



@snehainguva

observability and product release:  
**leveraging prometheus**  
to build and test new products



## about me

**software engineer @DigitalOcean**  
currently **network services**  
**<3 cats**



some **stats**



**90M+** timeseries

**85** instances of prometheus

**1.7M+** samples/sec



# the **history**



# ye' olden days

use nagios + various plugins to monitor

use collectd + statsd + graphite

openTSDB



# lovely prometheus

white-box monitoring

multi-dimensional data model

fantastic querying language





# glorious kubernetes

easily deploy and update services

scalability

combine with prometheus + alertmanager



# sneha joins networking

set up monitoring for VPC

working on DHCP

**how can we use prometheus even before release?**



the plan:

✓ observability DigitalOcean

**build --- instrument --- test --- iterate**

examples



**metrics:** *time-series of sampled data*

**tracing:** *propagating metadata through different requests, threads, and processes*

**logging:** *record of discrete events over time*



**metrics:**  
*what do we measure?*



# *four golden signals*



**latency:** *time to service a request*

**traffic:** *requests/second*

**error:** *error rate of requests*

**saturation:** *fullness of a service*



**U**tilization

**S**aturation

**E**rror rate





*“**USE** metrics often allow users to solve 80% of server issues with 5% of the effort.”*



the plan:

- ✓ observability DigitalOcean
- ✓ **build --- instrument --- test --- iterate**

examples



## **build:**

design the service

write it in go

use internally shared libraries



## build: doge/dorpc - shared rpc library

```
var DefaultInterceptors = []string{ StdLoggingInterceptor, StdMetricsInterceptor, StdTracingInterceptor }
```

```
func NewServer(opt ...ServerOpt) (*Server, error) {  
    opts := serverOpts{  
        name:          "server",  
        clientTLSAuth:  tls.VerifyClientCertIfGiven,  
        intercept:      interceptor.NewPathInterceptor(interceptor.DefaultInterceptors...),  
        keepAliveParams: DefaultServerKeepAlive,  
        keepAliveEnforce: DefaultServerKeepAliveEnforcement,  
    }  
    ...  
}
```



## **instrument:**

send logs to centralized logging

send spans to trace-collectors

set up prometheus metrics



# metrics instrumentation: go-client

```
func (s *server) initializeMetrics() {  
    s.metrics = metricsConfig{  
        attemptedConvergeChassis: s.metricsNode.Gauge("attempted_converge_chassis", "number of chassis  
converger attempting to converge"),  
        failedConvergeChassis: s.metricsNode.Gauge("failed_converge_chassis", "number of chassis that failed to  
converge"),  
    }  
}
```

```
func (s *server) ConvergeAllChassis(...) {  
    ...  
    s.metrics.attemptedConvergeChassis(float64(len(attempted)))  
    s.metrics.failedConvergeChassis(float64(len(failed)))  
    ...  
}
```



## Quick Q & A: Collector Interface

*// A collector must be registered.*

```
prometheus.MustRegister(collector)
```

```
type Collector interface {
```

*// Describe sends descriptors to channel.*

```
Describe(chan<- *Desc)
```

*// Collect is used by the prometheus registry on a scrape.*

*// Metrics are sent to the provided channel.*

```
Collect(chan<- Metric)
```

```
}
```



# metrics instrumentation: third-party exporters

Built using the collector interface

Sometimes we build our own

Often we use others:

[github.com/prometheus/\*\*mysqld\*\*\\_exporter](https://github.com/prometheus/mysqld_exporter)

[github.com/kbudde/\*\*rabbitmq\*\*\\_exporter](https://github.com/kbudde/rabbitmq_exporter)

[github.com/prometheus/\*\*node\*\*\\_exporter](https://github.com/prometheus/node_exporter)

[github.com/digitalocean/\*\*openvswitch\*\*\\_exporter](https://github.com/digitalocean/openvswitch_exporter)





