

“Network Monitoring and Management 2.0”

SANOG 36

Hervey Allen of the Network Startup Resource Center

www.ws.nsrc.org



These materials are licensed under the Creative Commons Attribution-NonCommercial 4.0 International license (<http://creativecommons.org/licenses/by-nc/4.0/>)

A few “Walls of Text”

I promise pictures after these initial slides...



NMM 2.0

Why?

The Why of NMM 2.0

- Finer-grained metrics (“real time network telemetry”)
 - Network telemetry streams vs. occasional data pulls
- Scaling (hyper scale)
 - Ability to measure monitor hyper-scale projects
 - Polling 10,000 devices/containers... that’s hard
 - Can have operational impact
- Portability:
 - Gather data once, use with multiple tools



NMM 2.0

How?



NMM 2.0

Traditional vs. Present Day Practices

Push vs. Pull or...

Network telemetry / push / passive vs. polling / pull

After this we would start talking about...

Monitoring vs. *Observing* (o11y)

A wonderful discussion at

<https://twitter.com/isotopp/status/1328653624470331392>

NMM 2.0

Traditional vs. Present Day Practices*

Push vs. Pull or...

Network telemetry / push / passive vs. polling / pull

- Traditional: standards-based like snmp or agents (Nagios, Check MK)
- Present: some push protocols:
 - Cisco compact Google Protocol Buffers
 - Google Protocol Buffers
 - Json
- Newer agents used with present day network monitoring stacks
 - Telegraf, beats, node exporter, Promtail, logstash, etc...

**Sort of... Depends on your needs, resources, goals, etc.*

NMM 2.0

Traditional vs. Present Day Practices

How we store our network metrics (**NoSQL** vs. **Relational**)

- Traditional: relational data stores for network metrics
 - MySQL, PostgreSQL, SQLite, Oracle, DB2, SQL Server, MariaDB, etc.
- Present: a few time series data stores or NoSQL databases:
 - Cassandra
 - CouchDB
 - Elasticsearch
 - InfluxDB
 - MongoDB
 - Prometheus
 - RRDTOol (Old school time series data store! Heavily used.)
 - TimescaleDB


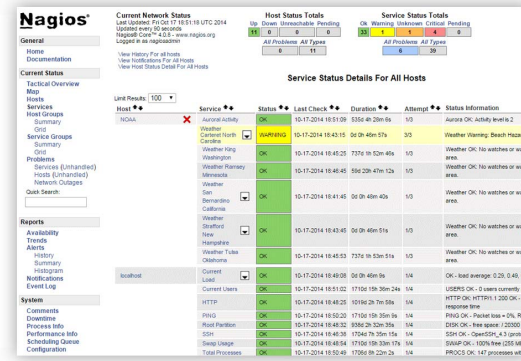
NMM 2.0

Traditional vs. Present Day Practices*

Dashboards vs. Monolithic interfaces to network metrics

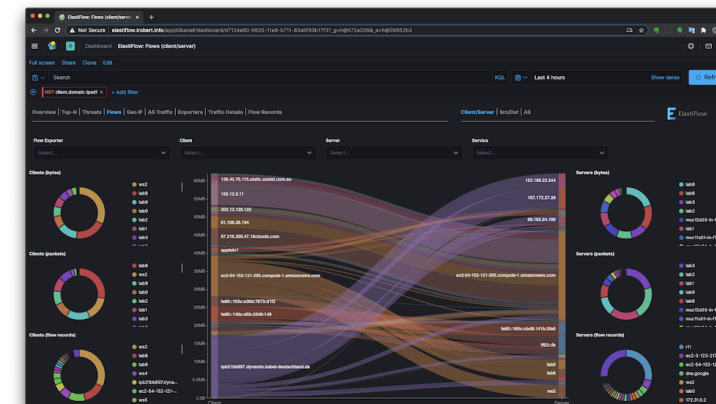
– Traditional: Constrained interfaces with less extensibility

- Nagios
- Cacti
- LibreNMS
- SmokePing



– Present: Dashboards massively configurable, harder to get started (for some)

- Chronograf, Grafana, Kibana*
 - *Elastiflow: a flow collection tool that use Kibana and Elasticsearch



