



KubeCon



CloudNativeCon

North America 2019



Great Cardinality Disasters of Our Time



KubeCon



CloudNativeCon

North America 2019

Bryan Boreham (@bboreham)

Chris Marchbanks (@csmarchbanks)



weaveworks

splunk>





Who are we?



KubeCon



CloudNativeCon

North America 2019

Bryan Boreham

Director of Engineering,
Weaveworks

Maintainer on Cortex and CNI

Part-time SRE on Weave Cloud -
GitOps and Prometheus-as-a-
Service

Chris Marchbanks

Senior Software Engineer, Splunk

Maintainer on Cortex,
Prometheus

Building internal observability
platform for Splunk Cloud

Why are we talking about “Cardinality”?



KubeCon



CloudNativeCon

North America 2019

Prometheus blows up.

CAUTION: Remember that every unique combination of key-value label pairs represents a new time series, which can dramatically increase the amount of data stored. Do not use labels to store dimensions with high cardinality (many different label values), such as user IDs, email addresses, or other unbounded sets of values.

Every metrics system runs into issues with Cardinality

How much is too much?



KubeCon



CloudNativeCon

North America 2019

For a single metric, queries are painfully slow above ~100,000 series.

E.g. an instant query like `count(metric_a)`

- 100,000 series in ~1.5 seconds
- 200,000 series in ~5 seconds

A single Prometheus can handle >10,000,000 series in memory

- Startup can take 15+ minutes
- Heap can exceed 100 GB

Disaster #1: Alerts



KubeCon



CloudNativeCon

North America 2019



Disaster #1: First thing we knew




KubeCon



CloudNativeCon

North America 2019

Memory usage



Disaster #1: The explanation



KubeCon



CloudNativeCon

North America 2019

User had a rule configured:

```
alert: InvalidTagSequenceTotal
expr:  invalidtag_sequence_total > 0
for:   1m
```

`count(invalidtag_sequence_total>0)` was over 40,000.

Prometheus creates two time-series per alert, so it was trying to update 80,000 series every 15 seconds.

Total of 300,000 alerts over the whole period.

Disaster #1: What to learn



KubeCon



CloudNativeCon

North America 2019

Careful with the query in alerts

```
expr: count(invalidtag_sequence_total > 0) > 0
```

Or perhaps:

```
expr: topk(10, invalidtag_sequence_total > 0)
```

Disaster #2: 100KB label values



KubeCon



CloudNativeCon

North America 2019

A bug in Istio wrote raw error messages to the err label, causing Prometheus label values to be 100KB unique values.

Example label value

```
Error adding/updating cluster(s) outbound|50052|clientservice.clientnamespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|clientservice.clientnamespace.svc.cluster.local: Invalid path:
/etc/certs/cert-chain.pem, outbound|50052|clientservice.clientnamespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|clientservice.clientnamespace.svc.cluster.local: Invalid path:
/etc/certs/cert-chain.pem, outbound|50052|clientservice.clientnamespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|8443|aservice.somenamespace.svc.cluster.local: Invalid path:
/etc/certs/cert-chain.pem, outbound|50051|anotherservice.anotherservicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|9091|service3.yetanother.servicename.namespace.svc.cluster.local: Invalid
path: /etc/certs/cert-chain.pem, outbound|50051|service4.yetanother.servicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|9000|service4.yetanother.servicename.namespace.svc.cluster.local:
Invalid path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid
path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path:
/etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path:
/etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path:
/etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path:
/etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path:
/etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path: /etc/certs/cert-chain.pem, outbound|50052|servicename.servicename.namespace.svc.cluster.local: Invalid path:
```

And continuing for many more lines...

Disaster #2: Full Prometheus Outage



KubeCon



CloudNativeCon

North America 2019


Prometheus instances all OOMed simultaneously.



Growth of the symbol table in Prometheus:

`prometheus_tsdb_symbol_table_size_bytes`

Symbol table size grew from 5 GB to 60 GB.



Disaster #2: What to learn



KubeCon



CloudNativeCon

North America 2019

Do not put raw messages into a label value!

