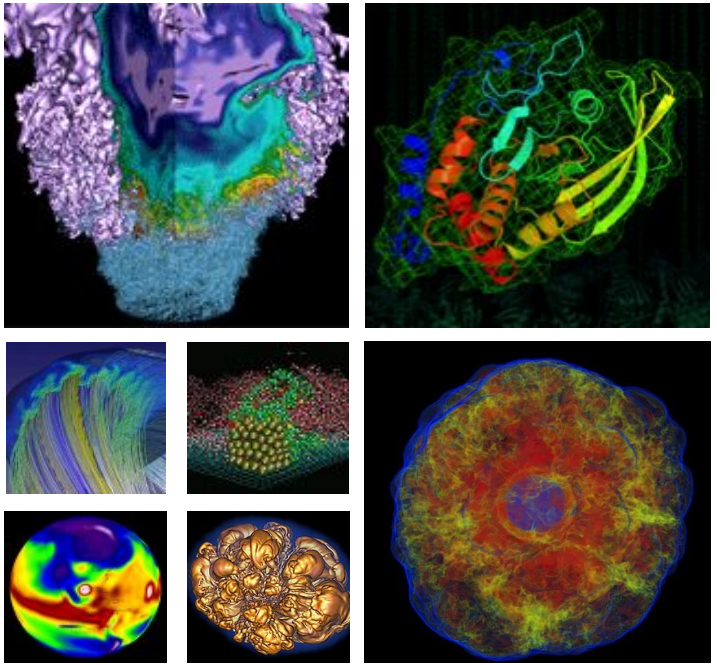

Siqi Deng
siqideng@lbl.gov
SRE@NERSC Operations

04/15/2022

Perlmutter

Phase 1:
1536 compute nodes with
6,000+ NVIDIA A100 GPU
100+ non-compute nodes

70.9 Pflop/s
No. 5 in the Top500 list in November, 2021

Phase 2:
3,000+ CPU-only compute nodes in 2022

Nagios®

elastic

VICTORIA METRICS

Grafana

Alertmanager

Data pipeline

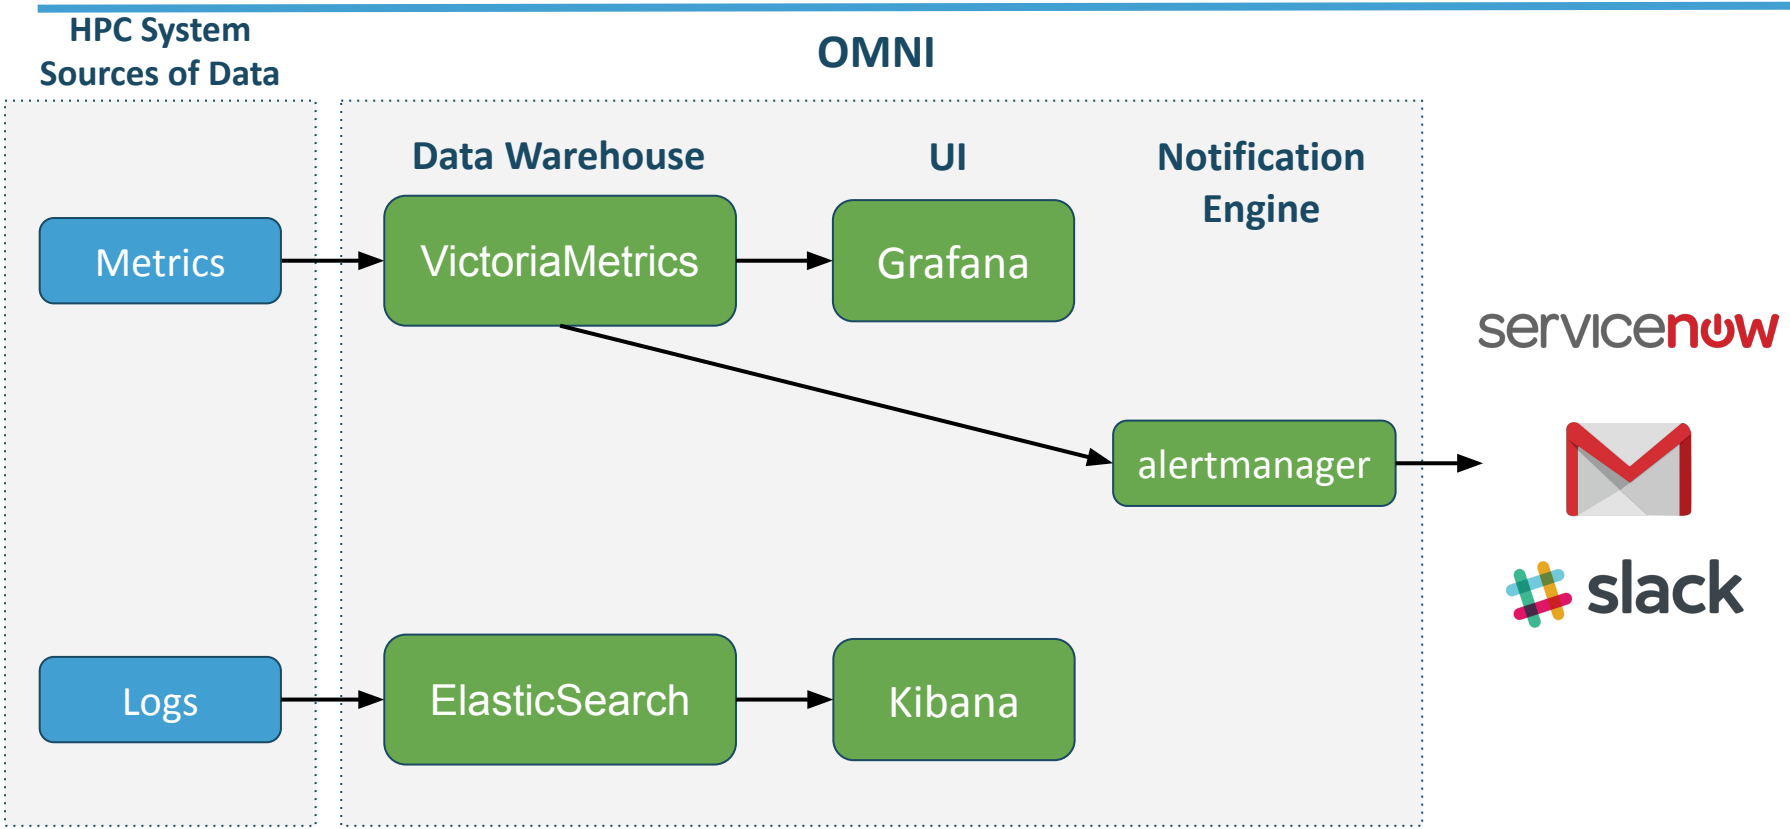

Why not stay using ElasticSearch only?