NMM 2.0

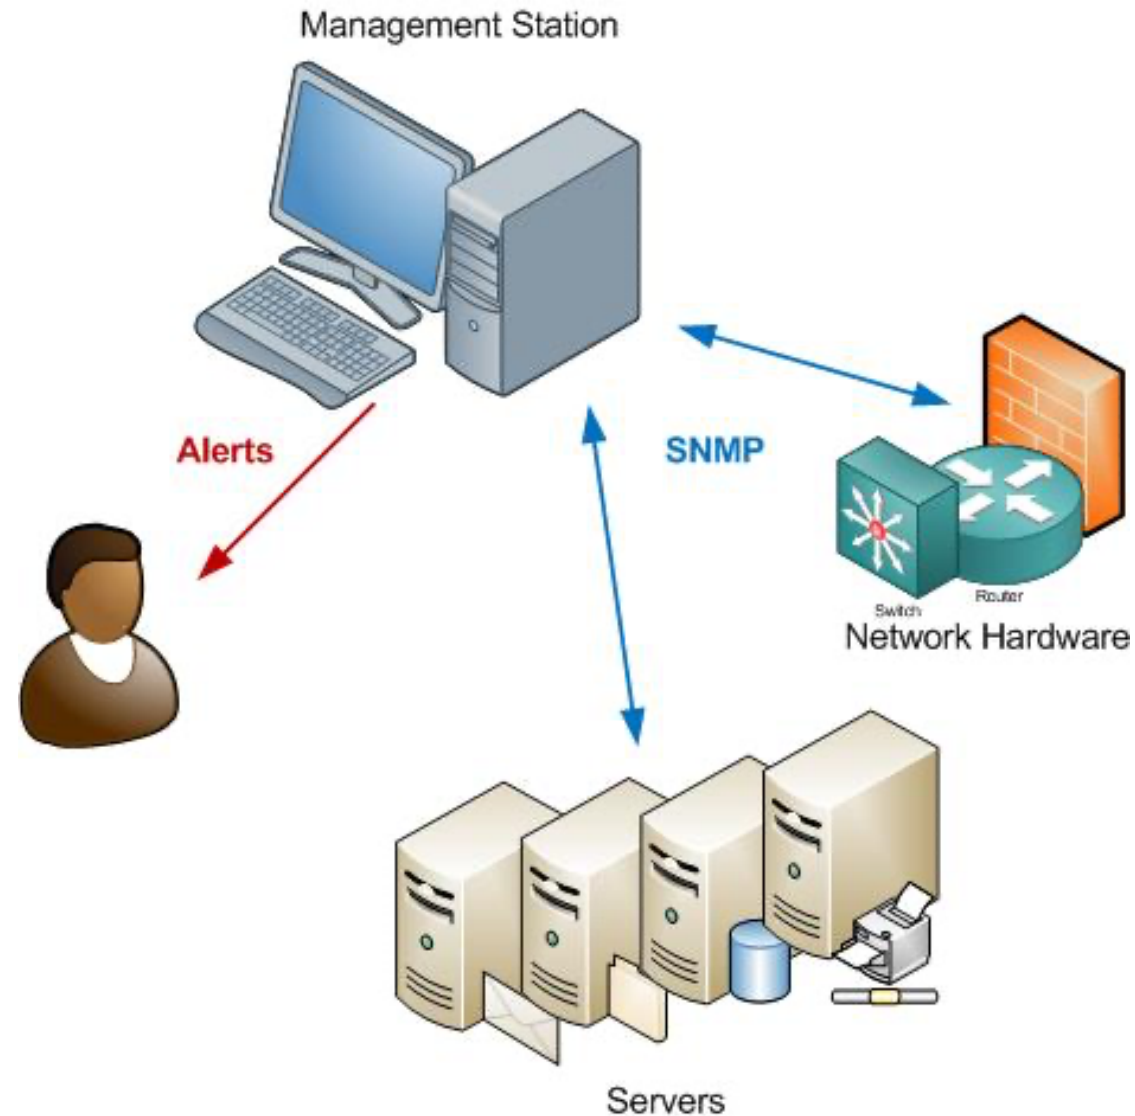
Traditional vs. Present Day Practices

Alerting

- Traditional: If available, built-in to the tool. Often minimal.
 - *SmokePing*: alerts.cfg with custom regex language
 - *Nagios*: template based. Very well implemented.
 - *Cacti*: plugins required. Variable.
 - *LibreNMS*: built-in. Not intuitive. Improving over time.
- Present: Often a separate tool or built-in to dashboard tool
 - *AlertManager* (Prometheus solution)
 - *Grafana* (visualizer/analyzer)
 - *Kapacitor* (TICK Stack)
 - *Kibana* (ELK Stack)

Stacks: ELK, TICK, Prometheus. We'll get to these! 😊


Classical Polling Model



“Network Telemetry” or “Push Model”



The Elastic Stack (ELK)

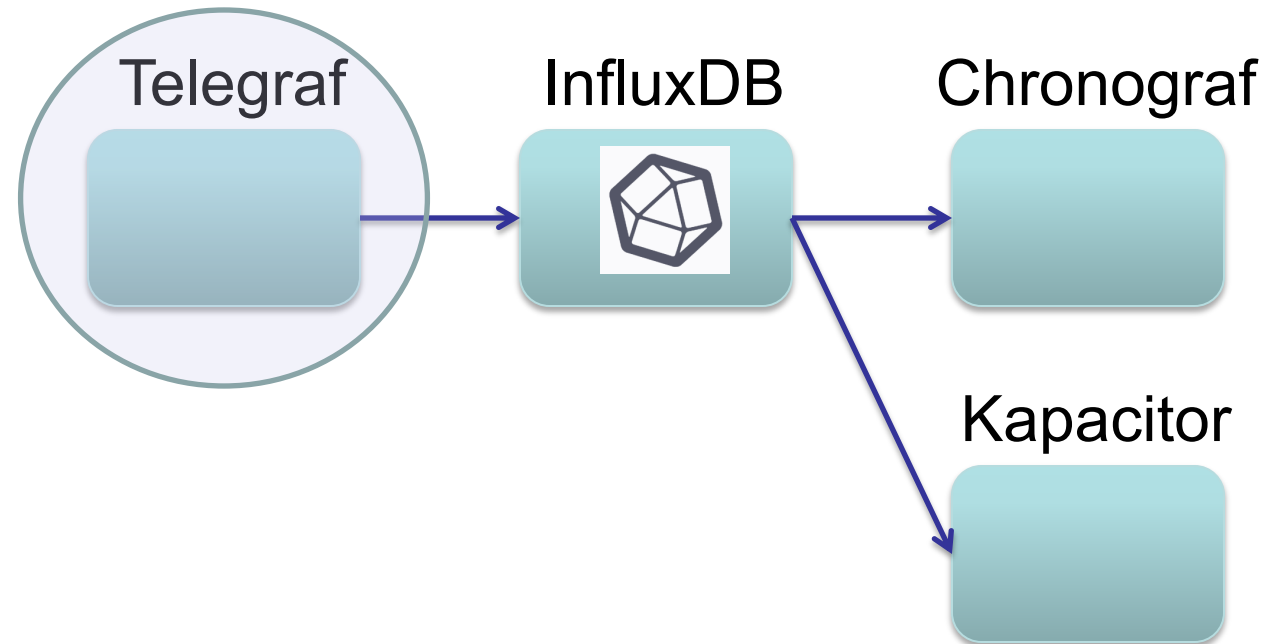


Present day network measurement “Stacks” are a group of software components that work together to form a monitoring and management solution.