Prefer a fixed size enum, and logging the raw messages.

Until the source can be fixed, use a relabeling rule to drop the offending metric.

Even well-used projects can have bugs in them.

Thank you to the Istio team for quickly fixing our issue.

Disaster #3: Buckets




KubeCon



CloudNativeCon

North America 2019

A Prometheus Histogram counts the number of samples in each bucket.



Disaster #3: Buckets




KubeCon



CloudNativeCon

North America 2019



Disaster #3: Buckets




KubeCon



CloudNativeCon

North America 2019



20ms

40ms

60ms

80ms

100ms
SLO

160ms





Disaster #3: kube-dns histogram



KubeCon



CloudNativeCon

North America 2019

github.com/kubernetes/dns/pkg/sidecar/dnsprobe.go

```
prometheus.HistogramOpts{
    Name:      "kubedns_latency_ms",
    Buckets:   prometheus.LinearBuckets(0, 10, 500),
}
```

Five. Hundred. Buckets.

Bryan fixed it for you.




KubeCon



CloudNativeCon

North America 2019

pkg/sidecar/dnsprobe.go 

```
@@ -95,7 +95,7 @@ func (p *dnsProbe) registerMetrics(options *Options) {  
    Subsystem: dnsProbeSubsystem,  
    Name:      p.Label + "_latency_ms",  
    Help:      "Latency of the DNS probe request " + p.Label,  
-    Buckets:  prometheus.LinearBuckets(0, 10, 500),  
+    Buckets:  prometheus.ExponentialBuckets(0.25, 2, 16), // from 0.25ms to 8 seconds  
})  
prometheus.MustRegister(p.latencyHistogram)
```

Disaster #4: All the metric names!



KubeCon



CloudNativeCon

North America 2019

A team added Prometheus metrics using a library that didn't expose tags. What can you do? Try putting them in names?

```
my_service_http_<route>_<tenant>_<status_code>_total
```

Queries like:

```
sum({__name__=~"my_service_http_getInfo_.*_total"})
```

Do not actually do this...

Disaster #4: Sad browsers



KubeCon



CloudNativeCon

North America 2019

Prometheus /graph UI becomes unusable

Prometheus (and Grafana) pulls back all metric names for autocomplete

Some improvements have been made, but still work to be done



Aw, Snap!

Something went wrong while displaying this webpage.

[Learn more](#)

Reload

How many metric names?!



KubeCon



CloudNativeCon

North America 2019



commented on Sep 3



```
curl http://prometheus:9090/api/v1/label/__name__/values > metricnames.json
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           % Dload  % Upload   Total     Spent    Left      Speed
100  121M    0  121M    0     0  1191k      0  --:--:--  0:01:44  --:--:-- 1551k

cat metricnames.json | tr "' ' ' | tr -d ", " | wc -w
2129511
```

We have about 2 million metric names.

Disaster #4: What to learn



KubeCon



CloudNativeCon

North America 2019

Do not put labels into a metric name.

The overall cardinality will be the same, but leads to many inefficiencies.

If needed, you can use relabelling rules to parse labels out of names - but it is better to fix them at the source.

How to troubleshoot?



KubeCon



CloudNativeCon

North America 2019

Extreme cardinality can be unpleasant to deal with.

What are the techniques to apply when you think it is happening?

```
count by(__name__)({"__name__=~".+"})
```

How to troubleshoot?



KubeCon



CloudNativeCon

North America 2019

In larger instances, `__name__` queries stop working. Instead,

1. How many series are in memory?
 - `prometheus_tsdb_head_series`
2. How much space do my unique strings take?
 - `prometheus_tsdb_symbol_table_size_bytes`
3. In 2.14.0 use the status page! <http://prometheus:9090/status>
 - Summary cardinality stats for a variety of issues

Troubleshooting example



KubeCon



CloudNativeCon

North America 2019

My Symbol table size was greater than 1 GB in an instance.