Prometheus Ecosystem

VICTORIA METRICS

Grafana

Alertmanager

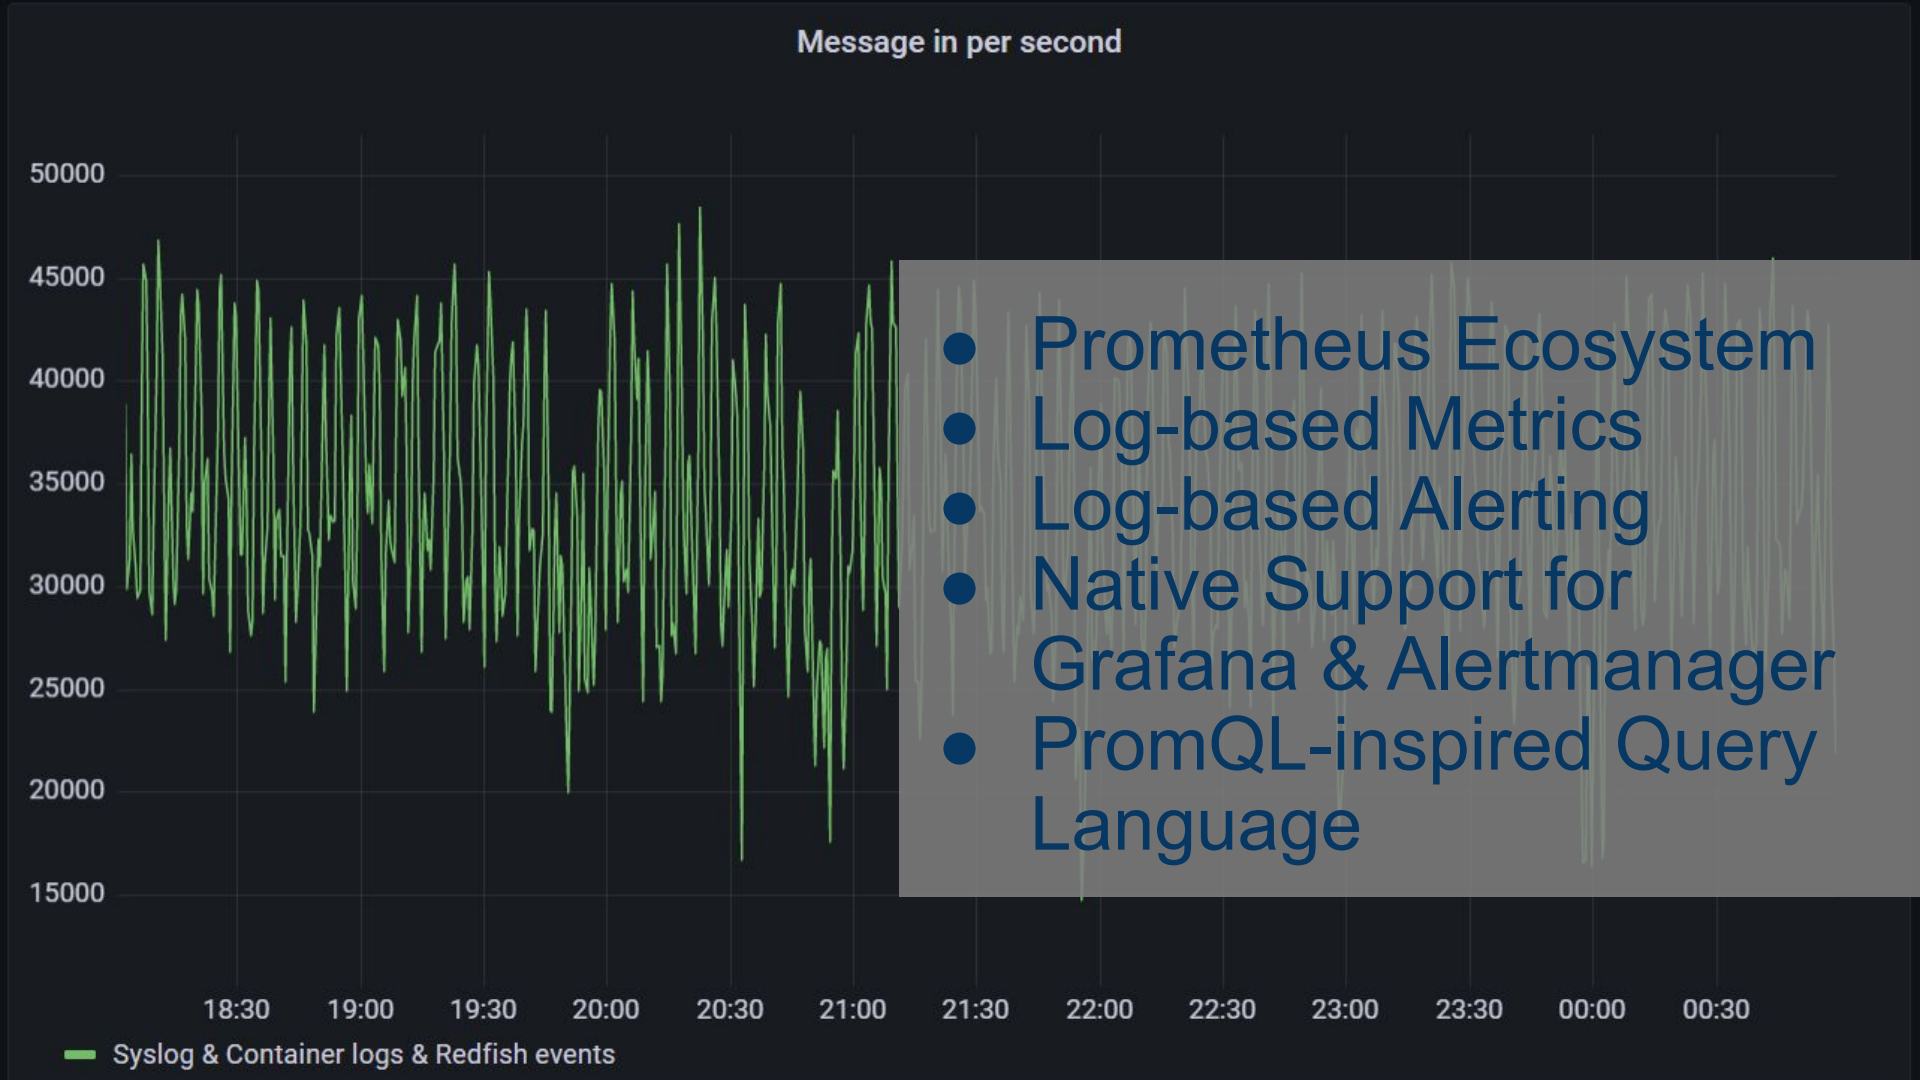

- Prometheus Ecosystem
- Log-based Metrics
- Log-based Alerting
- Native Support for Grafana & Alertmanager
- PromQL-inspired Query Language