# metrics instrumentation: in-service collectors

```
type RateMap struct {
    mu      sync.Mutex
    ...
    rateMap map[string]*rate
}

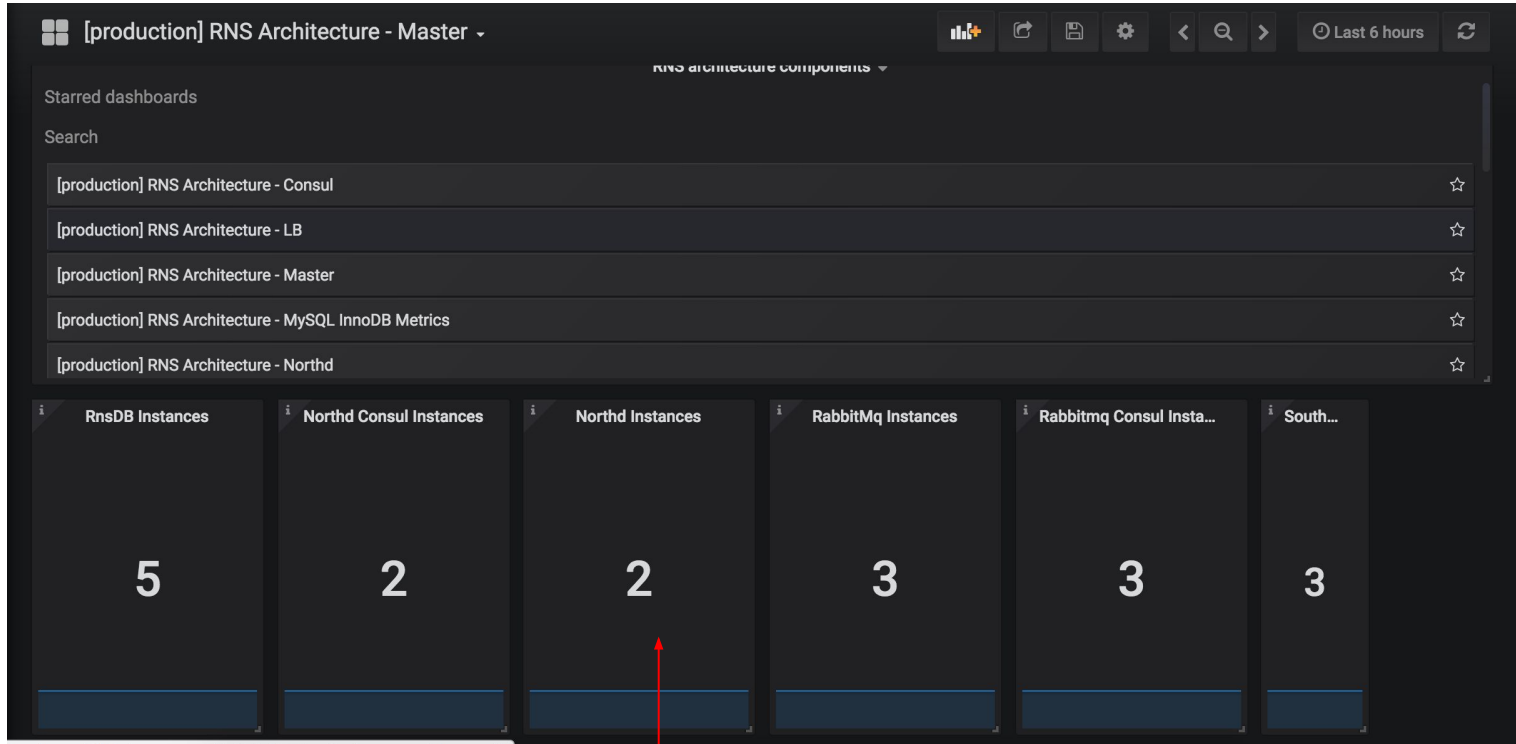
var _ prometheus.Collector = &RateMapCollector{}

func (r *RateMapCollector) Describe(ch chan<- *prometheus.Desc) {
    ds := []*prometheus.Desc{ r.RequestRate}
    for _, d := range ds {
        ch <- d
    }
}

func (r *RateMapCollector) Collect(ch chan<- prometheus.Metric) {
    ...
    ch <- prometheus.MustNewConstHistogram( r.RequestRate, count, sum, rateCount)
}
```



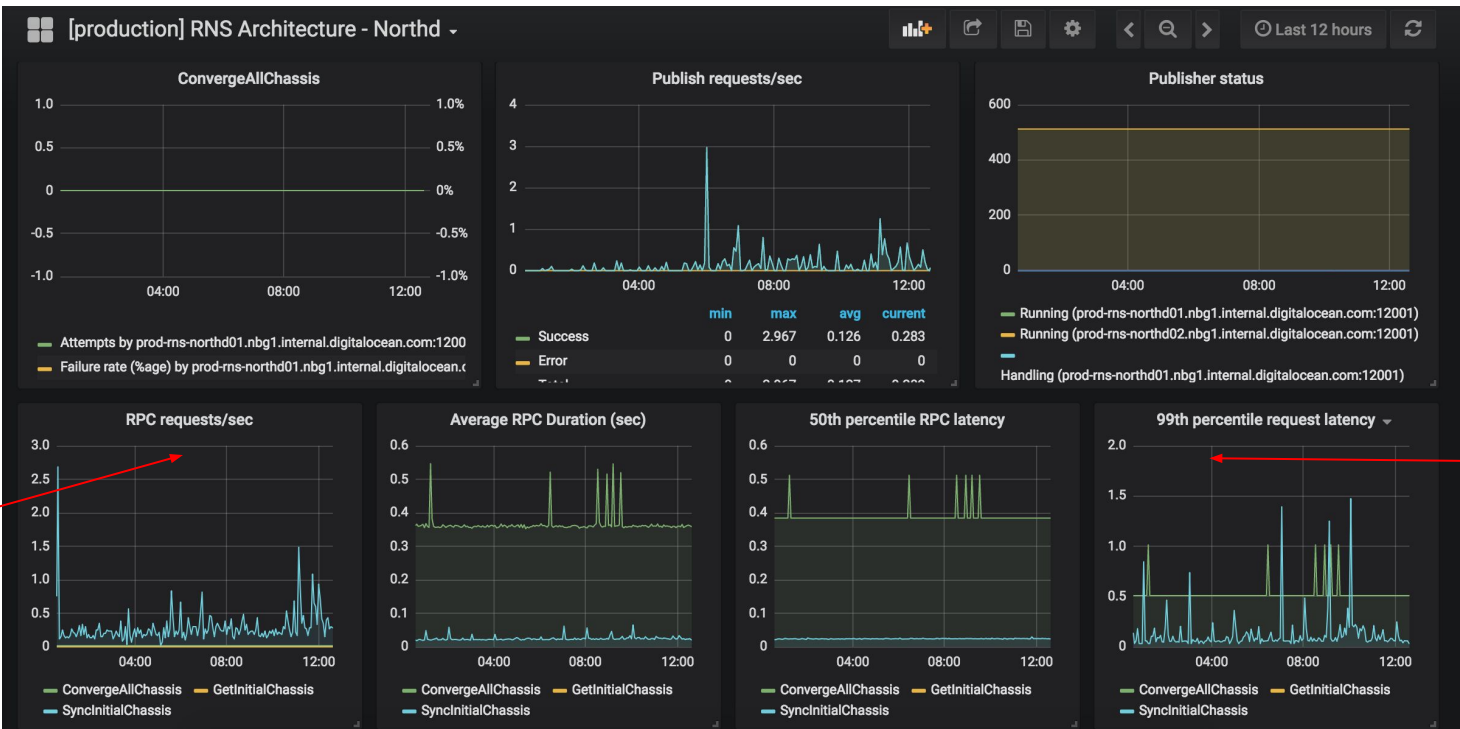
# metrics instrumentation: dashboards #1