Prometheus Alerts Graph Status ▾ Help

Label Names With Highest Cumulative Label Value Length

| Name | Length |
|--------------|----------|
| logfile | 11012076 |
| cluster | 645665 |
| container_id | 633567 |

What series have logfile?



KubeCon



CloudNativeCon

North America 2019

Query: `count({logfile!=""}) by (job, __name__)`

Result:

| Element | Value |
|--|-------|
| <code>mtail_log_rotations_total</code> <code>{job="mtail"}</code> | 783 |
| <code>mtail_log_lines_total</code> <code>{job="mtail"}</code> | 48407 |

Results



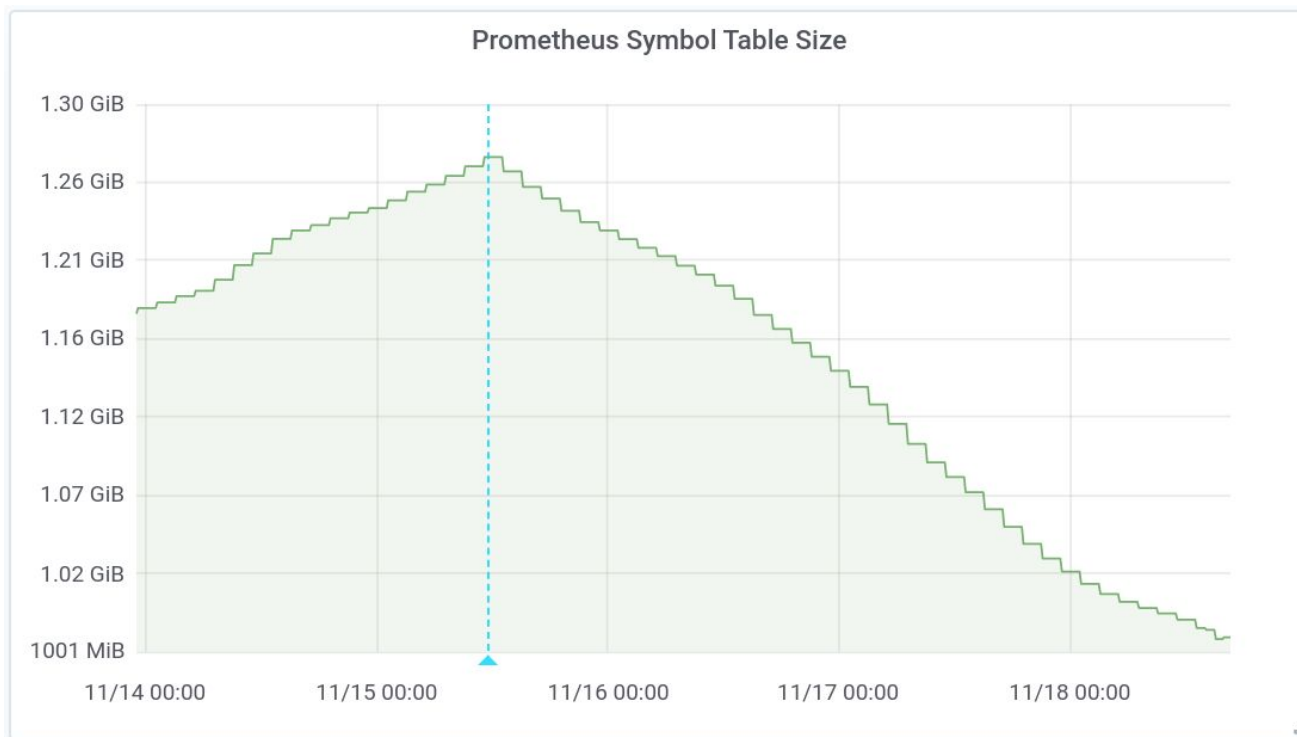
KubeCon



CloudNativeCon

North America 2019

Reduce the number of containers mtail is following.



What cardinality can you find?



Questions?



KubeCon



CloudNativeCon

North America 2019

Bryan Boreham (@bboreham)

Chris Marchbanks (@csmarchbanks)



weaveworks

splunk>