Data pipeline

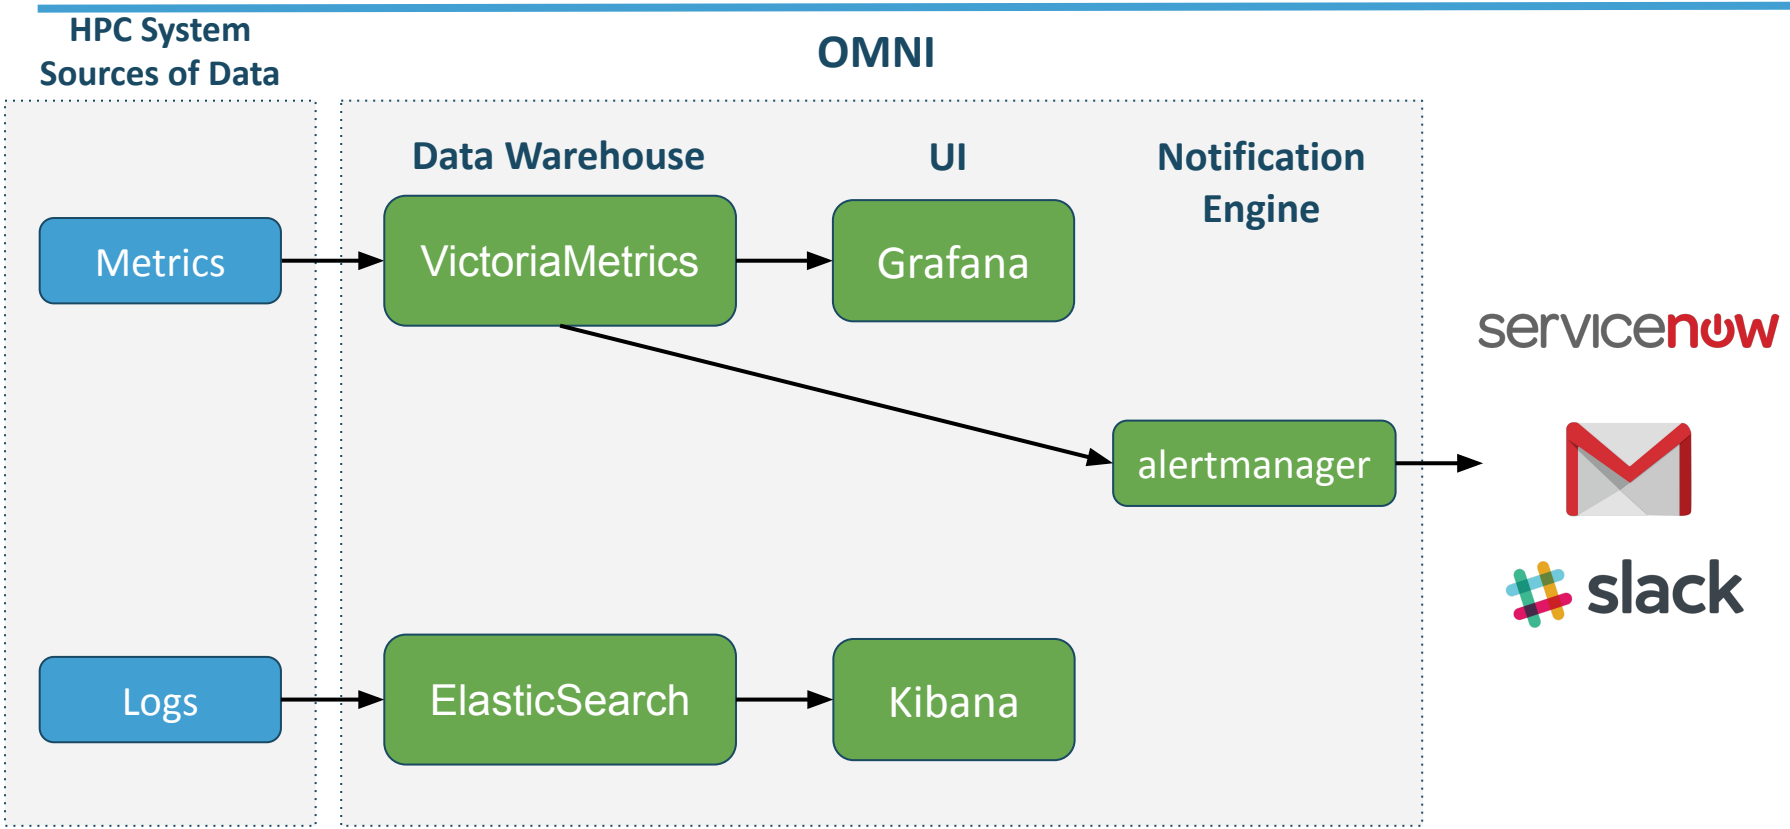

Data pipeline

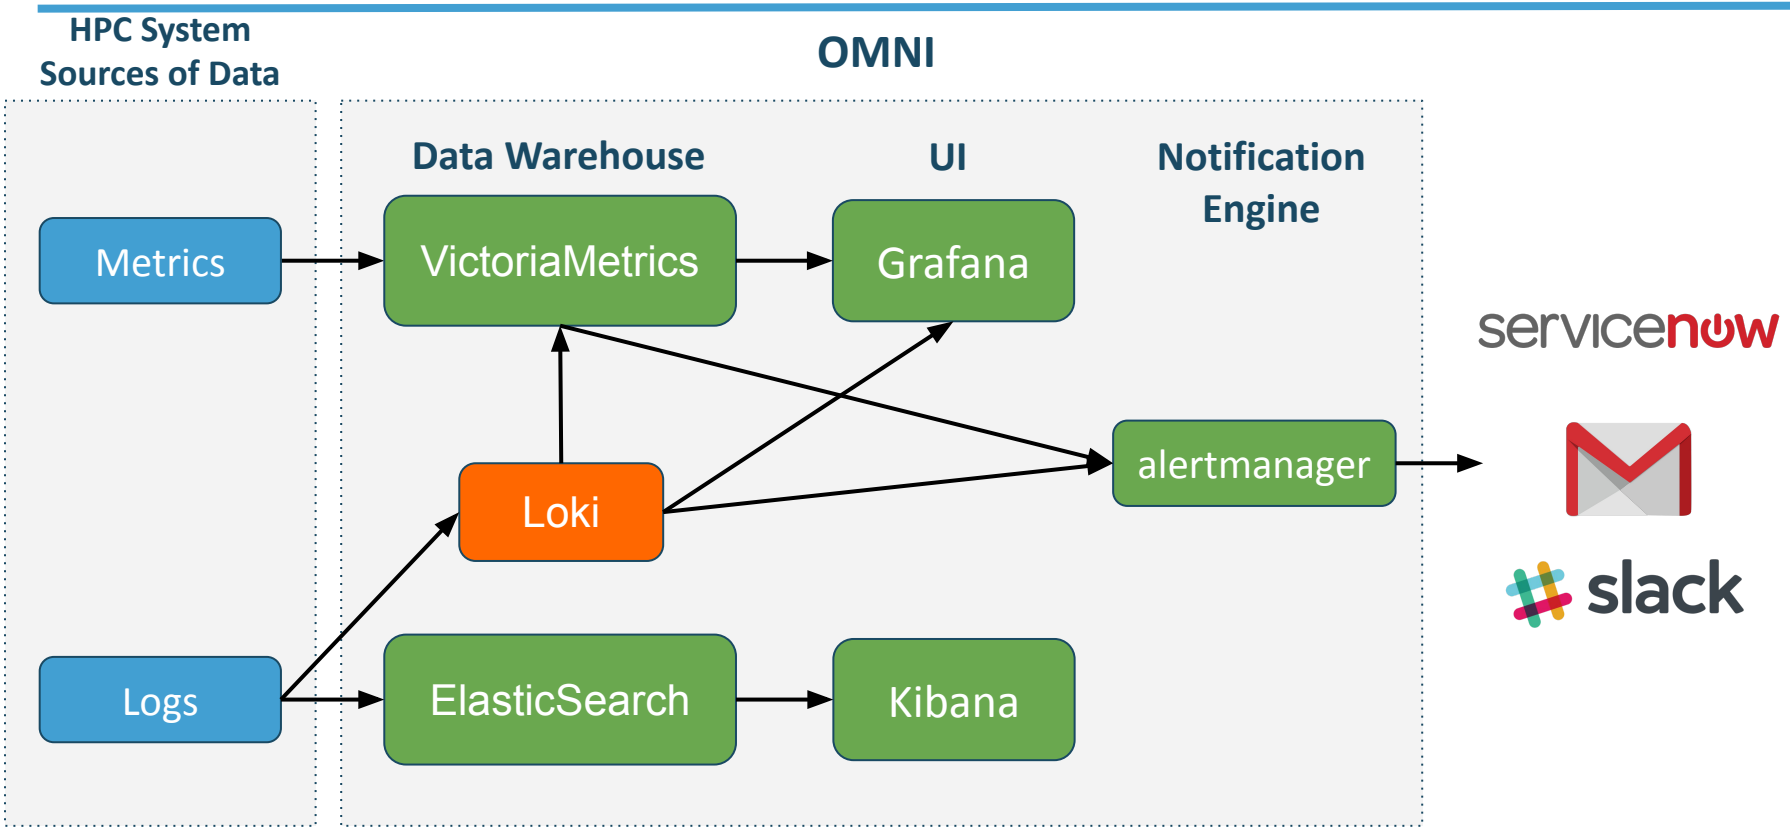

- A log aggregation system inspired by Prometheus

- Like Prometheus, but for Logs

How Loki works

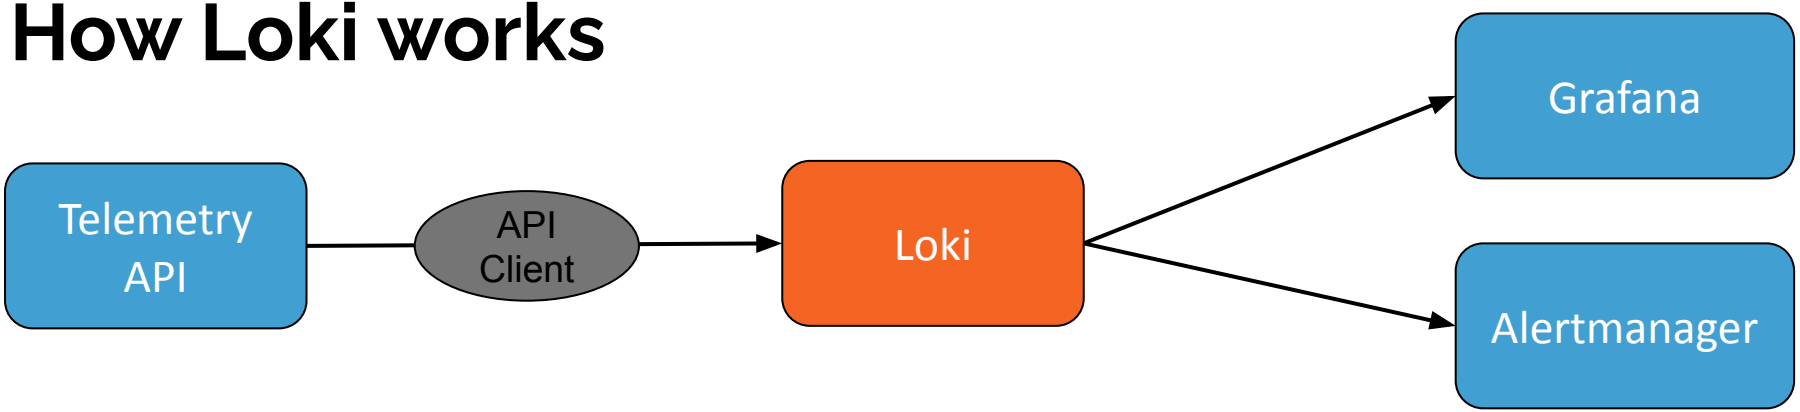


```
2021-08-02T13:58:50Z

{namesapce:"sma",pod:"rsyslog-aggregator-1"...}

rsyslogd: ... error ...
```