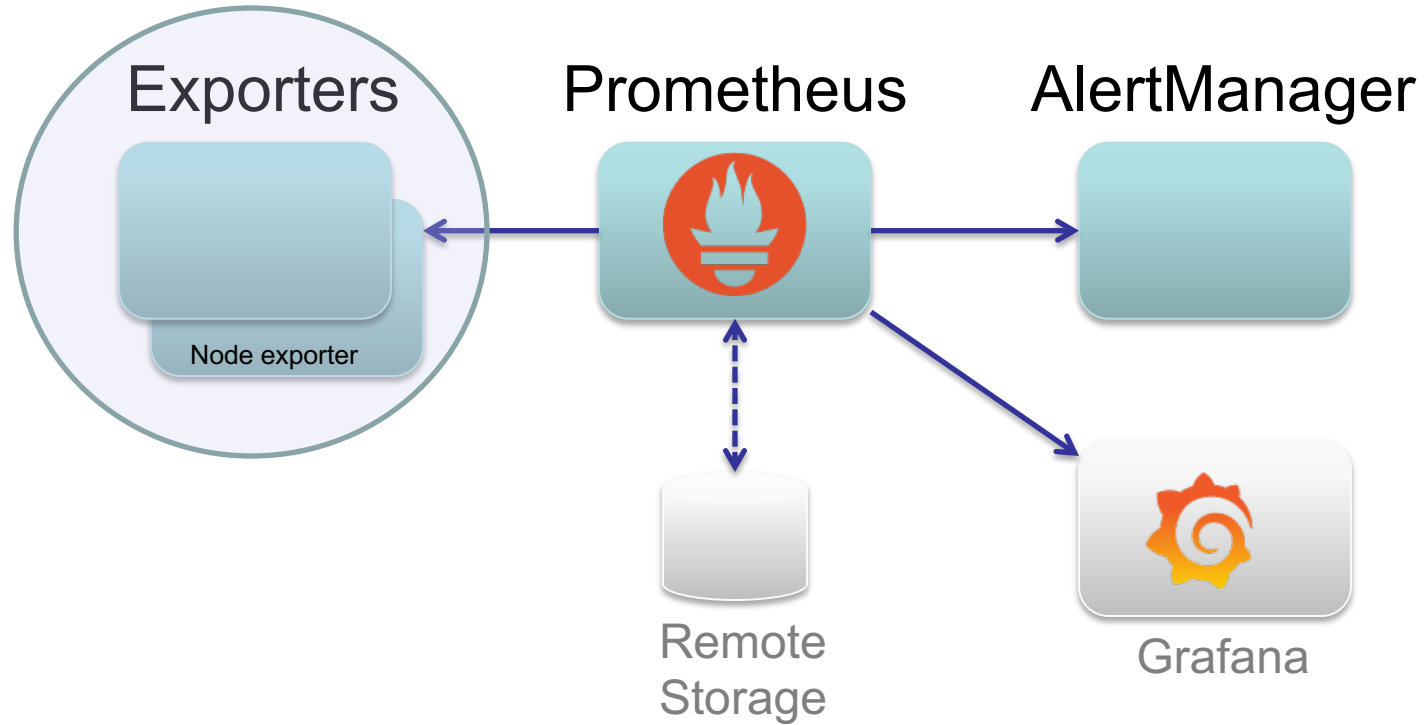
Typical stacks include (more or less):

- Mechanism(s) to push data to a data store (agents, protocols, both)
- A time series or NoSQL data store
- An engine to query the data store and present results in a graphical format in a dashboard format.
- A built-in or separate alerting component that works with the data store
- Note that many components are interchangeable between stacks

The TICK Stack



Prometheus



Typical Relational Store (MySQL)

```
CREATE TABLE `device_metrics` (  
  `id` int(11) NOT NULL AUTO_INCREMENT,  
  `timestamp` int(11) NOT NULL,  
  `metric1` smallint(6) NOT NULL,  
  `metric2` int NOT NULL,  
  `metric3` float NOT NULL DEFAULT '0',  
  PRIMARY KEY (`id`),  
  UNIQUE KEY `idposition_UNIQUE` (`id`)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8
```



What this looks like

Field	Type	Null	Key	Default	Extra
id	int(11)	NO	PRI	NULL	auto_increment
timestamp	int(11)	NO		NULL	
metric1	smallint(6)	NO		NULL	
metric2	int(11)	NO		NULL	
metric3	float	NO		0	

This is moderately efficient vs. putting every metric in to a different table. But, you still only get one data set per row.



What this looks like with inserted data

```
SELECT * FROM device_metrics;
```

id	timestamp	metric1	metric2	metric3
1	1610232093	29001	1800789199	79.86
2	1610232094	29002	1800789200	79.98
3	1610232095	29003	1800789201	77.67
4	1610232065	29004	1800789223	78.32
5	1610232097	29077	1800789456	80.01
6	1610232098	29232	1800723455	79.11



Table Growth

id	timestamp	metric1	metric2	metric3
1	1610232093	29001	1800789199	79.86
2	1610232094	29002	1800789200	79.98
3	1610232095	29003	1800789201	77.67
4	1610232065	29004	1800789223	78.32
5	1610232097	29077	1800789456	80.01
6	1610232098	29232	1800723455	79.11

← A new data point every second!

- With “push” model and agents much more telemetry data.
- Querying and displaying large numbers of metrics become inefficient in a relational model.

How to get to this? →
(Grafana)



Inefficiencies of relations...

Inserting, Updating and Selecting, or...

- Adding data
- Changing data
- Getting data

Each row increases

- Index size
- Compute

NoSQL / Time Series data stores allow for very large sets of metrics in sequence and ability to query these metrics at large scale

Time series data stores / NoSQL

A few ways to store time series data (there are many):

- timestamp, metric, timestamp, metric

or

- timestamp, metric, metric, ..., timestamp, metric, metric, ...

or

- metric, metric, metric, metric, metric, ... timestamp

Per row. Each row can have *many* columns.

- For example, Cassandra DB can support up to 2 billion columns per row!


- Nice discussion on what is time series data:

<https://www.influxdata.com/what-is-time-series-data/>

NMM 2.0

The Datastores


The Elastic Stack (ELK)