state metrics

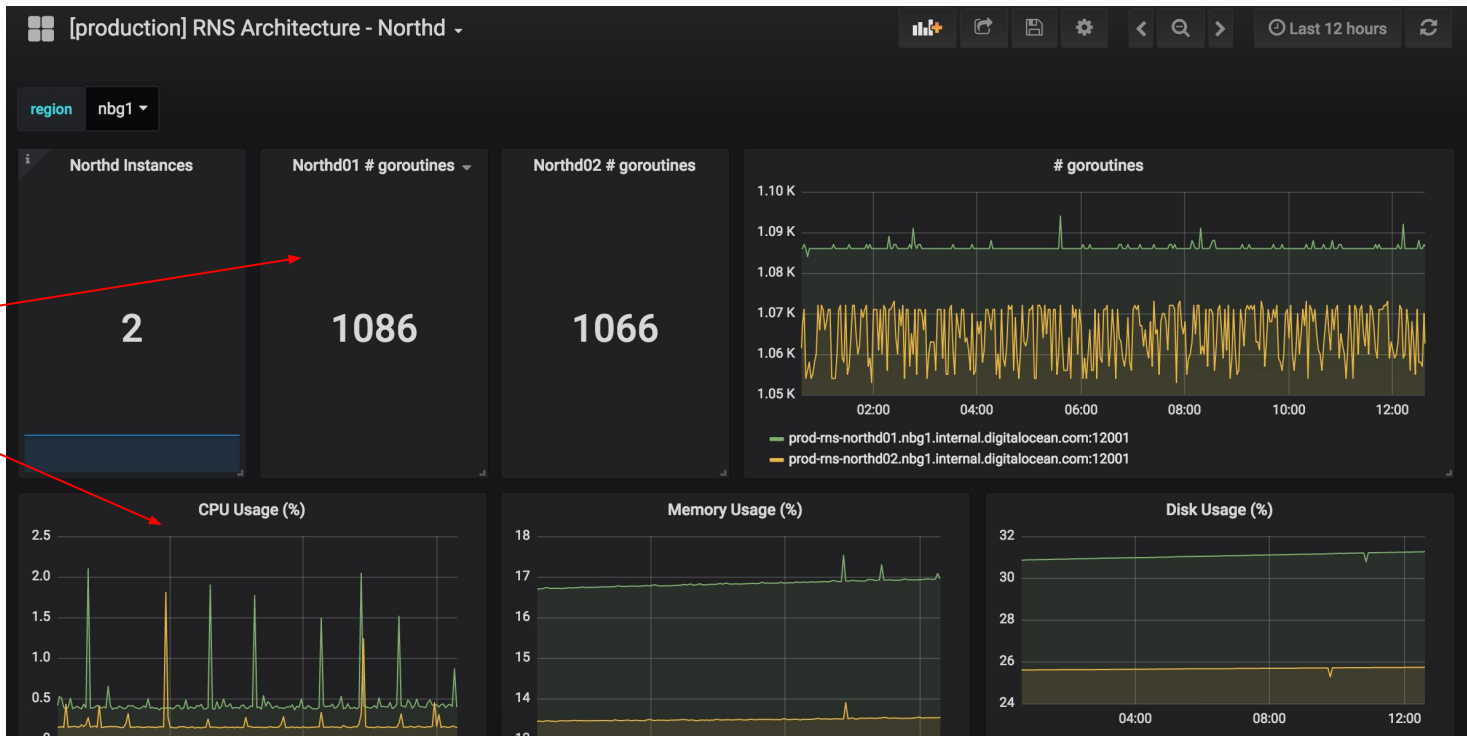


# metrics instrumentation: dashboard #2





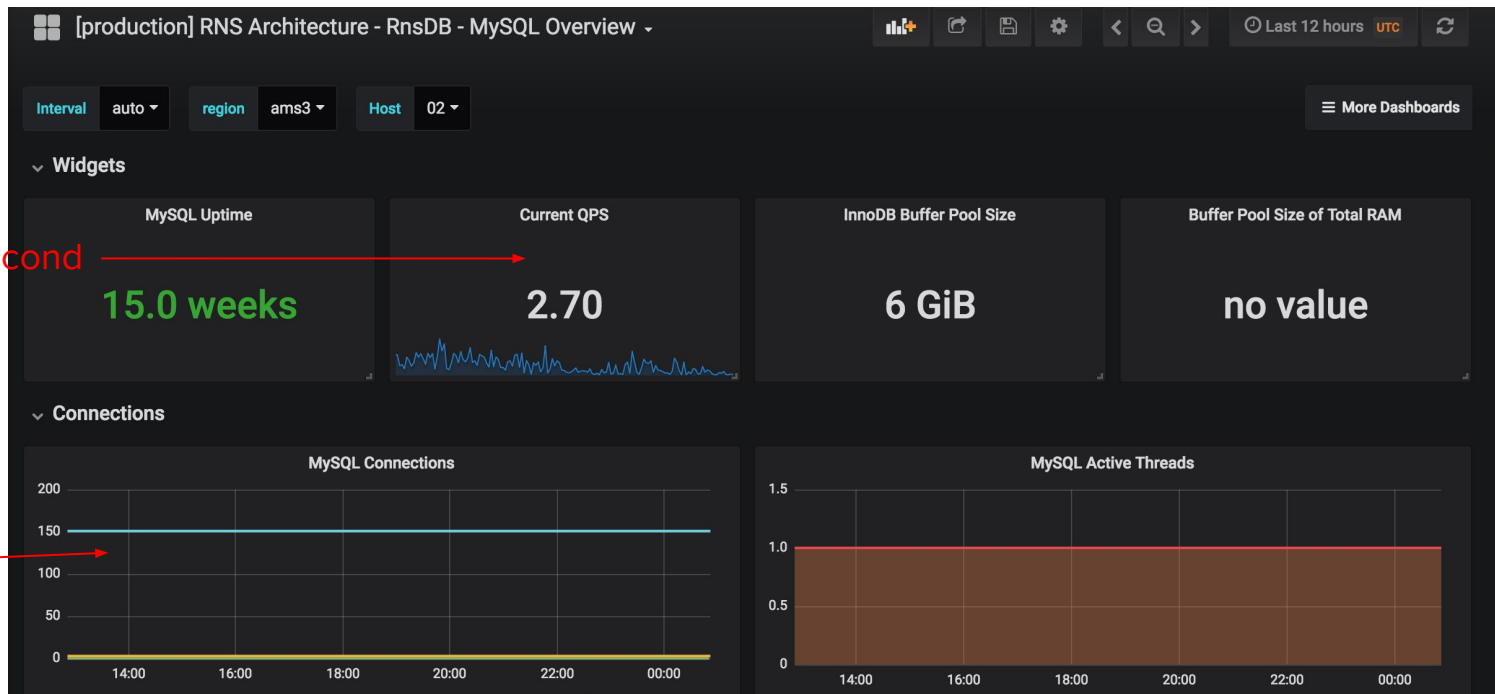
# metrics instrumentation: dashboard #3



utilization  
metrics



# metrics instrumentation: dashboard #4





# metrics instrumentation: dashboard #5



saturation  
metric





## **test:**

load testing:

grpc-clients and goroutines

chaos testing:

take down a component of a system

integration testing:

how does this feature integrate with the cloud?



## testing: identify key issues

how is our latency? —————→ *use tracing to dig down*

is there a goroutine leak? —————→ *use a worker pool*

does resource usage increase with traffic? —→ *use cpu and memory profiling*

is there a high error rate? —————→ *check logs for types of error*

how are our third-party services?





## testing: tune metrics + alerts

*do we need more labels for our metrics?*

*should we collect more data?*