How Loki stores logs

```
2021-08-02T13:58:50Z  {namespace:"sma"...} rsyslogd: ...error...
```

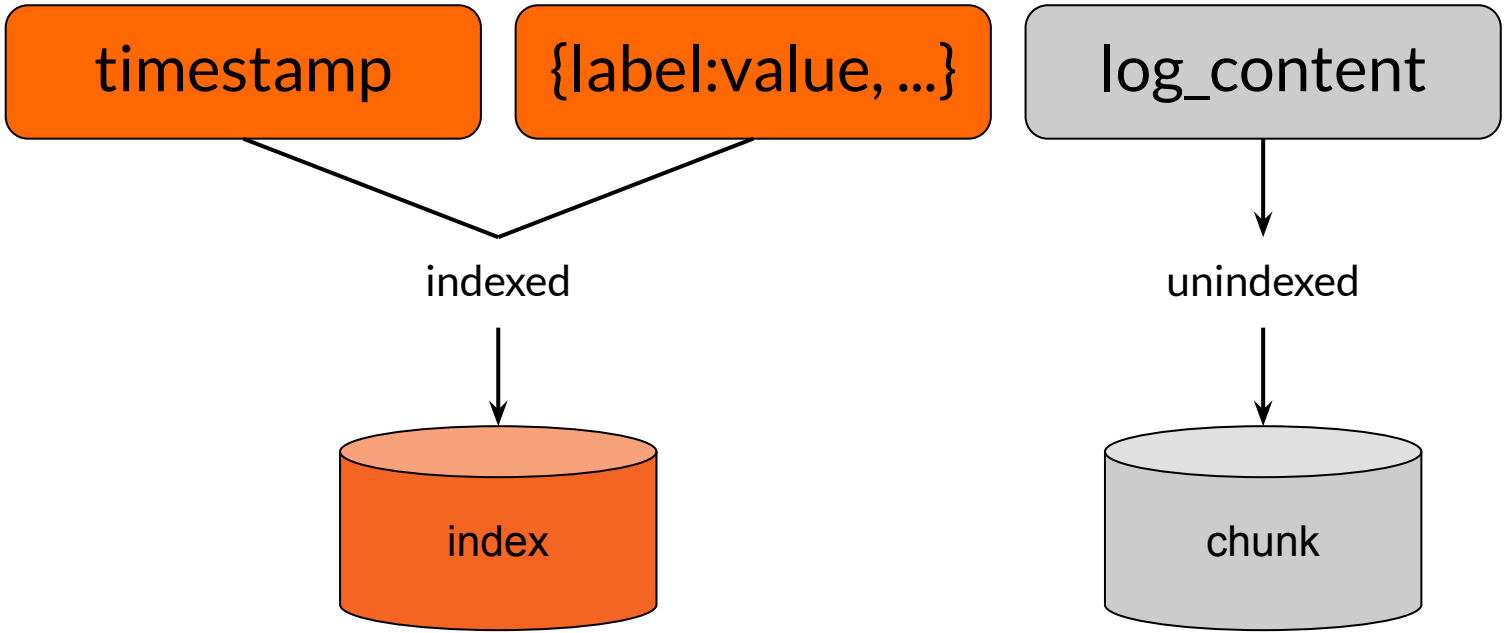

Log streams and chunks

```
2021-08-02T13:58:50Z {namespace:"sma",pod:"rsyslog-aggregator-1"...} log_msg 1
2021-08-02T13:58:51Z {namespace:"sma",pod:"rsyslog-aggregator-1"...} log msg 2
2021-08-02T13:58:52Z {namespace:"sma",pod:"rsyslog-aggregator-1"...} log_msg 3
```

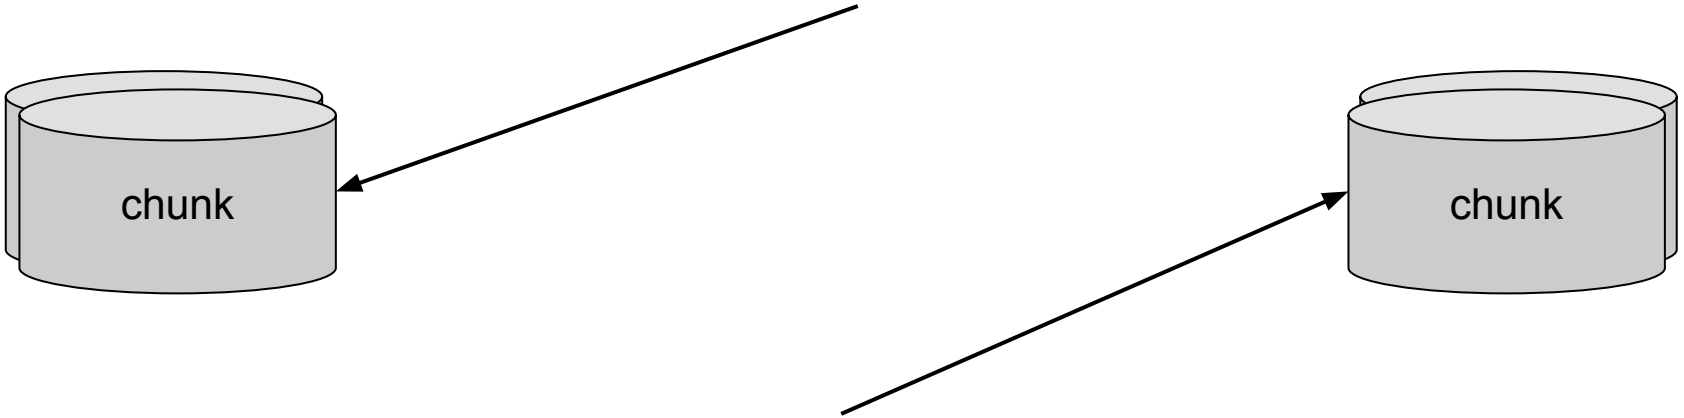


```
2021-08-02T13:57:10Z {namespace:"sma",pod:"cluster-kafka-0"...} log msg 1
2021-08-02T13:58:20Z {namespace:"sma",”pod:"cluster-kafka-0"...} log msg 2
2021-08-02T13:59:30Z {namespace:"sma",pod:"cluster-kafka-0"...} log_msg 3
```

Chunks

```
{namespace:"sma",
pod:"rsyslog-aggregator-1"...}
```

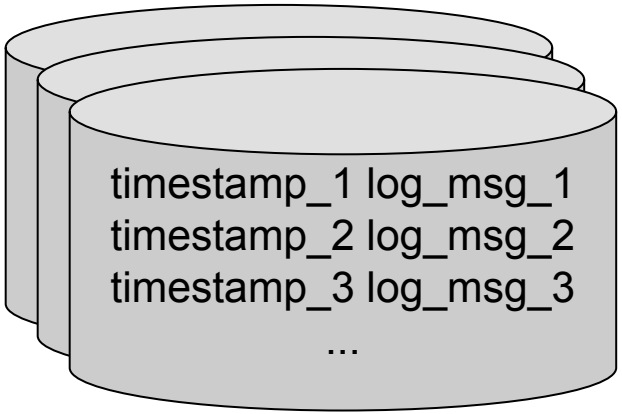


```
{namespace:"sma",
pod:"cluster-kafka-0"...}
```

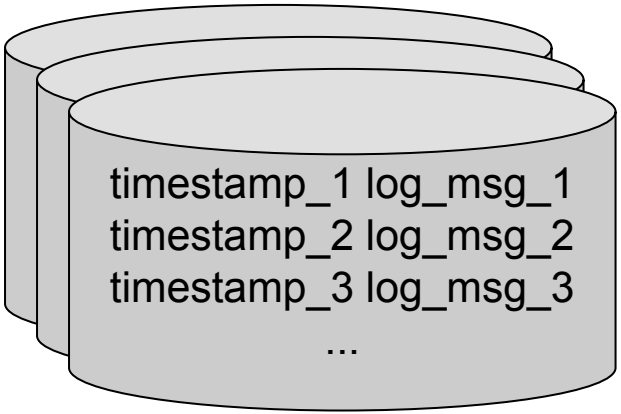

How a query works

Logcli/Grafana

Label matchers [Filter expressions]

Query:

{namespace="sma"} |= "error"

between timestamp_2

and timestamp_3

How logs look like in Loki/Grafana

log_content

{label:value, ...}

timestamp

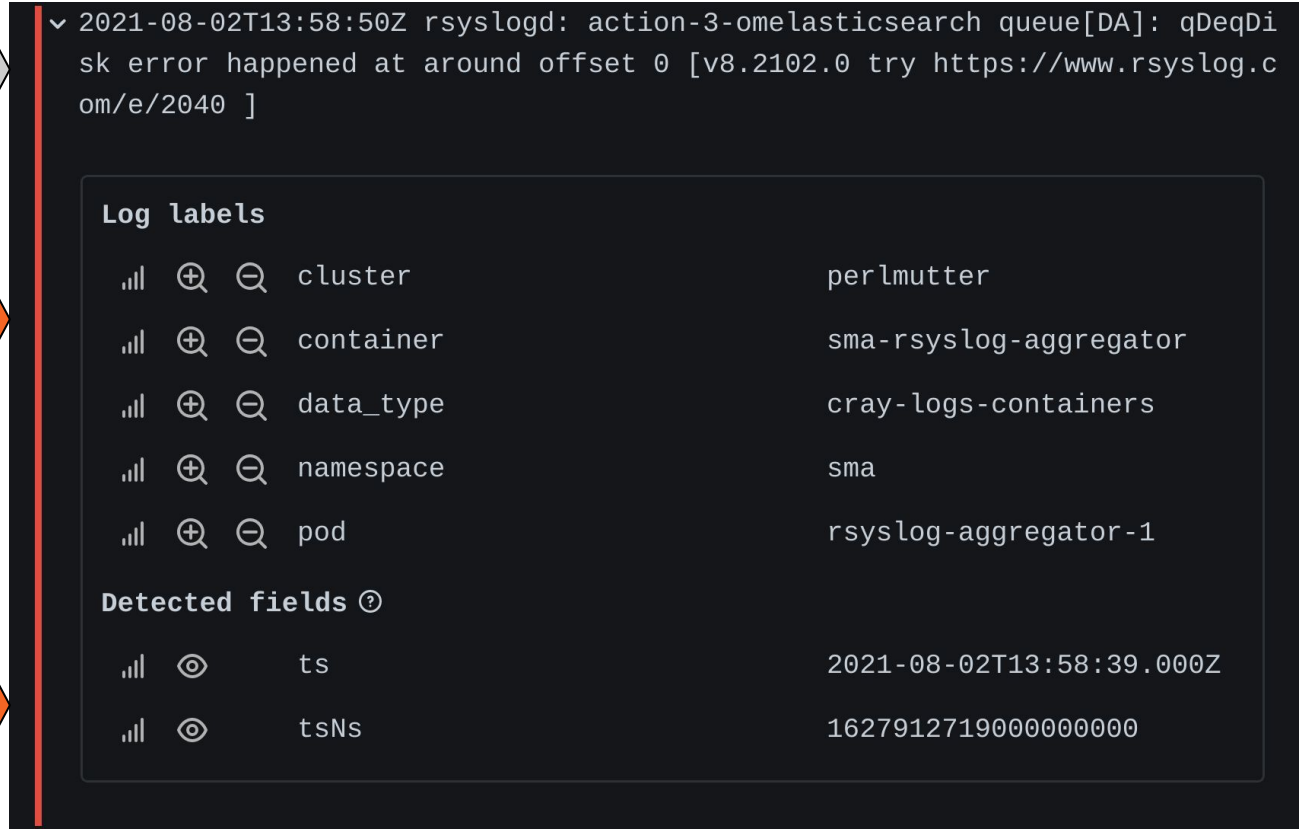

Selecting log streams with LogQL

{namespace="sma", pod=~"rsyslog-*"}	|= "error" != "timeout"
Label matchers	Filter expressions

- = exactly equal to a string
- != not equal to a string
- =~ regex matches
- !~ regex does not match

- |= contain a string
- != does not contain a string
- |~ regex matches
- !~ regex does not match

LogQL functions

func({label matchers} filter expressions [time range])

Convert logs into metrics

count_over_time({pod="rsyslog-aggregator-1"}[10m])

- rate(log-range): calculates the number of entries per second
- absent_over_time(log-range): returns an empty vector if the range vector passed to it has any elements and a 1-element vector with the value 1 if the range vector passed to it has no elements

LogQL functions

Same functions from PromQL

```
sum(count_over_time({namespace="sma"}[1m])) by (pod)
```

avg, max, min, topk, bottomk, etc ...

Alerting Rules in Loki

Query:

sum(count_over_time({hostname=~"ncn-.+"} |= "error"[1m])) by (hostname)

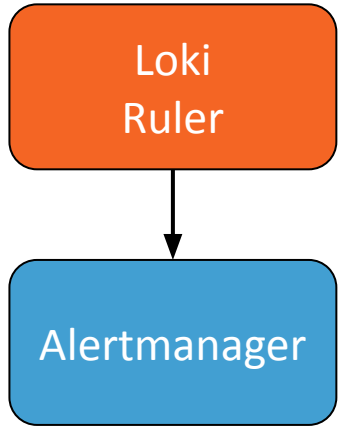

Alerting Rule:

```
- alert: HighLogErrorRate
  expr: sum(count_over_time({hostname=~"ncn-.+"} |= "error"[1m])) by (hostname) > 100
  for: 5m
```

Like Prometheus

- Has native support in Grafana
- Use the same labels
- Use PromQL-inspired language for queries
- Has native support for Alertmanager