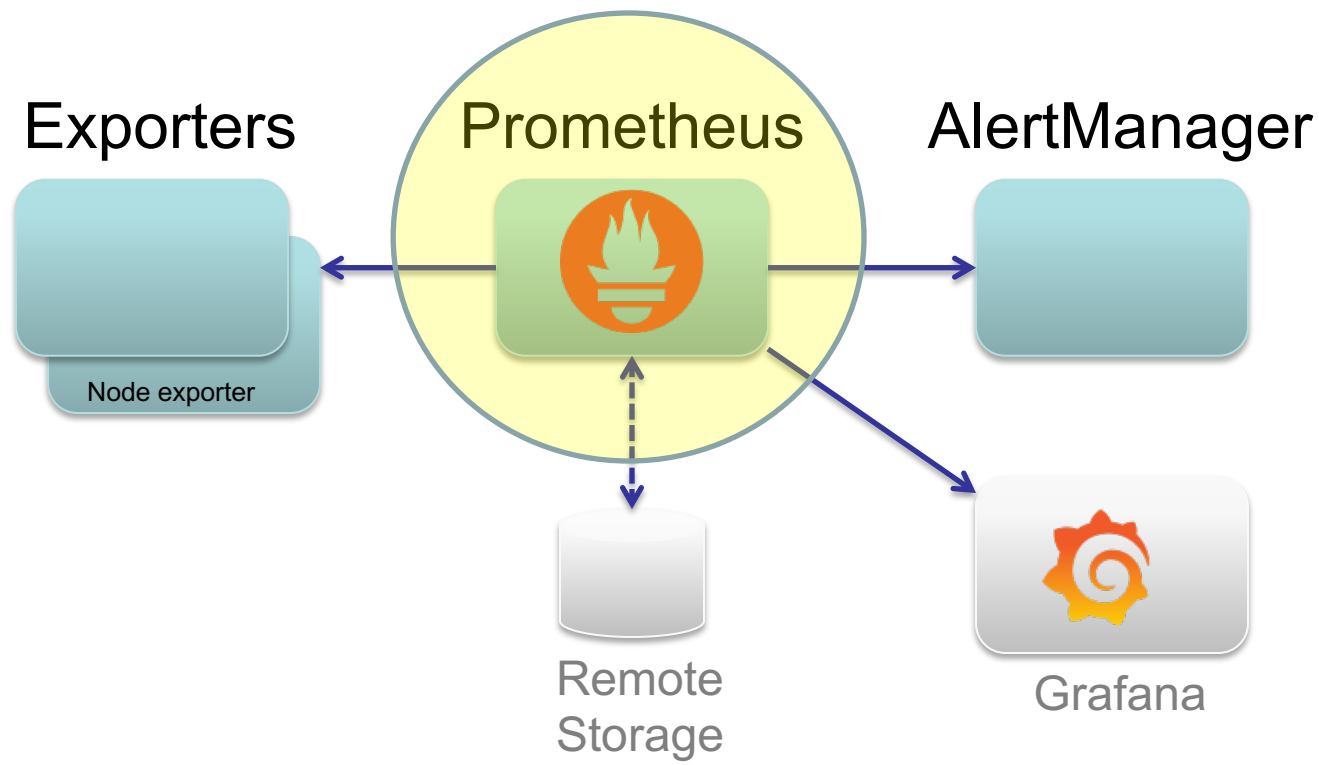
("The BLEK Stack" doesn't sound as good)



The TICK Stack



Prometheus




NMM 2.0

The Dashboards



The Elastic Stack (ELK)



Not sure whether to use Logstash or Beats?


Beats are lightweight data shippers that you install as agents on your servers to send specific types of operational data to Elasticsearch. Beats have a small footprint and use fewer system resources than Logstash.

Logstash has a larger footprint, but provides a broad array of input, filter, and output plugins for collecting, enriching, and transforming data from a variety of sources.


<https://www.elastic.co/guide/en/beats/filebeat/current/diff-logstash-beats.html>



The TICK Stack



Prometheus




*Grafana was designed to work as a UI for analyzing metrics. As such, it can work with multiple time-series data stores, including built-in integrations with Graphite, Prometheus, InfluxDB, MySQL, PostgreSQL, and Elasticsearch, and additional data sources using plugins. For each data source, Grafana has a specific query editor that is customized for the features and capabilities that are included in that data source (<https://logz.io/blog/grafana-vs-kibana/>).