***State-based alerting:*** *Is our service up or down?*

***Threshold alerting:*** *When does our service fail?*



# testing: documentation

*set up operational playbooks*  
*document recovery efforts*

## VPC - Recovery test notes

Created by jheimann, last modified by amigliaccio on Jul 19, 2018

[ [Installation](#) ] [ [Northd](#) ] [ [Southd](#) ] [ [Consul](#) ] [ [Rabbit](#) ] [ [Percona](#) ] [ [Alpha read replica](#) ] [ [Hvflowd](#) ]

### Northd

[VPC-387 - Simulate north / south failures](#) **DONE**

Two northd instances per region, in active-active failover configuration. Active-active operation is essential as the service is currently designed to be the sole operator of a single virtual chassis; any concurrent operation on a single chassis will produce a deadlock.

- `stage-rns-northd01.<region>`
- `stage-rns-northd02.<region>`



**iterate!**

(but really, let's look at some examples...)



## the plan:

- ✓ observability DigitalOcean
- ✓ **build --- instrument --- test --- iterate**
- ✓ examples

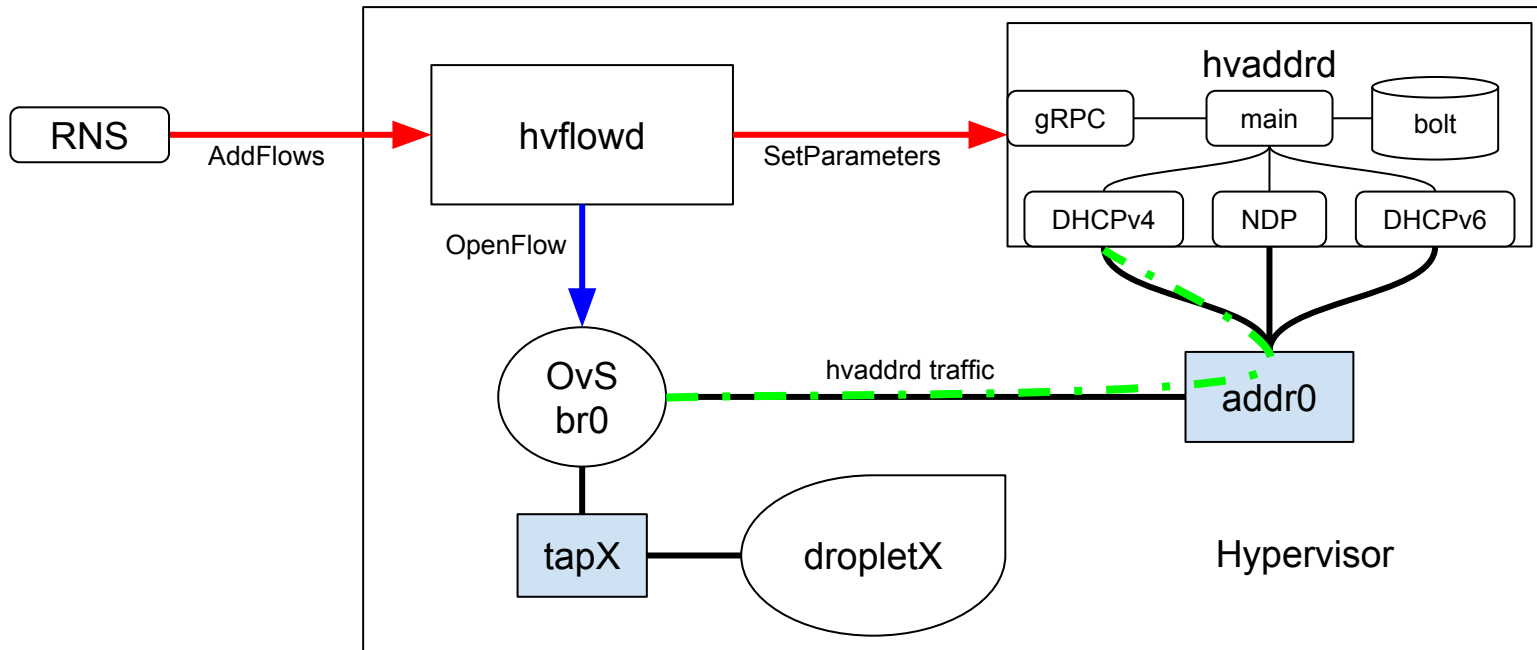


# product #1: DHCP

(hvaddr)