Data pipeline

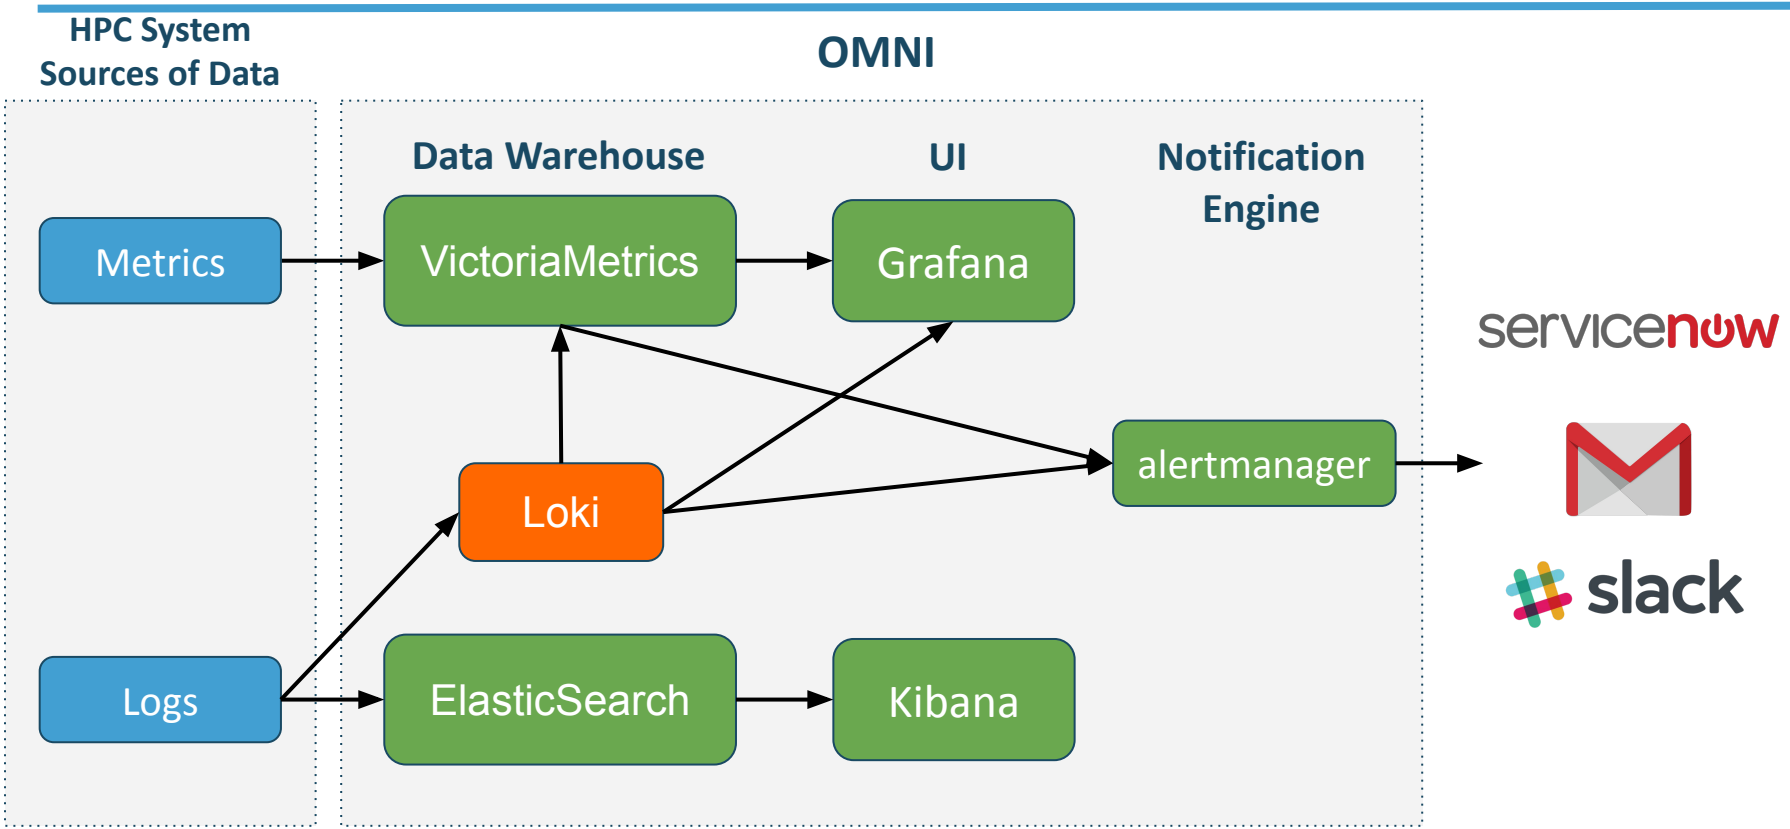

Performance

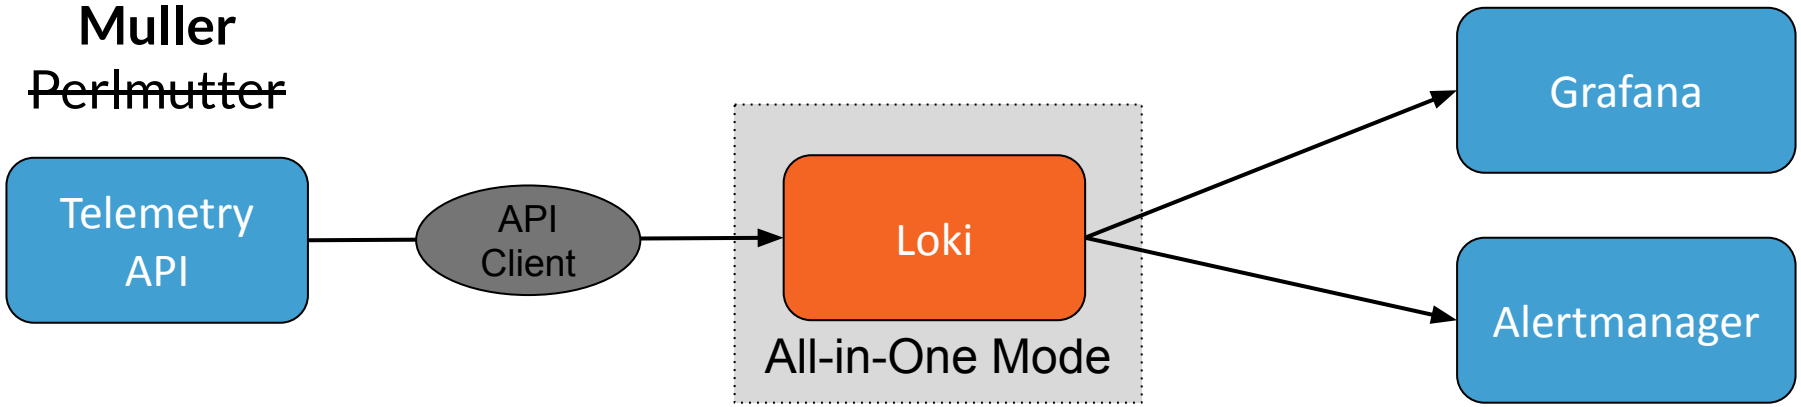

Distributed Mode

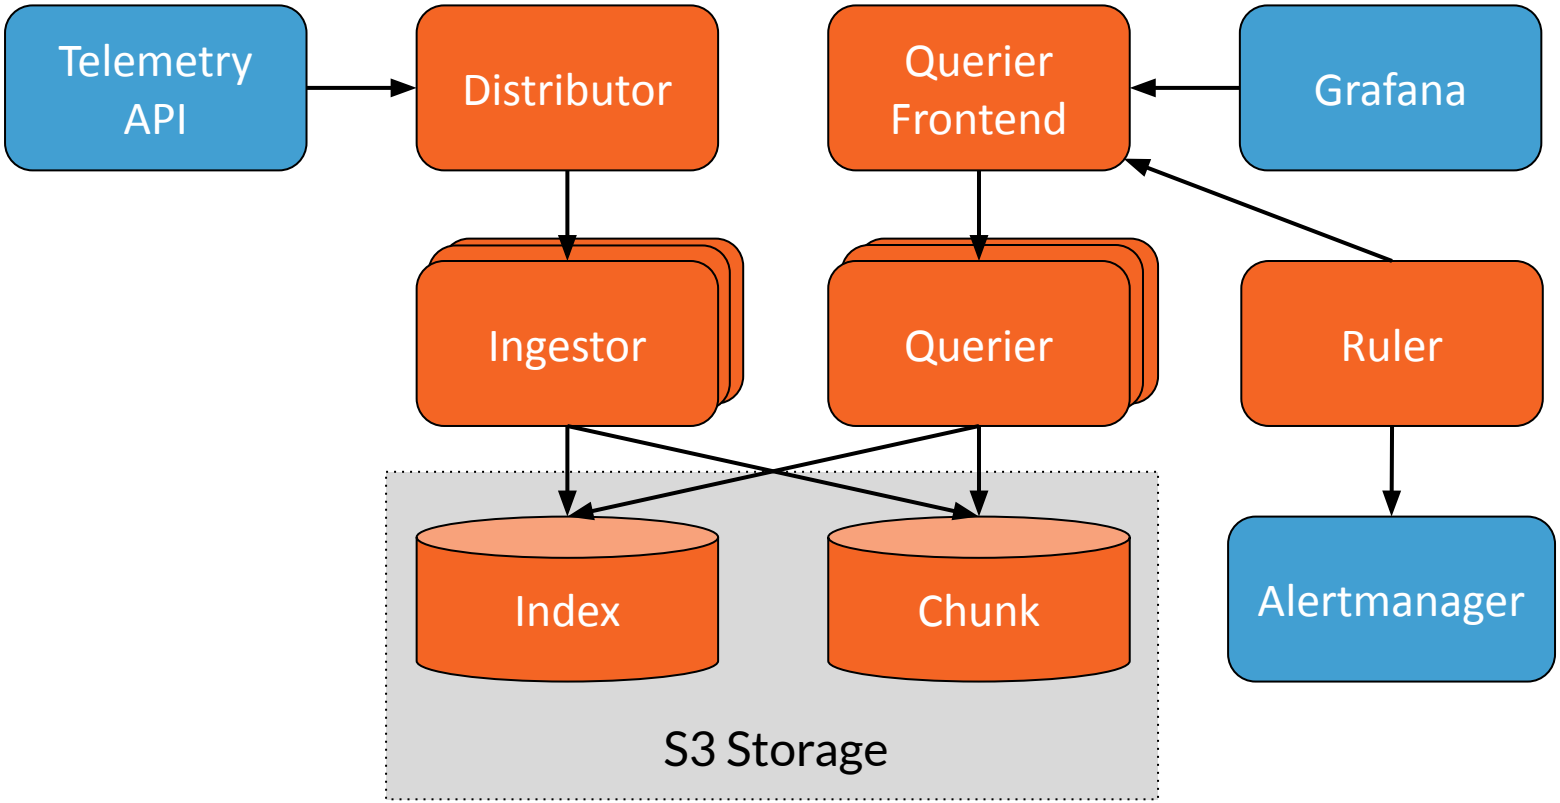

Reducing Labels

Raw Data from Telemetry API:

{ "procid":"5374",
"timereported":"2021-06-11T03:41:15",
"message":" Auth returned 1:",
"hostname":"ncn-s017",
"tag":"python3[5374]",
"priority":"30",
"severity":"info", "
facility":"daemon" }

Log for Loki:

2021-06-11T03:41:15

{hostname :"ncn-s017",
tag:"python3[5374]",
priority:"30",
Severity:"info",
"facility":"daemon" }