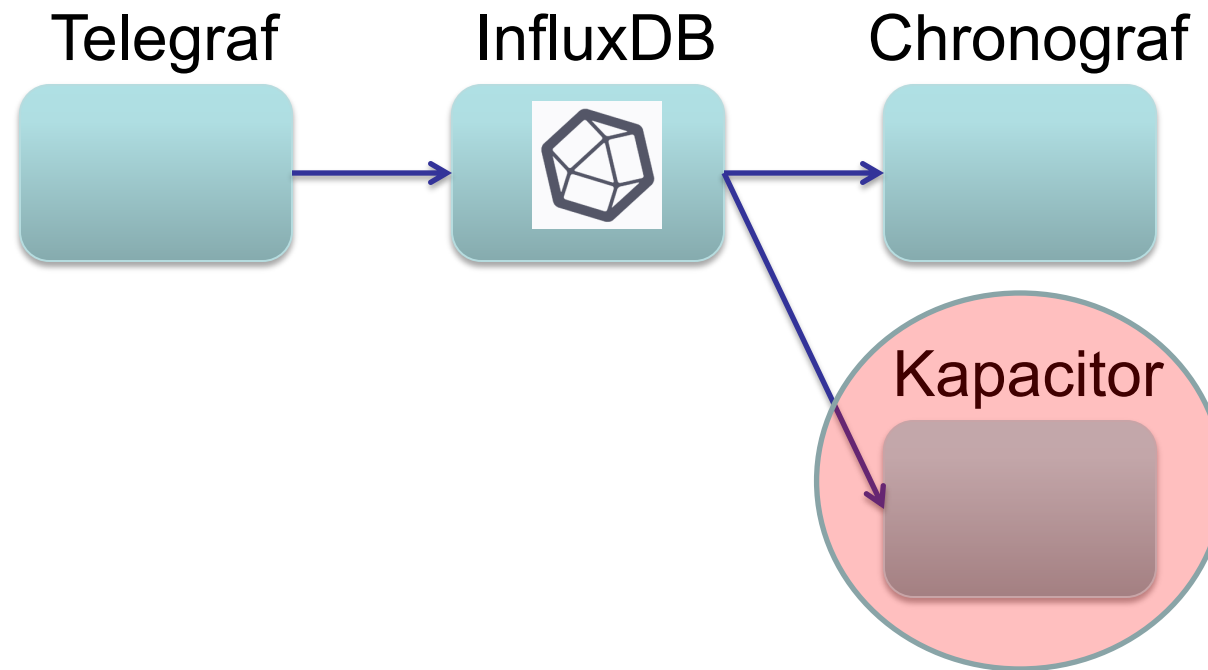
NMM 2.0

Alerting

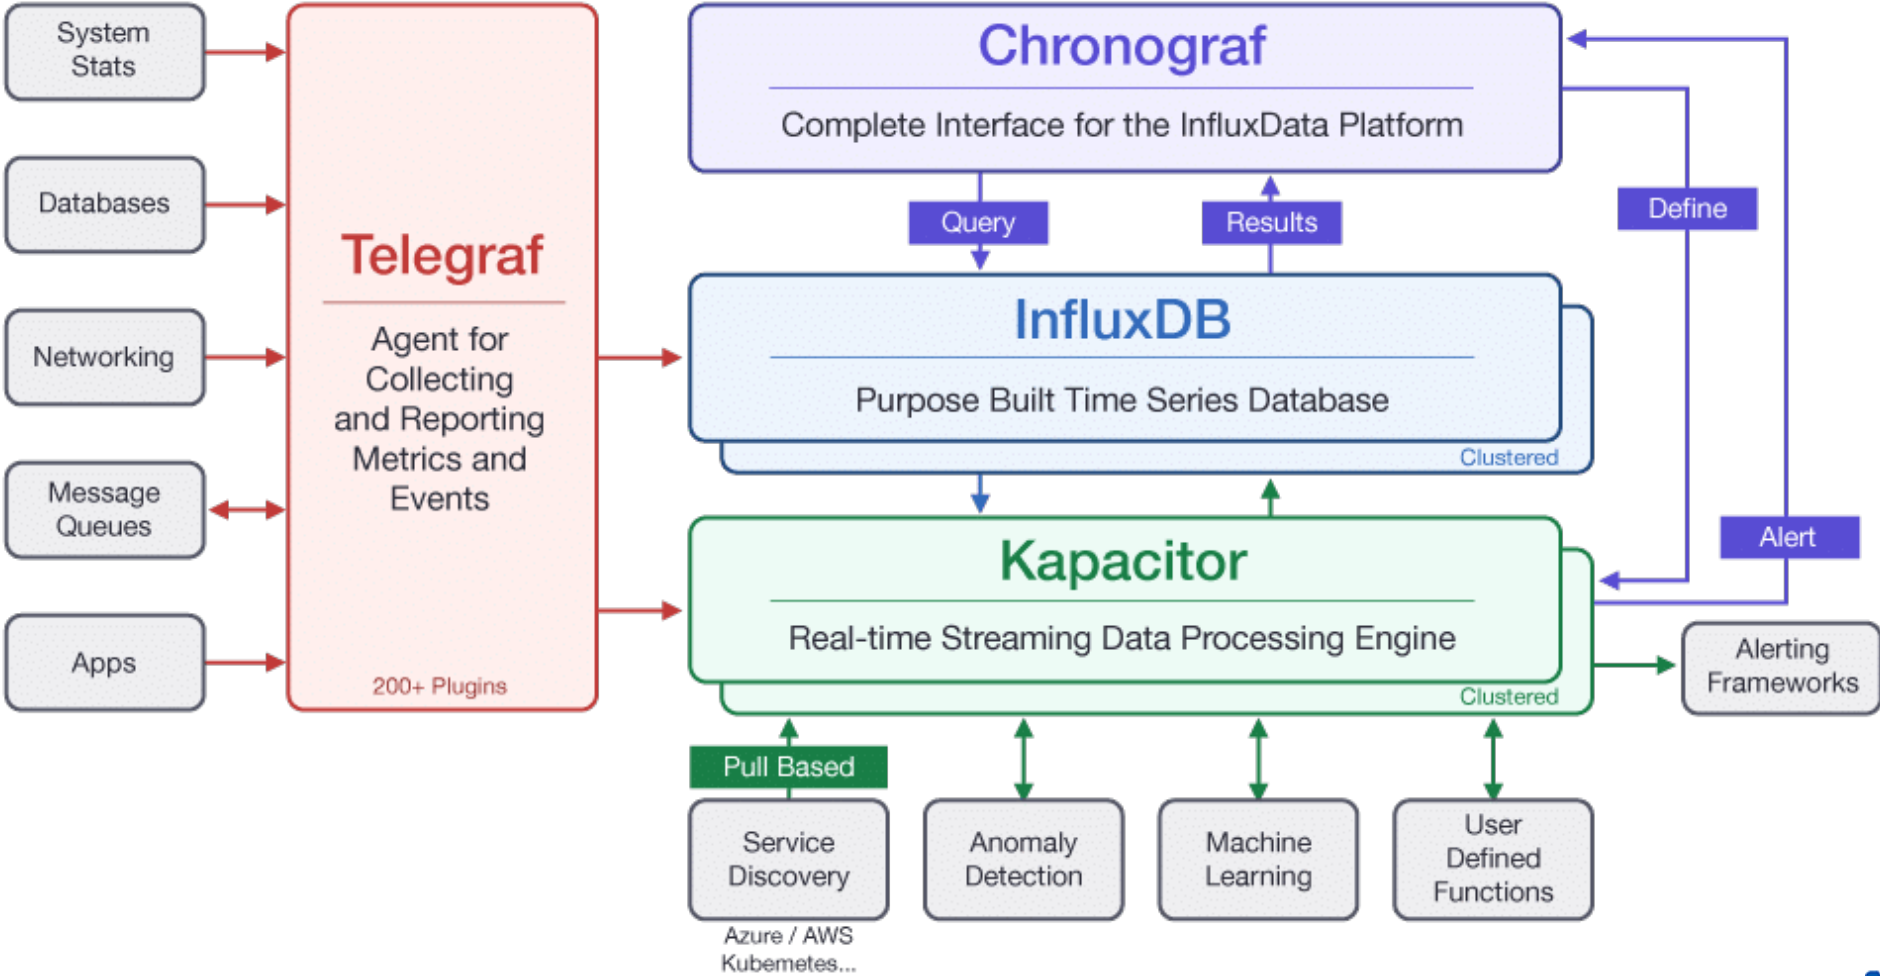
The Elastic Stack (ELK)



