Grafana Labs

Lessons Learned Building a Global Synthetic Monitoring System

Suraj

Software Engineer, Grafana Labs

Grafana Labs

Poll: Are you familiar with Synthetic Monitoring?

Grafana Labs

Poll: Are you using Synthetic Monitoring?

Grafana Labs

What is Synthetic Monitoring?

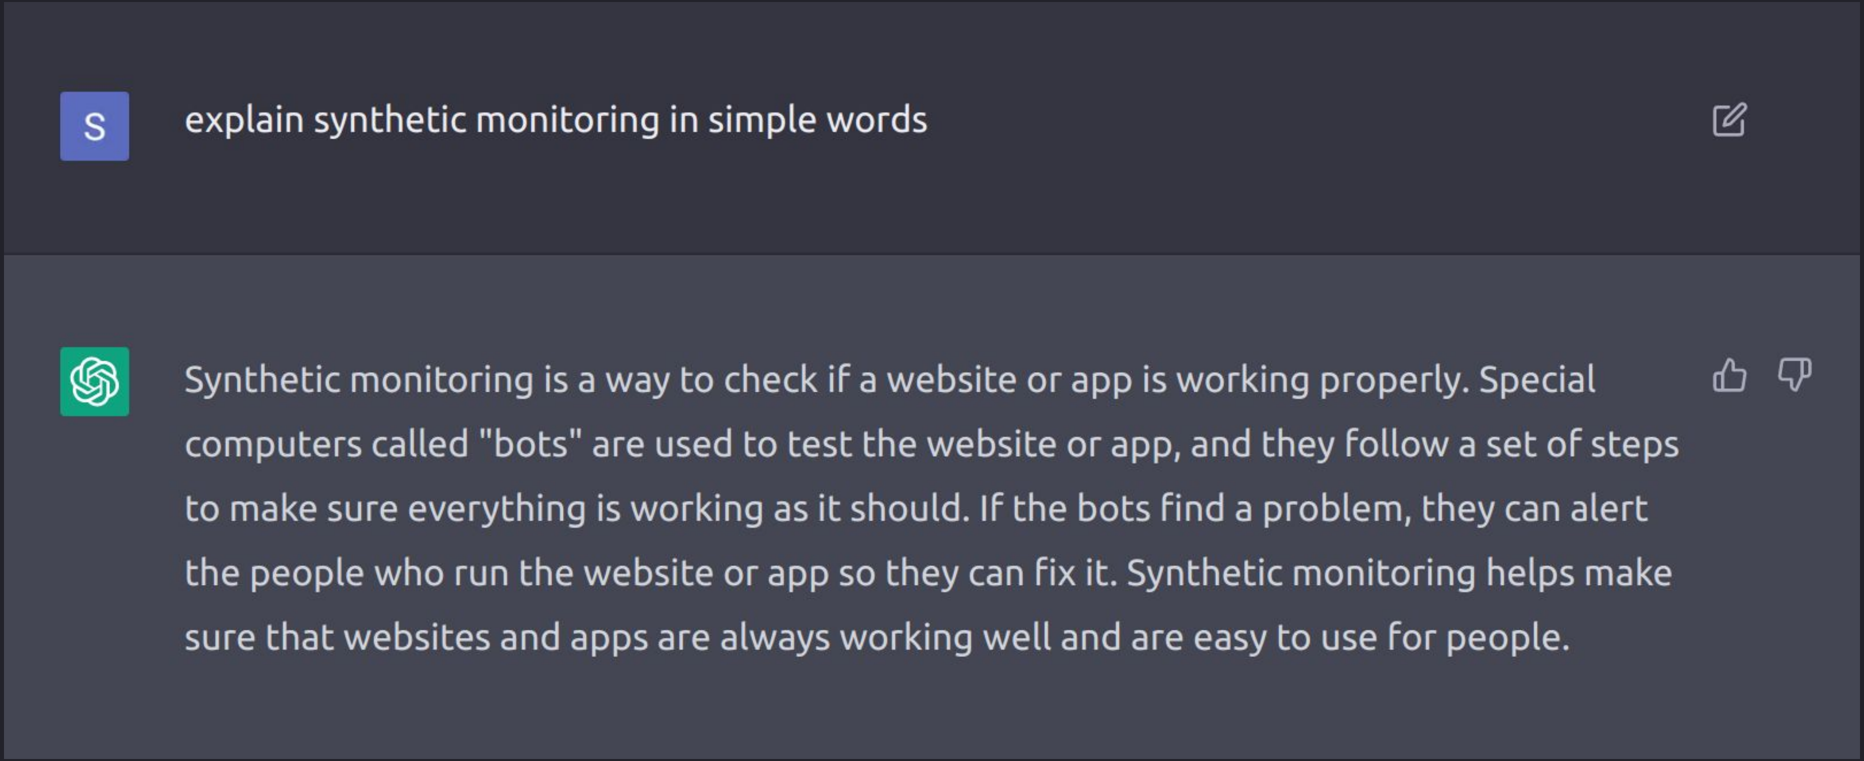

explain synthetic monitoring in simple words

Synthetic monitoring is a way to check if a website or app is working properly. Special computers called "bots" are used to test the website or app, and they follow a set of steps to make sure everything is working as it should. If the bots find a problem, they can