" Auth returned 1:"

Overusing Labels

{hostname :"ncn-s017", tag:"python3[5374]", priority:"30", Severity:"info", "facility":"daemon" }

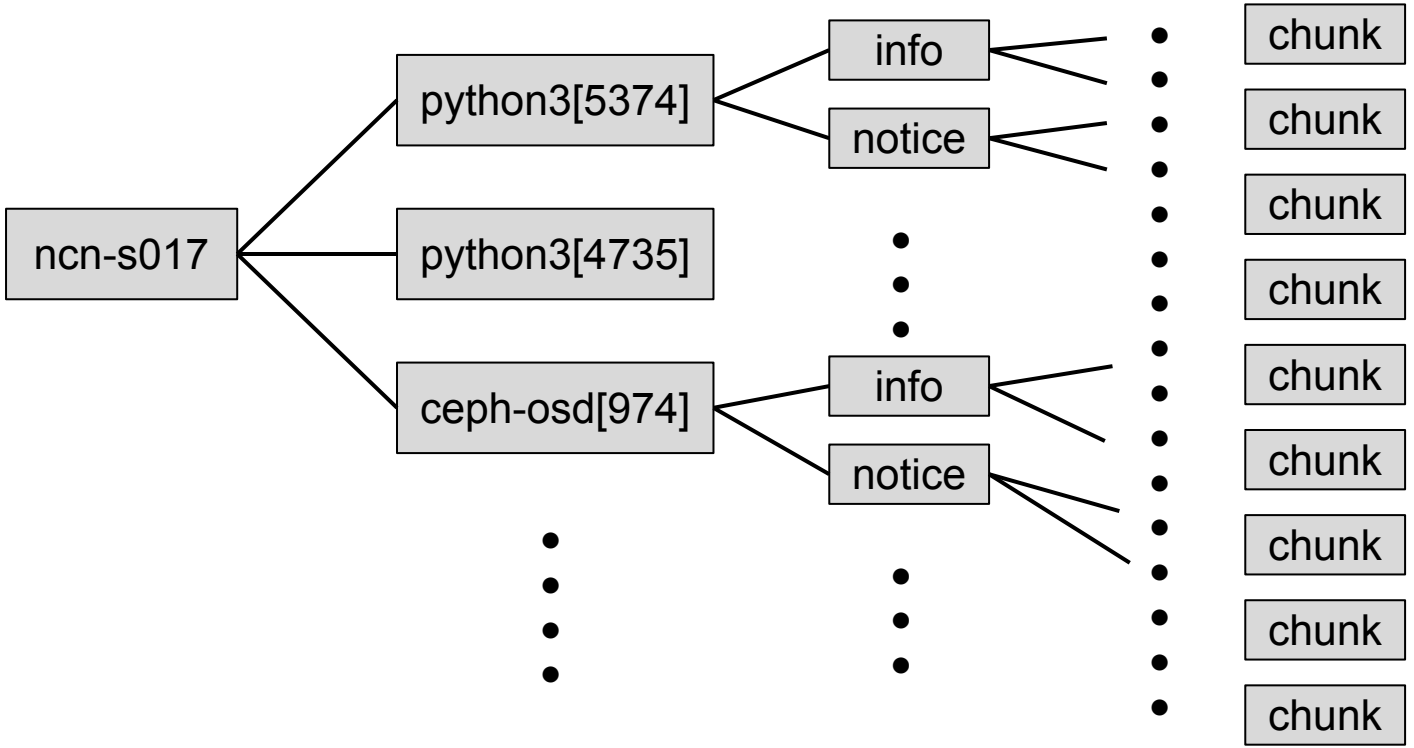

Reducing Labels

Before:

2021-08-11T03:41:15

{hostname :"ncn-s017",
tag:"python3[5374]",
priority:"30",
Severity:"info",
"facility":"daemon" }

" Auth returned 1:"

After:

2021-08-11T03:41:15

{hostname :"ncn-s017"}

"tag=python3[5374] priority=30
severity=info facility=daemon
Auth returned 1:"

Query: {hostname :"ncn-s017"} |= "tag=python3" |= "severity=info"

Use Cases

Syslog & Redfish Event Monitoring:

- Leak Detection
- HsnLinkFlap Detection
- Switch Offline Detection
- GPFS Long Waiter Detection
- GPFS Deadlock Detection

Shasta-alerter APP 1:14 PM

[FIRING:524] GPFSdeadlockDetected

Alert: - `critical`

Description: perlmutter-lmem01 reports GPFS deadlock

Details:

- **alertname:** `GPFSdeadlockDetected`
- **IP:** `10.249.0.190`
- **alertgroup:** `GPFS monitoring`
- **cluster:** `perlmutter`
- **hostname:** `lmem01`
- **message:** `[N] sdrServ: Received deadlock notification from 10.249.0.190`
- **rule_editor:** `siqideng`
- **severity:** `critical`