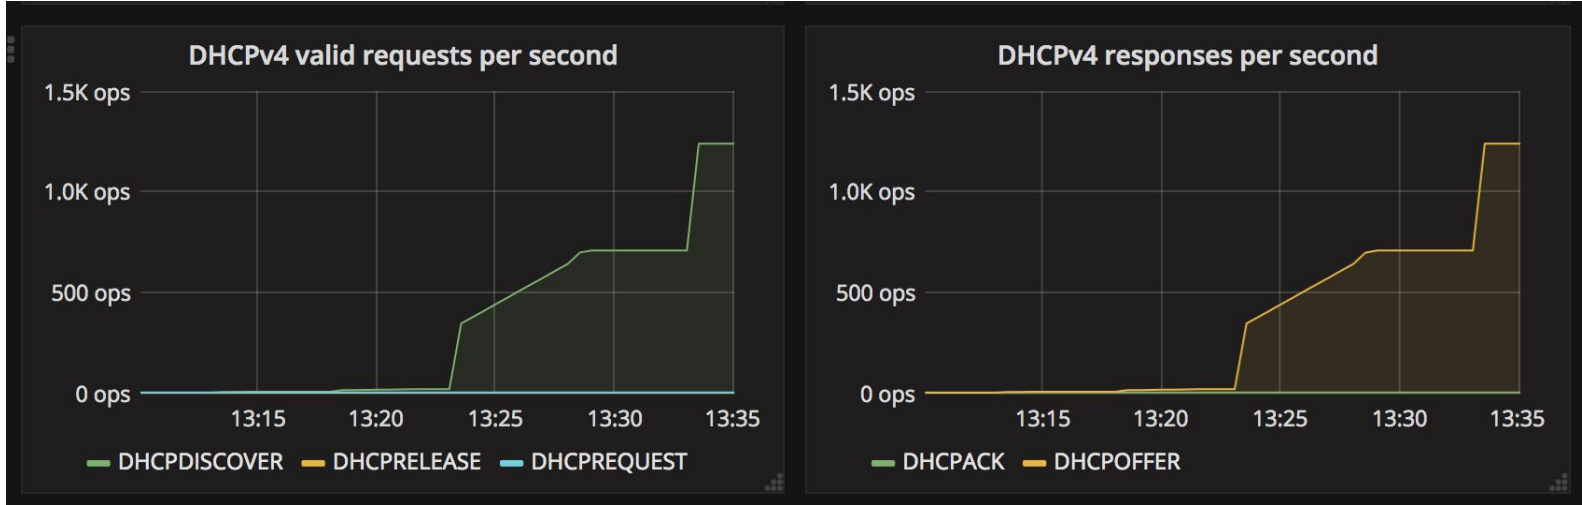


# product #1: DHCP






# DHCP: load testing





## DHCP: load testing (2)



**Matt Layher** 8:16 AM

managed to OOM hvaddr with an early hvaddrhammer tool:

```
[2586276.456342] Task in /hvaddr killed as a result of  
limit of /hvaddr
```





# DHCP: custom conn collector

```
package dhcp4conn —————→ Implements the net.conn interface
                    and allows us to process ethernet
                    frames for validation and other
                    purposes.


var _ prometheus.Collector = &collector{}

// A collector gathers connection metrics.

type collector struct {
    ReadBytesTotal    *prometheus.Desc
    ReadPacketsTotal *prometheus.Desc
    WriteBytesTotal   *prometheus.Desc
    WritePacketsTotal *prometheus.Desc
}
```



# DHCP: custom conn collector





# DHCP: goroutine worker pools

```
workC := make(chan request, Workers) —————> Uses buffered channel to process
for i := 0; i < Workers; i++ { requests, limiting goroutines and
    go func() { resource usage.
        defer workWG.Done()
        for r := range workC {
            s.serve(r.buf, r.from)
        }
    }()
}
```



# DHCP: rate limiter collector

```
type RateMap struct {  
    mu      sync.Mutex  
    ...  
    rateMap map[string]*rate  
}
```



ratemap calculates the exponentially weighted moving average on a per-client basis and limits requests

```
type RateMapCollector struct {  
    RequestRate *prometheus.Desc  
    rm          *RateMap  
    buckets    []float64  
}
```