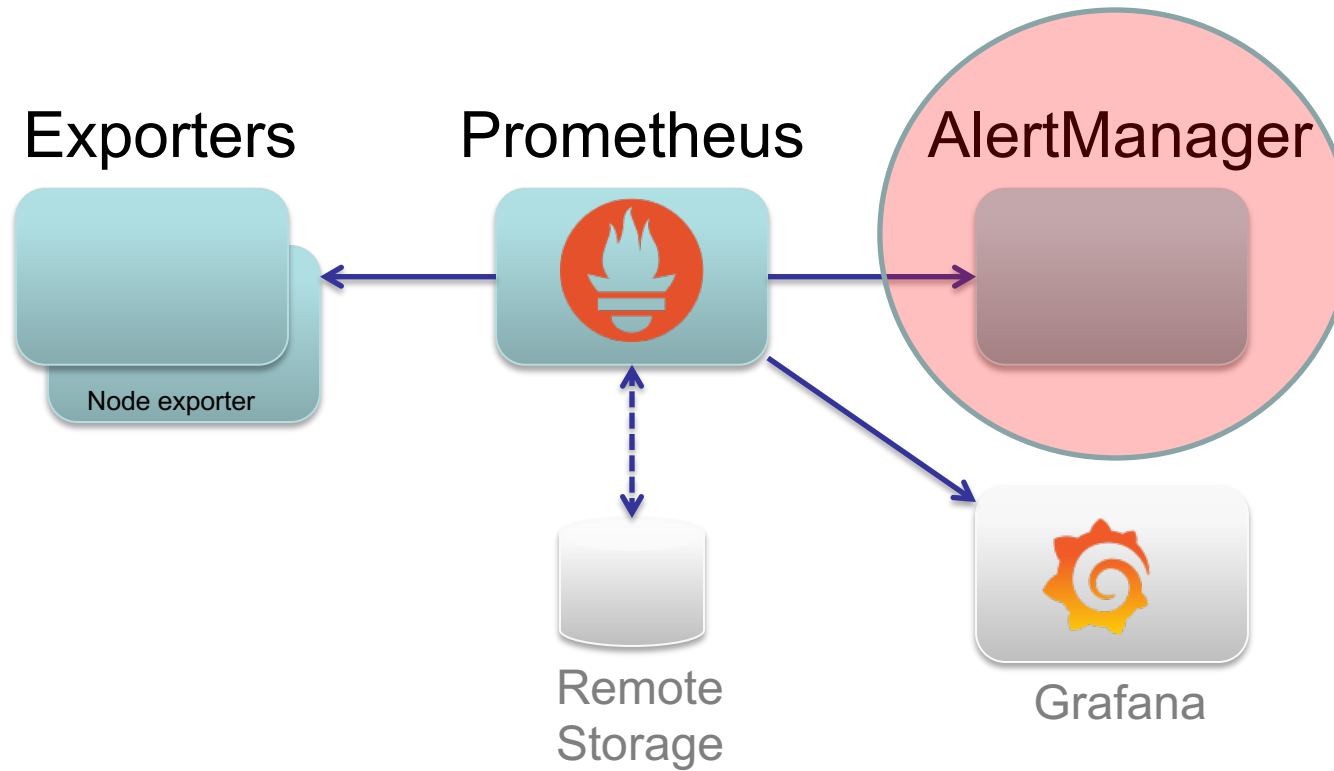
The TICK Stack



TICK stack detail



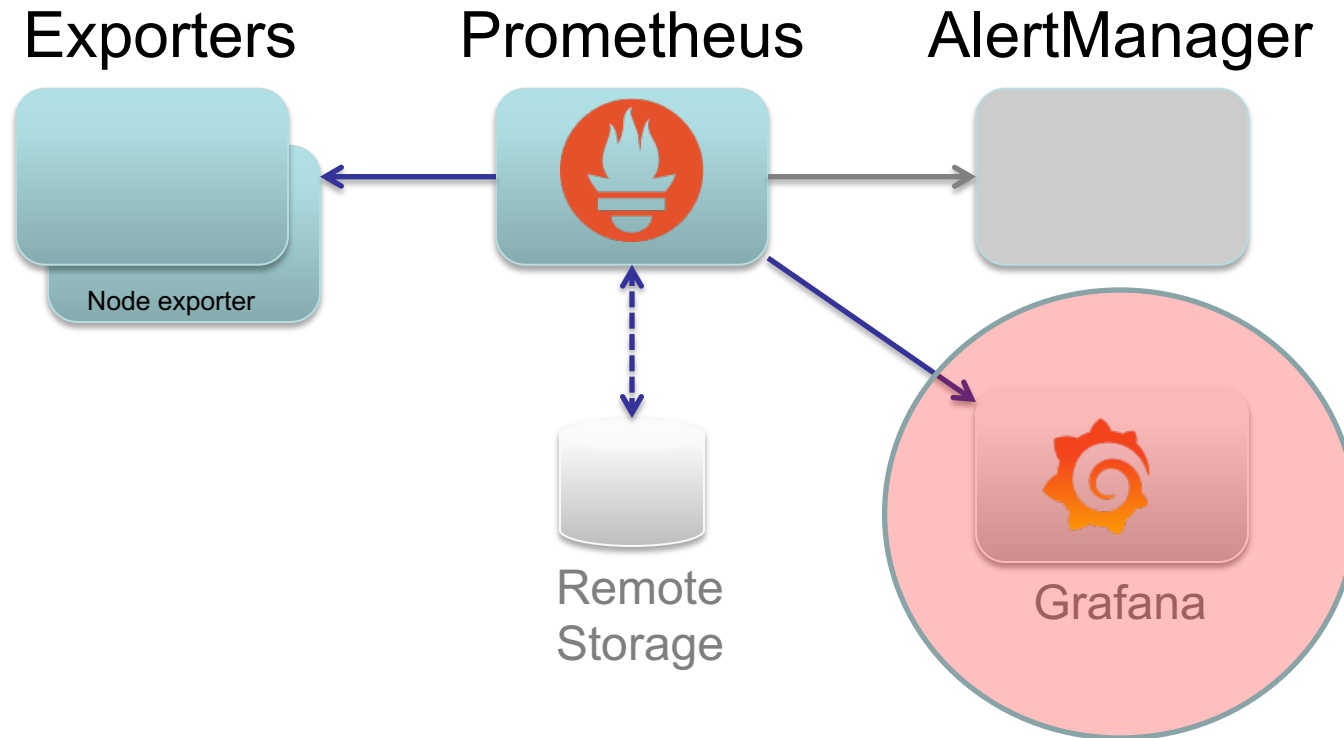
Prometheus



Or... →



Prometheus



NMM 2.0

Traditional vs. Present Day Practices

Putting it all together

- Presentation of data often requires more resources
 - Disk and CPU
 - Fine-grained telemetry (seconds or less vs. minutes) == more data on disk
 - Large data stores and complex dashboards can == more CPU
- Regex knowledge
 - You figure out what you want to know (some preconfigured dashboards as well)
 - Stack Based. Multiple software projects working together



logstash



elasticsearch



kibana



node_exporter



Prometheus



Promtail



Grafana loki



kibana



elasticsearch



logstash



FLOWS

node_exporter



Prometheus



SYSTEM METRICS

& ALERTS

Grafana



Grafana

influxdb



telegraf™

NETWORK

AND APP

METRICS

Grafana Loki



LOGS



Promtail



logstash



node_exporter

COLLECT



telegraf



Promtail



elasticsearch



Prometheus

STORE



influxdb



Grafana loki




kibana

VIEW



Grafana





Takes the following flow protocols

- ✓ Netflow
- ✓ IPFix
- ✓ SFlow




logstash



elasticsearch



kibana





node_exporter



Prometheus



Generate alerts for reachability and metrics

localhost:9090/alerts

Prometheus Alerts Graph Status Help

Alerts

HighErrorRate (1 active)

```

alert: HighErrorRate
expr: rate(hello_requests_total{status="500"}[1m]) > 0
for: 1m
labels:
  severity: hipchat
annotations:
  summary: High request latency