Demo

Connecting
Logs and Metrics
Alerts and Graphs

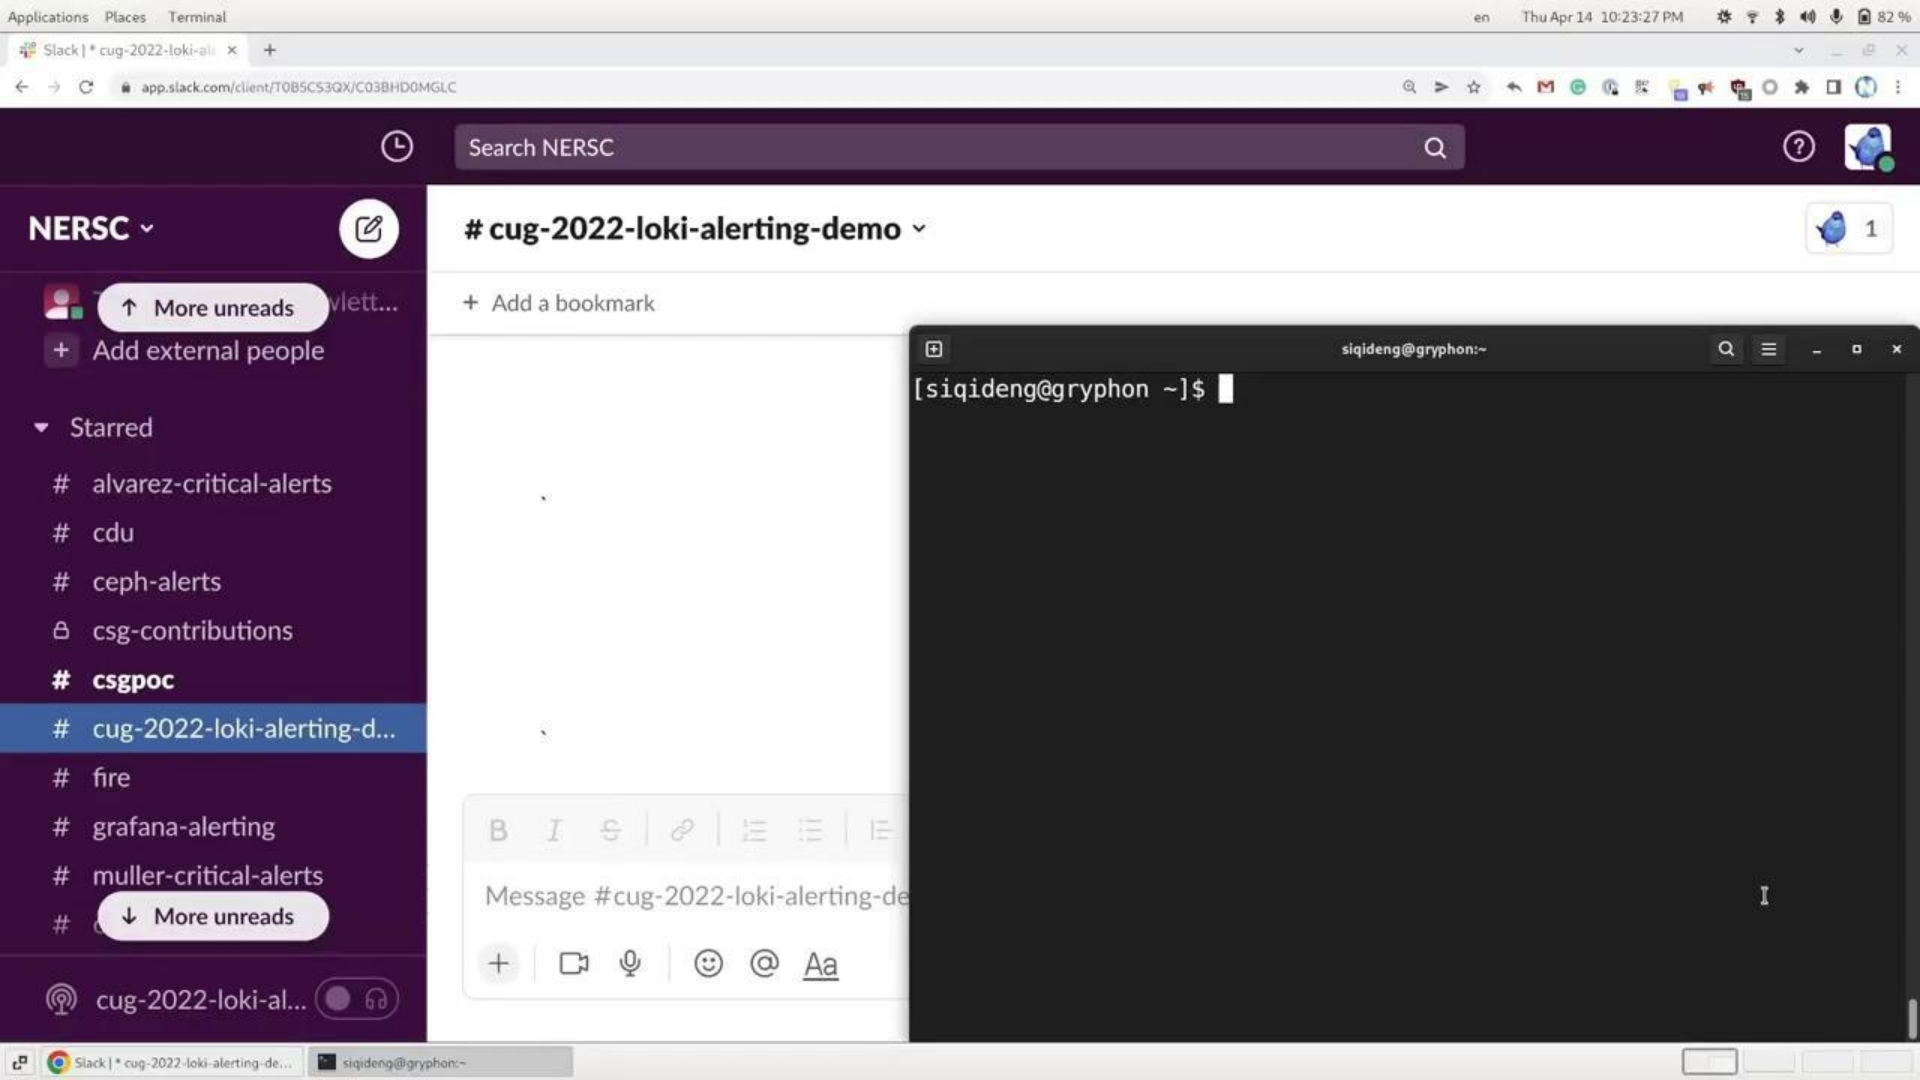
Applications Places Terminal
en
Thu Apr 14 10:23:27 PM
82 %
Slack | * cug-2022-loki-al
app.slack.com/client/T0B5CS3QX/C03BHD0MGLC
Search NERSC
NERSC
cug-2022-loki-alerting-demo
1
More unreads
Add external people
Add a bookmark
Starred
alvarez-critical-alerts
cdu
ceph-alerts
csg-contributions
csgpoc
cug-2022-loki-alerting-d...
fire
grafana-alerting
muller-critical-alerts
More unreads
cug-2022-loki-al...
Message #cug-2022-loki-alerting-de
siqideng@gryphon:~
[siqideng@gryphon ~]$
Slack | * cug-2022-loki-alerting-de...
siqideng@gryphon:~

Data pipeline

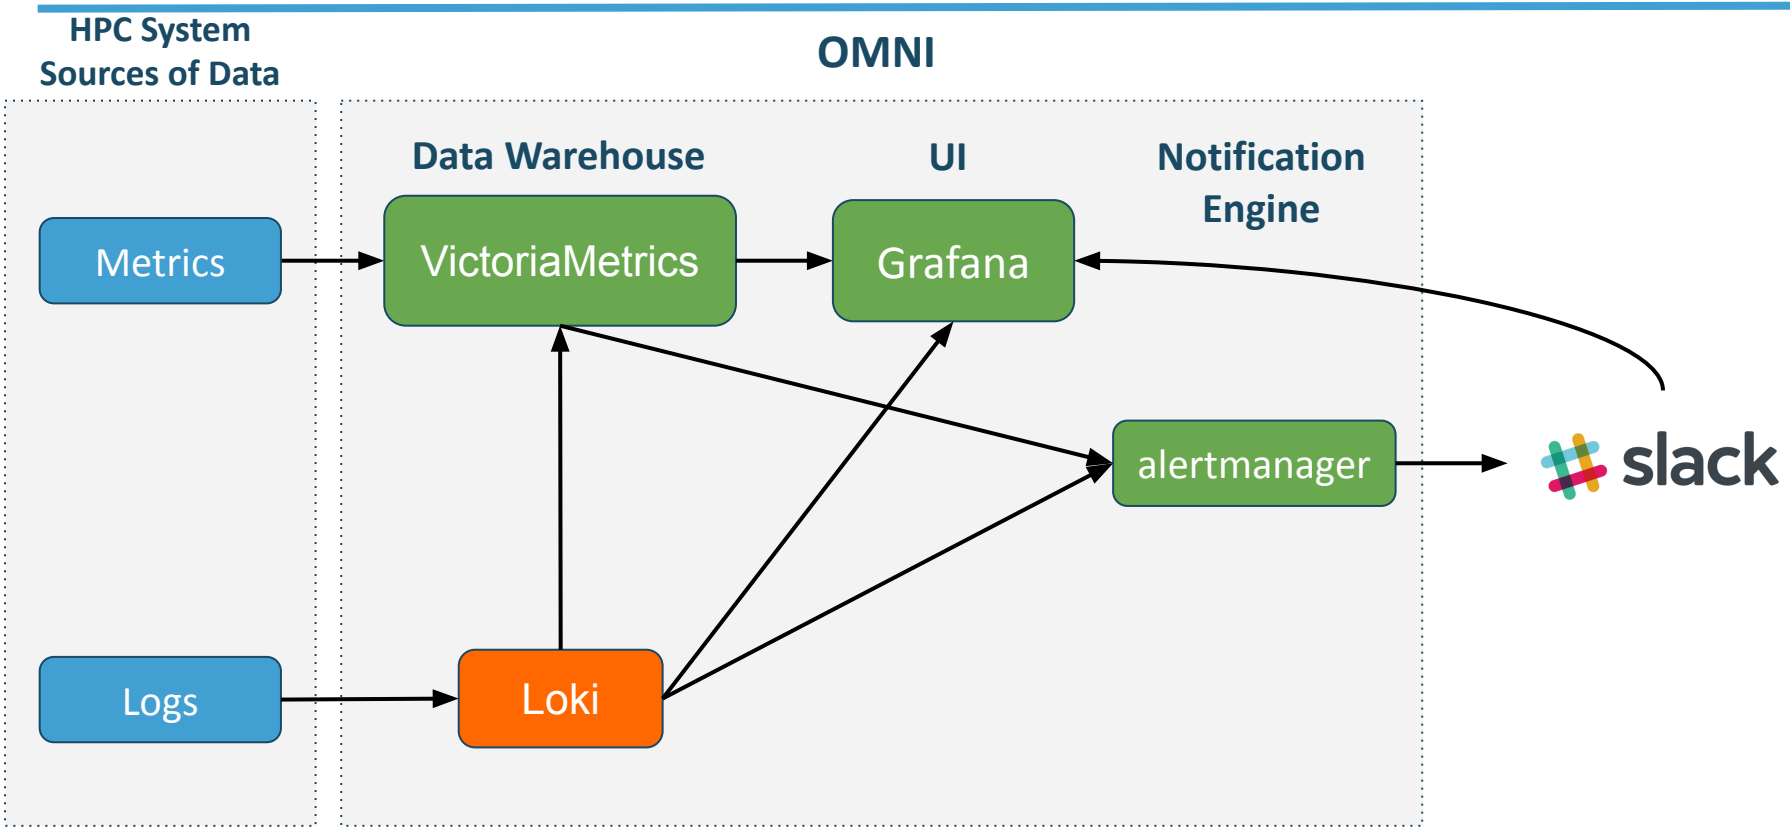

Q&A

Siqi Deng
SRE@NERSC