alert the people who run the website or app so they can fix it. Synthetic monitoring helps make sure that websites and apps are always working well and are easy to use for people.

Grafana Labs

Why do we need Synthetic Monitoring?

Where does it fit in observability landscape?

Grafana Labs

Terms used in this talk

Probe / Agent

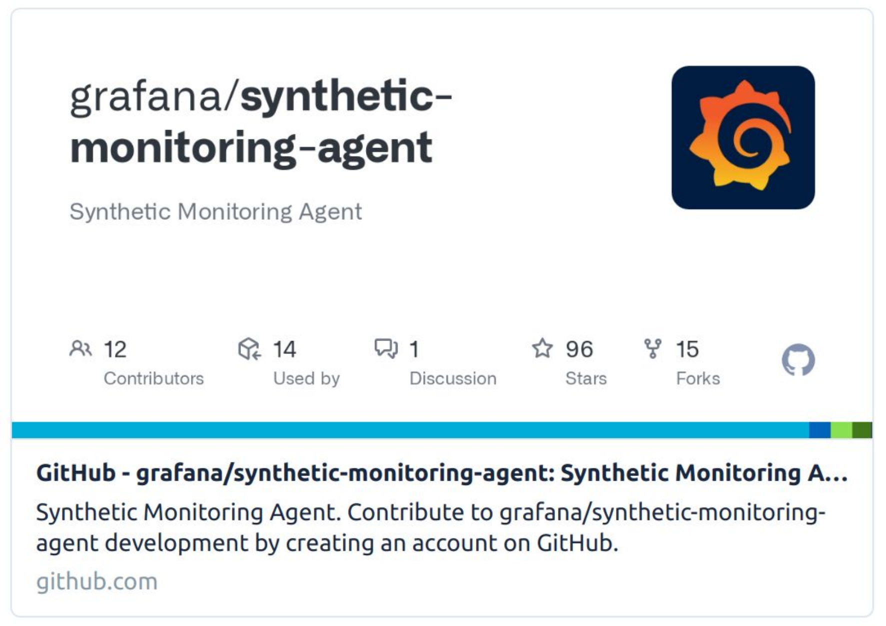

grafana/**synthetic-monitoring-agent**

Synthetic Monitoring Agent

GitHub - grafana/synthetic-monitoring-agent: Synthetic Monitoring A...

Synthetic Monitoring Agent. Contribute to grafana/synthetic-monitoring-agent development by creating an account on GitHub.