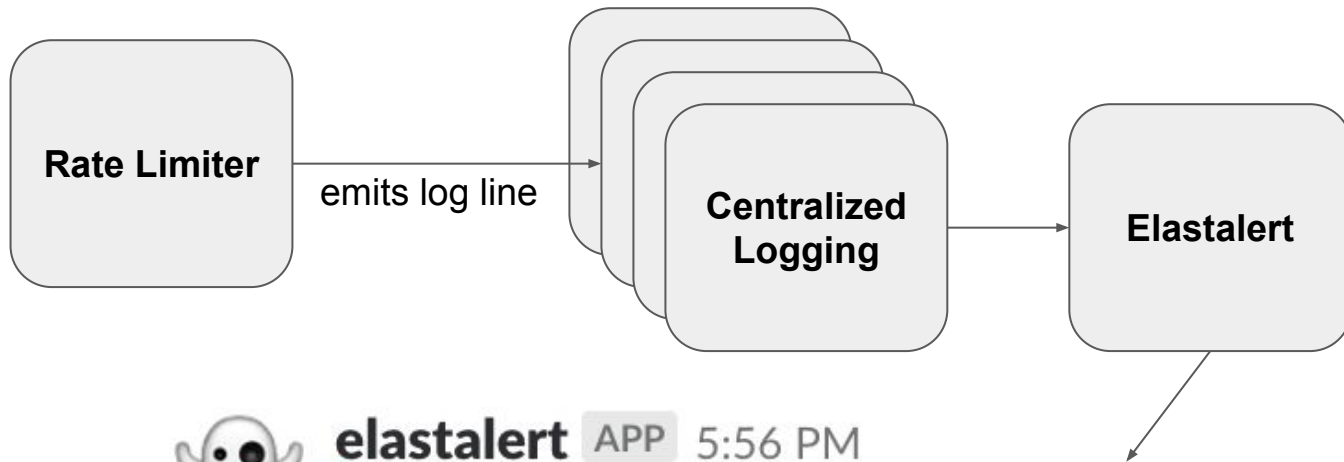


collector gives us a snapshot of rate distributions

```
func (r *RateMapCollector) Collect(ch chan<- prometheus.Metric) {  
    ...  
    ch <- prometheus.MustNewConstHistogram(  
        r.RequestRate,  
        count, sum,  
        rateCount)  
}
```



# DHCP: rate alerts

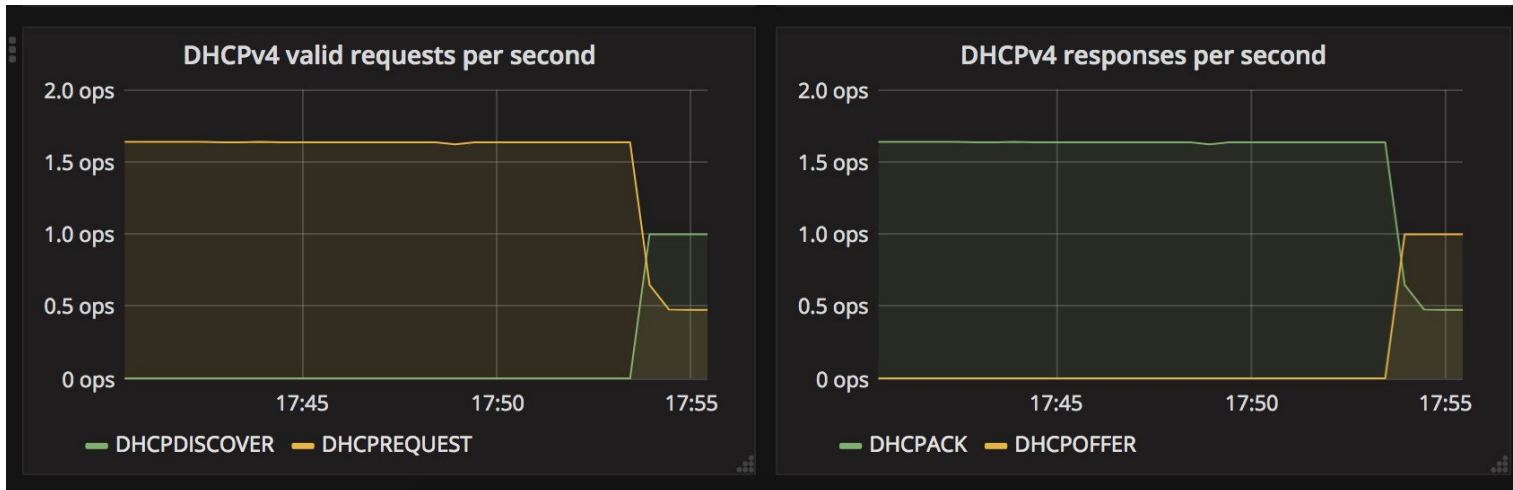


**elastalert** APP 5:56 PM

stage2 hvaddr abusive activity detected  
hvaddr abusive activity detected on host