```

Labels	State	Active Since	Value
alertname="HighErrorRate" instance="go-mux-example:8080" job="mux" severity="hipchat" status="500"	FIRING	2017-11-26 03:50:55.980225997 +0000 UTC	0.48888888888888889

6 alerts	0 Prometheus servers down	0 hosts down	6 disks nearly full
-------------	------------------------------	-----------------	------------------------

Time	Message	Alertname	Severity
2017-07-26 15:48:38	IP: 10.56.20.122 -> var : -10.20 / Usage: 17.78 %	DiskNearlyFull	1
2017-07-26 04:34:37	IP: 10.92.3.42 -> root : -1.03 / Usage: 68.92 %	DiskNearlyFull	1
2017-07-26 02:21:37	IP: 10.89.26.188 -> tmp : -1.04 / Usage: 55.29 %	DiskNearlyFull	1
2017-07-25 18:24:37	IP: 10.10.28.12 -> rootfs-var-lib-docker : -1.00 / Usage: 84.58 %	DiskNearlyFull	1
2017-07-25 07:15:38	IP: 10.91.232.42 -> opt : -1.06 / Usage: 69.91 %	DiskNearlyFull	1
2017-07-25 02:59:38	IP: 10.79.55.29 -> opt : -1.05 / Usage: 70.16 %	DiskNearlyFull	1