github.com

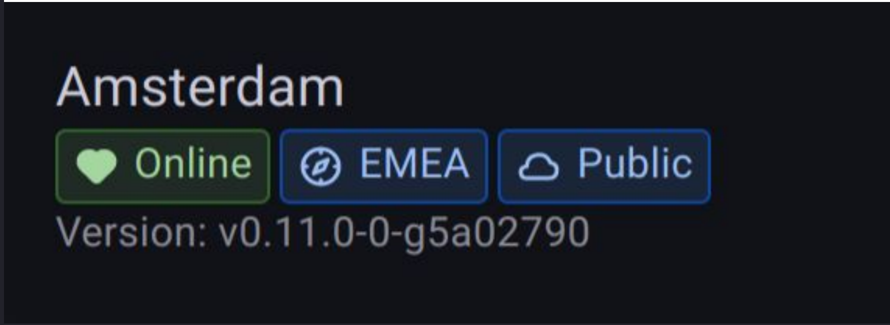

Amsterdam

Version: v0.11.0-0-g5a02790

Grafana Labs

Public Probe

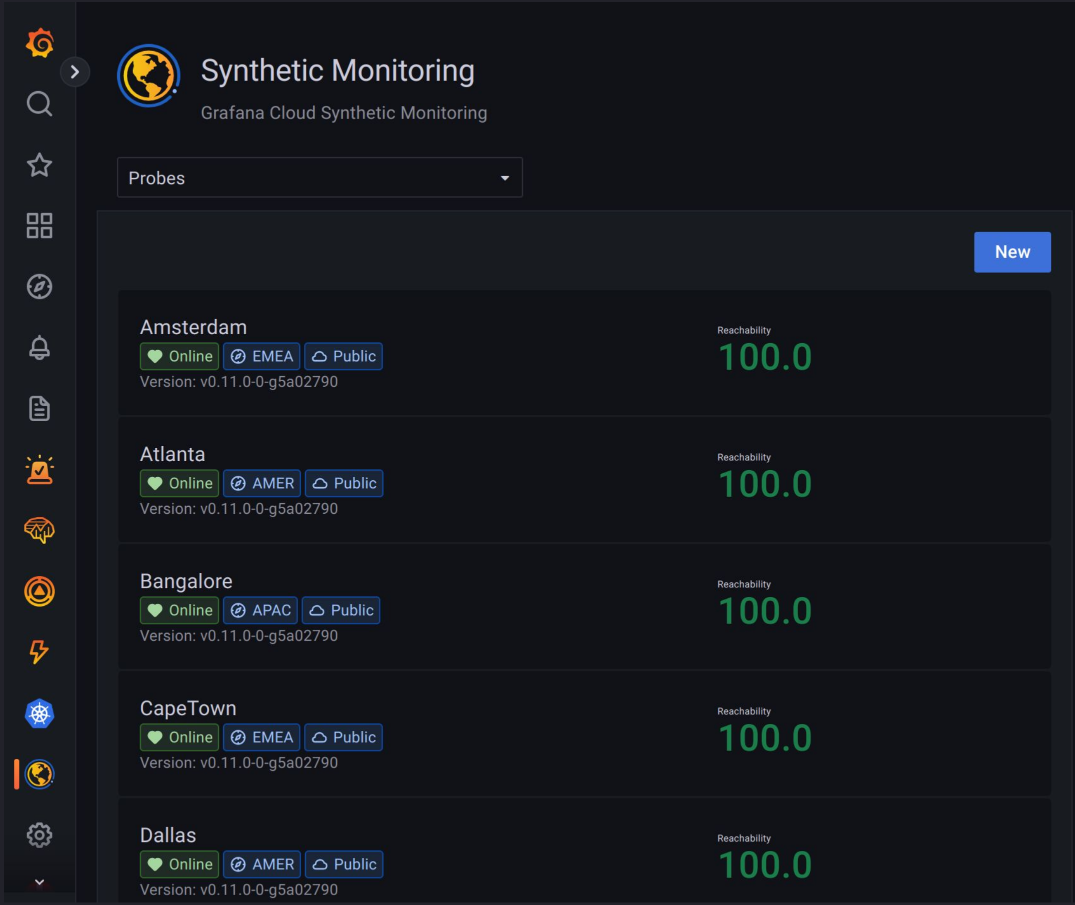

Synthetic Monitoring

Grafana Cloud Synthetic Monitoring

Probes

New

Amsterdam
Online EMEA Public
Version: v0.11.0-0-g5a02790
Reachability 100.0

Atlanta
Online AMER Public
Version: v0.11.0-0-g5a02790
Reachability 100.0

Bangalore
Online APAC Public