# DHCP: the final result






# product #2: VPC



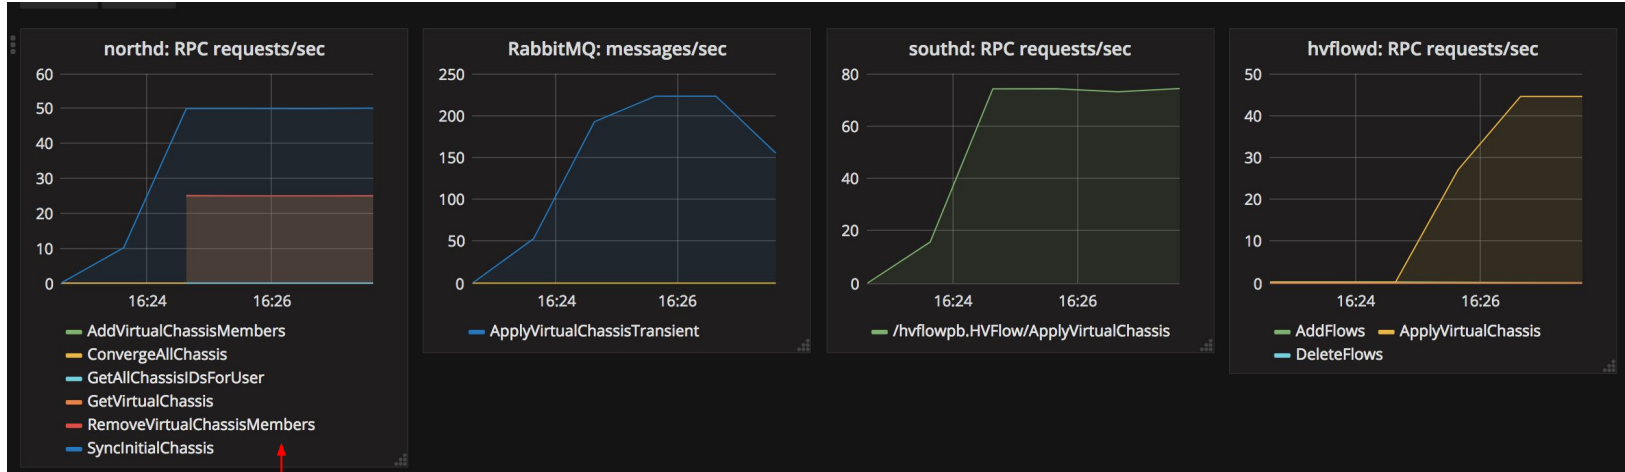
# product #2: VPC







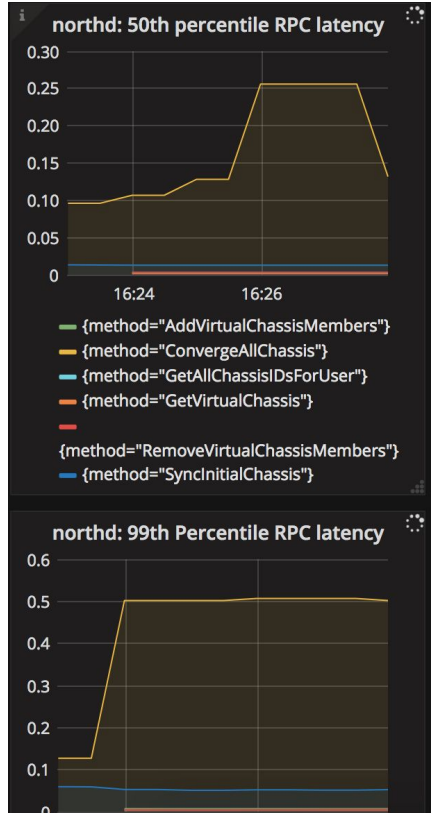
# VPC: load-testing



load tester repeatedly makes some RPC calls




# VPC: latency issues (1)



as load testing continued, started to notice latency in different rpc calls



# VPC: latency issues (2)

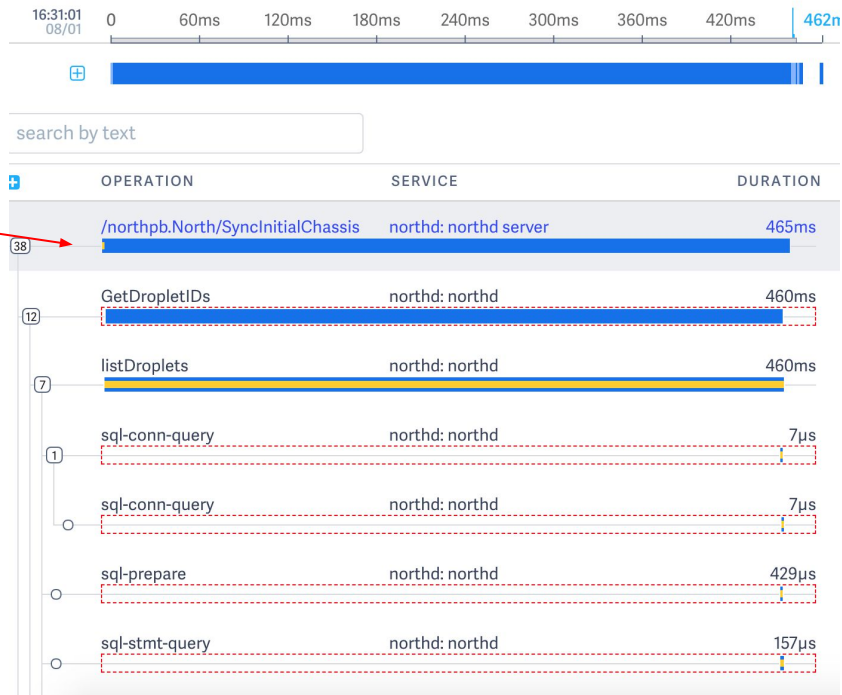


use tracing to take a look at the /SyncInitialChassis call



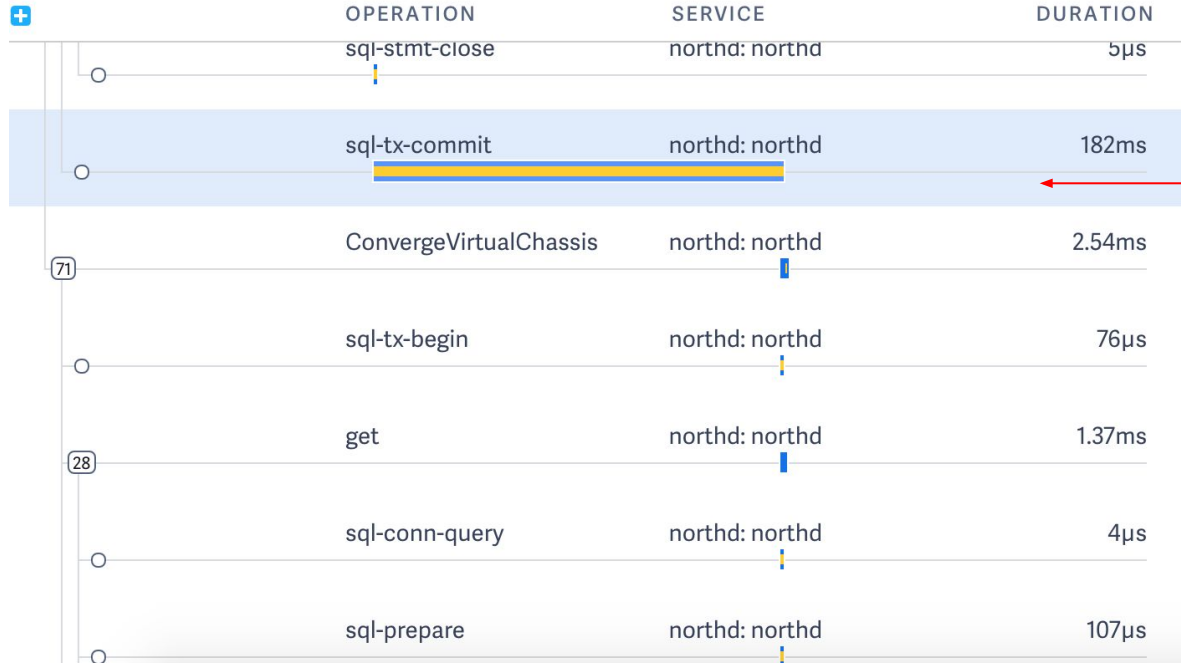
# VPC: latency issues (3)