Promtail

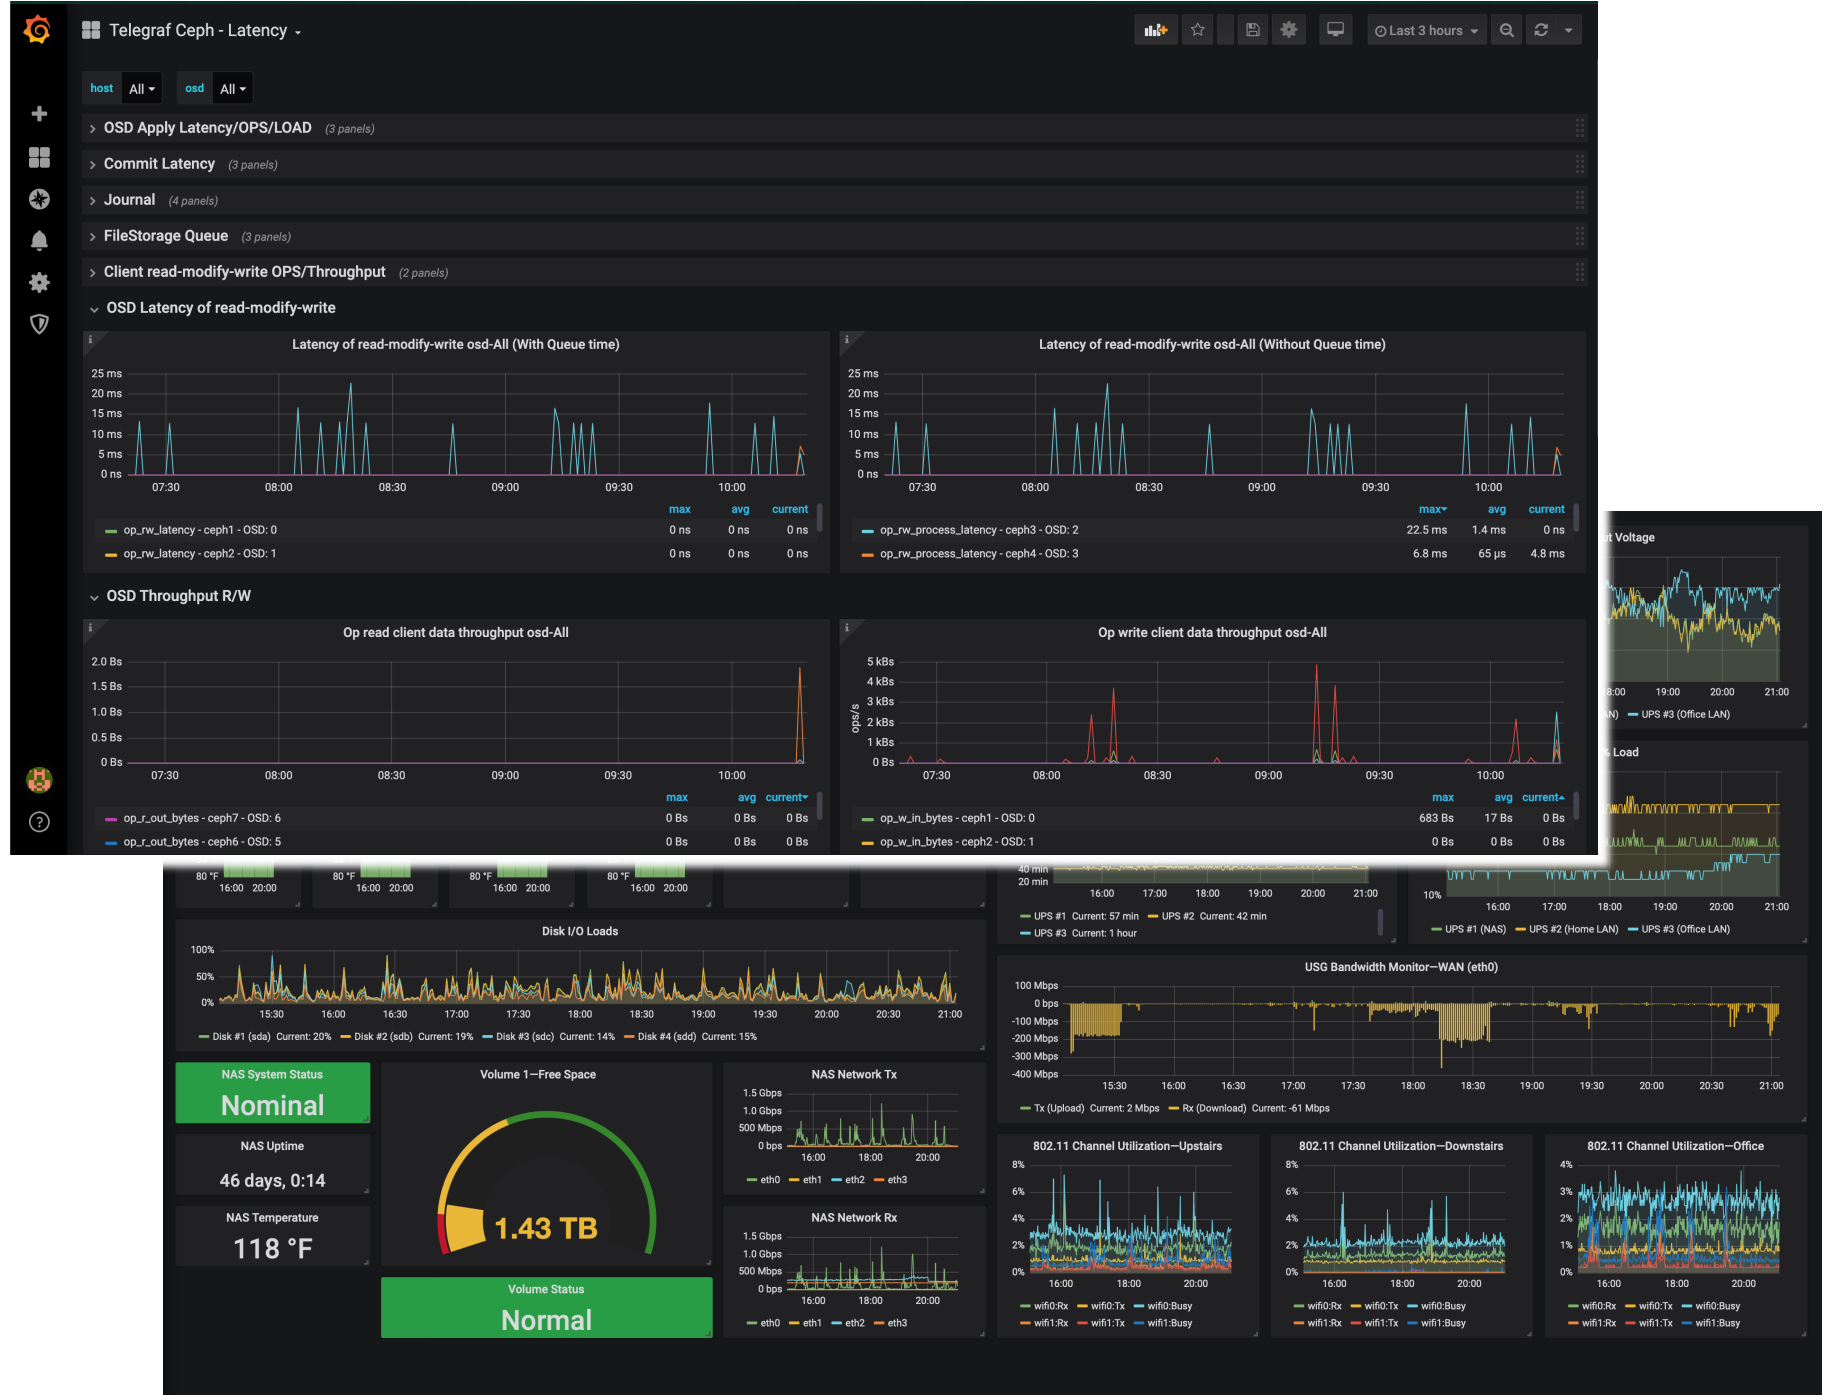


Grafana loki



- Streaming logs from files
- Works with Prometheus
- Kubernetes build available

The screenshot shows the Grafana Loki interface. At the top, there's a navigation bar with 'Explore', 'Loki', 'Metrics', and 'Logs' tabs. A search bar contains the query: `{job="varlogs"} |= "vtp" |= "UFW BLOCK"`. Below the search bar, there's a 'Log labels' section and an 'Add query' button. The main area displays a 'Logs' section with a bar chart showing log volume over time. Below the chart, there are toggle switches for 'Time', 'Unique labels', 'Wrap lines', and 'Dedup' (set to 'none'). The bottom section shows a list of log entries with their timestamps and details, such as: `> 2020-02-16 10:24:12 (no unique labels) Feb 15 23:24:12 vtp kernel: [353761.981013] [UFW BLOCK] IN=enx00e04c0612da OUT= MAC= SRC=fe80:0000:0000:0000:02e0:4cff:fe06:12da DST=ff12:0000:0000:0000:0000:0000:8384 LEN=370 TC=0 HOPLIMIT=1 FLOWLBL=284857 PROTO=UDP SPT=42765 DPT=21027 LEN=330`



Thanks!

Questions?



UNIVERSITY OF OREGON



References

- Cisco Telemetry with Google Protocol Buffers
<https://blogs.cisco.com/sp/streaming-telemetry-with-google-protocol-buffers>
- Cisco Model Driven Telemetry
<https://www.cisco.com/c/en/us/solutions/service-provider/cloud-scale-networking-solutions/model-driven-telemetry.html>
- Graphite
<https://graphiteapp.org/>
- InfluxDB
<https://www.influxdata.com/>
- Kafka
<https://docs.confluent.io/current/streams-ksql.html>



References

- Logz.io (Information on *Elastic Stack*, others)
<https://logz.io/>
- Monitoring vs. Observing
<https://twitter.com/isotopp/status/1328653624470331392>
- Prometheus
<https://prometheus.io/>
- Splunk
<https://www.splunk.com/>
- Tick Stack on CentOS
<https://www.digitalocean.com/community/tutorials/how-to-monitor-system-metrics-with-the-tick-stack-on-centos-7>
- TimescaleDB
<https://www.timescale.com/>



References from Dean 😊

- to docker-compose stacks
 - <https://github.com/robcowart/elastiflow>
 - <https://github.com/grafana/loki>
 - <https://github.com/nicolargo/docker-influxdb-grafana>
 - <https://github.com/vegasbrianc/prometheus>
- Other
 - <https://peter.run/blog/2019-07-28-visualising-latency-variance-in-grafana-in-2019/>
 - <https://hveem.no/visualizing-latency-variance-with-grafana>