Version: v0.11.0-0-g5a02790
Reachability 100.0

CapeTown
Online EMEA Public
Version: v0.11.0-0-g5a02790
Reachability 100.0

Dallas
Online AMER Public
Version: v0.11.0-0-g5a02790
Reachability 100.0

Private Probe

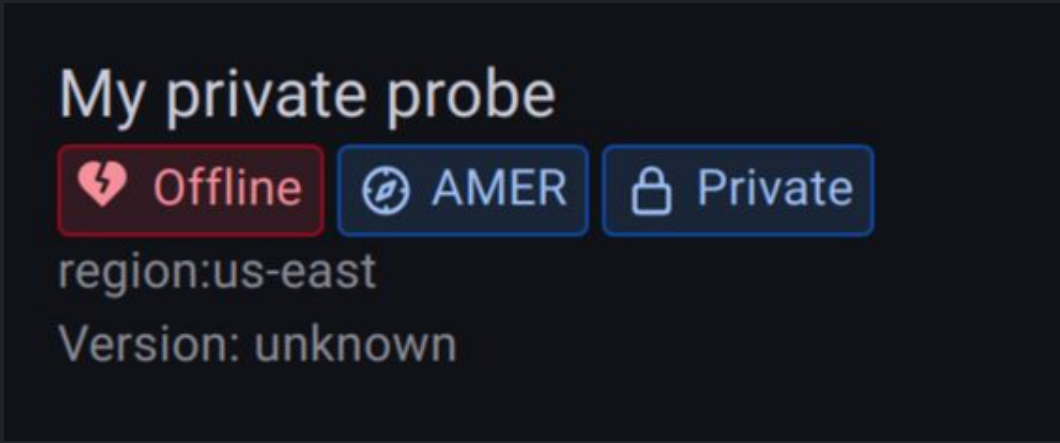

My private probe

region:us-east

Version: unknown

Grafana Labs

Check

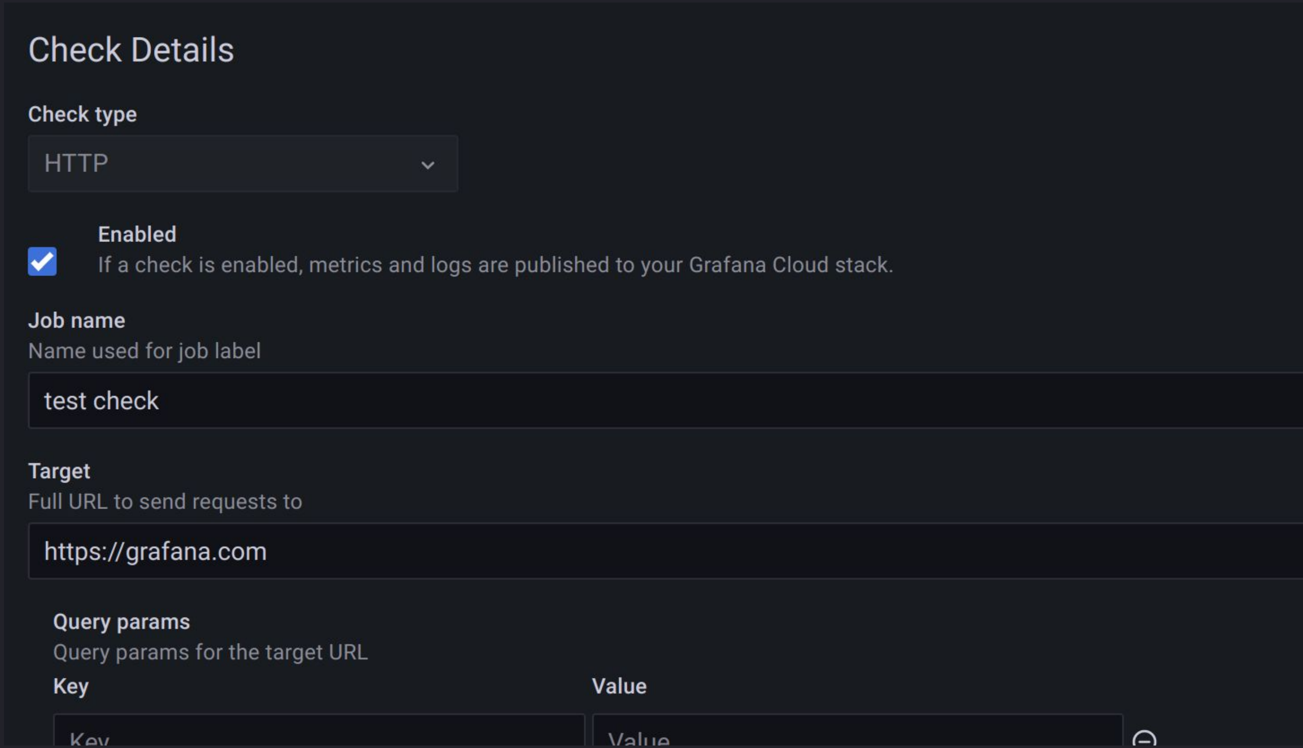

Check Details

Check type

HTTP

☑ **Enabled**