Note that spans for some traces were being dropped. Slowing down the load tester, however, eventually ameliorated that problem.





# VPC: latency issues (4)

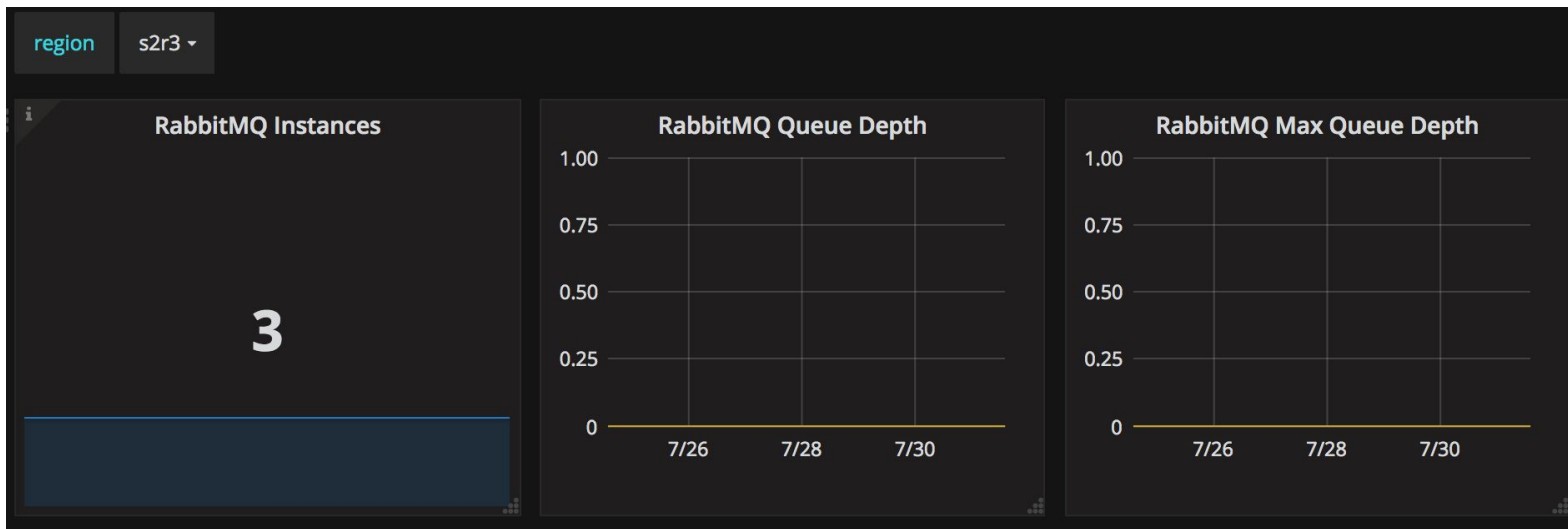


“The fix was to be smarter and do the queries more efficiently. The repetitive loop of queries to rnsdb really stood out in the lightstep data.”

- Bob Salmi



# VPC: remove component



can queue be replaced with simple request-response system?

Queues inevitably run in two states: full, or empty. If your queue is running full, you haven't pushed enough work to the edges, and if it is running empty, it's working as a slow load balancer.

source: <https://programmingisterrible.com/post/162346490883/how-do-you-cut-a-monolith-in-half>




# VPC: chaos testing

Induce northd failure  
and ensure failover  
works

Drop primary and  
recovery from  
secondary

Induce south service failure  
and see how rabbit  
responds





# VPC: add alerts (1)

state-based  
alerts

The screenshot shows a Slack message interface. On the left, a red text label "state-based alerts" has a red arrow pointing to a vertical red line that separates the message header from the body. The message header includes the title "StageNorthdDown" in blue, followed by "status: firing", "severity: warning", "source: Prometheus", and "description: northd instances down to 1 in s2r3". To the right of the message is a "Share message ..." button and a row of icons: a smiley face with a plus sign, a speech bubble with a plus sign, a right-pointing arrow, a star, and a three-dot menu. Below the message, a timestamp "7:43 PM" is visible. The message body contains the title "StageSouthdDown" in blue, followed by "status: firing", "severity: warning", "source: Prometheus", and "description: southd instances have dropped to : 1 in s2r3". At the bottom, a user profile picture is partially visible next to the name "Sa" and a timestamp "4 PM". The user's message reads "^^ thats me testing in s2r3".





# VPC: add alerts (2)

threshold alert



**AlertManager** APP 7:42 PM

**StageNorthdRpcErrorRateHigh**

status: firing

severity: warning

source: Prometheus

description: northd rpc error rate high 0.58 /s in s2r3



**AlertManager** APP 5:18 PM

**StageRnsMysqlReplicasDelayed**

status: firing

severity: warning

source: Prometheus Prometheus



# conclusion



## **what?**

*four golden signals, USE metrics*

## **when?**

*as early as possible*

## **how?**

*combine with profiling, logging, tracing*



**thanks!**

@snehainguva