If a check is enabled, metrics and logs are published to your Grafana Cloud stack.

Job name
Name used for job label

test check

Target
Full URL to send requests to

https://grafana.com

Query params
Query params for the target URL

Key	Value
Key	Value

Grafana Labs

Target

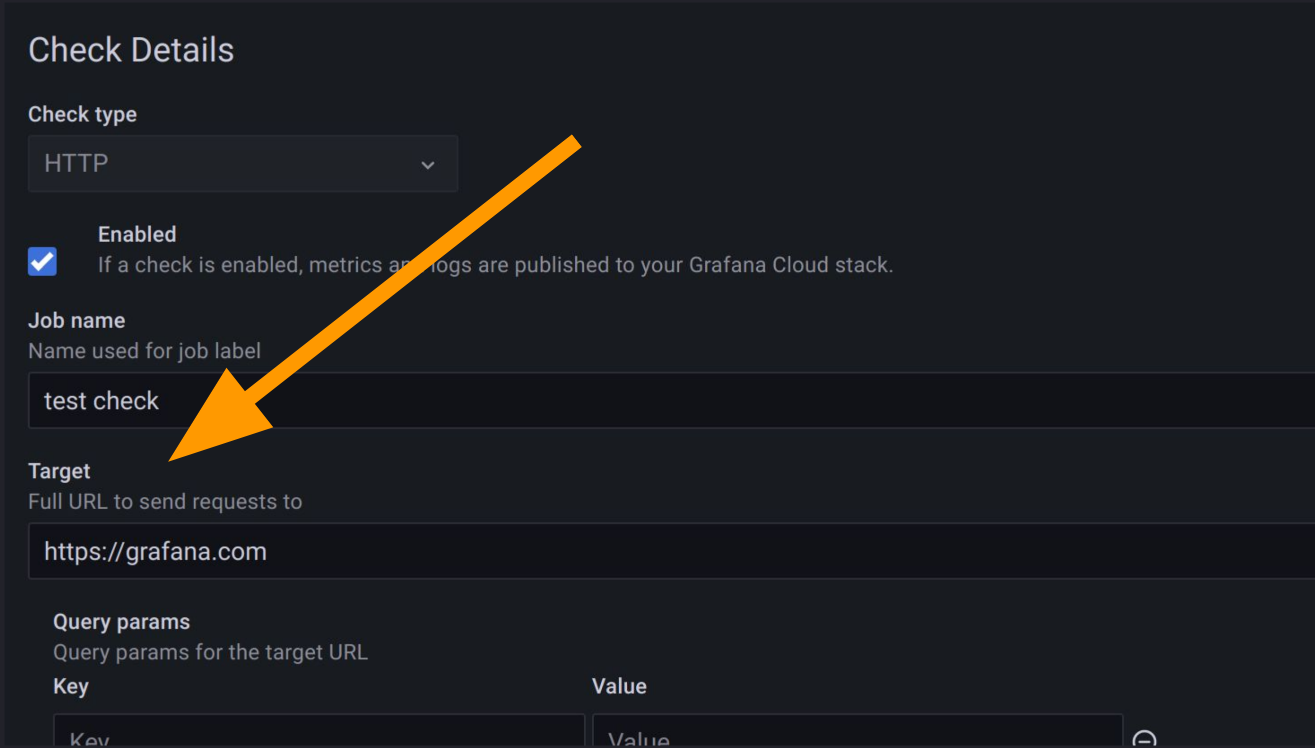

Check Details

Check type

HTTP

☑ Enabled

If a check is enabled, metrics and logs are published to your Grafana Cloud stack.

Job name

Name used for job label

test check

Target

Full URL to send requests to

https://grafana.com

Query params

Query params for the target URL

Key	Value
Key	Value

Grafana Labs

Region

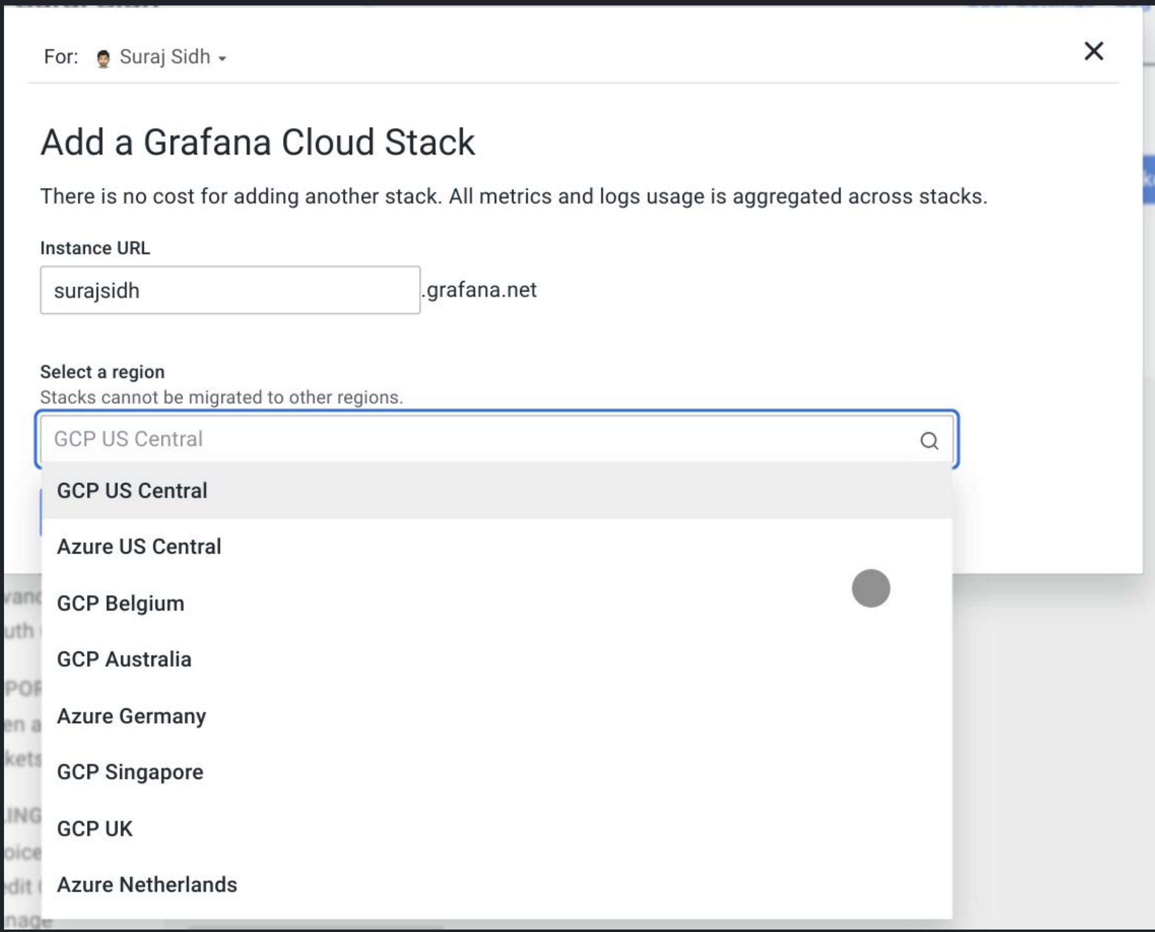

For: Suraj Sidh

Add a Grafana Cloud Stack

There is no cost for adding another stack. All metrics and logs usage is aggregated across stacks.

Instance URL

surajsidh .grafana.net

Select a region

Stacks cannot be migrated to other regions.

GCP US Central

- GCP US Central
- Azure US Central
- GCP Belgium
- GCP Australia
- Azure Germany
- GCP Singapore
- GCP UK
- Azure Netherlands

PoP - Point of Presence

Pop clusters: small kubernetes clusters in locations around the world

Grafana Labs

Grafana Cloud Synthetic Monitoring

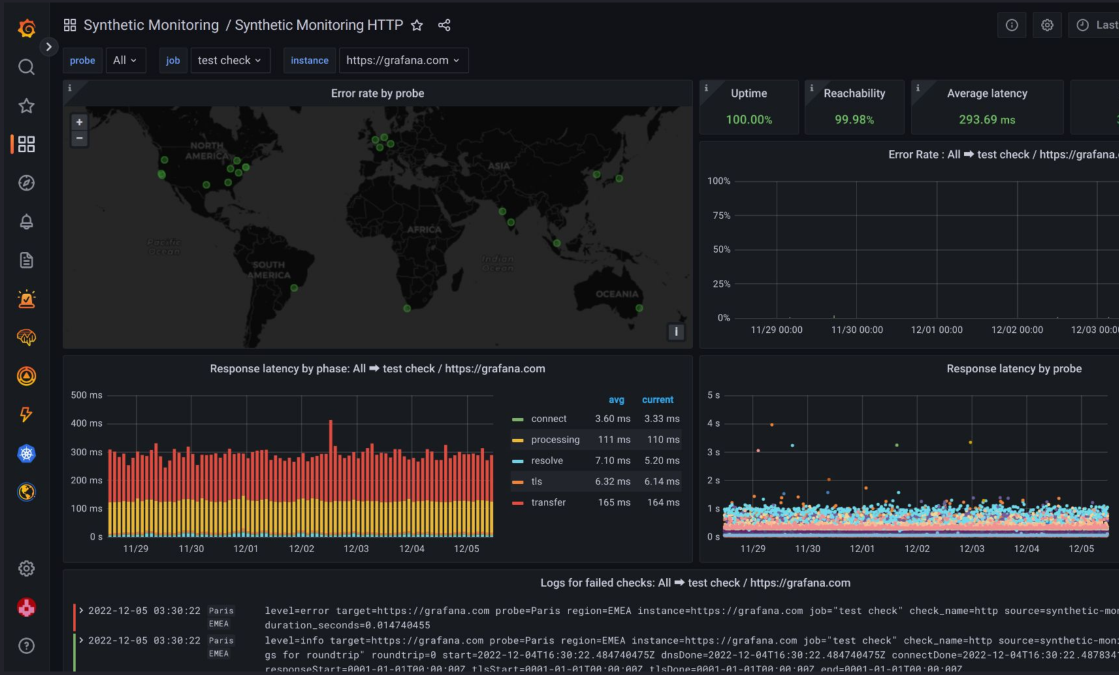

	avg	current
connect	3.60 ms	3.33 ms
processing	111 ms	110 ms
resolve	7.10 ms	5.20 ms
tls	6.32 ms	6.14 ms
transfer	165 ms	164 ms

```
> 2022-12-05 03:30:22 Paris   level=error target=https://grafana.com probe=Paris region=EMEA instance=https://grafana.com job="test check" check_name=http source=synthetic-mo
                      EMEA    duration_seconds=0.014740455
> 2022-12-05 03:30:22 Paris   level=info target=https://grafana.com probe=Paris region=EMEA instance=https://grafana.com job="test check" check_name=http source=synthetic-mon
                      EMEA    gs for roundtrip" roundtrip=0 start=2022-12-04T16:30:22.484740475Z dnsDone=2022-12-04T16:30:22.484740475Z connectDone=2022-12-04T16:30:22.487834
                              responseStart=0001-01-01T00:00:00Z tlsStart=0001-01-01T00:00:00Z tlsDone=0001-01-01T00:00:00Z end=0001-01-01T00:00:00Z
```

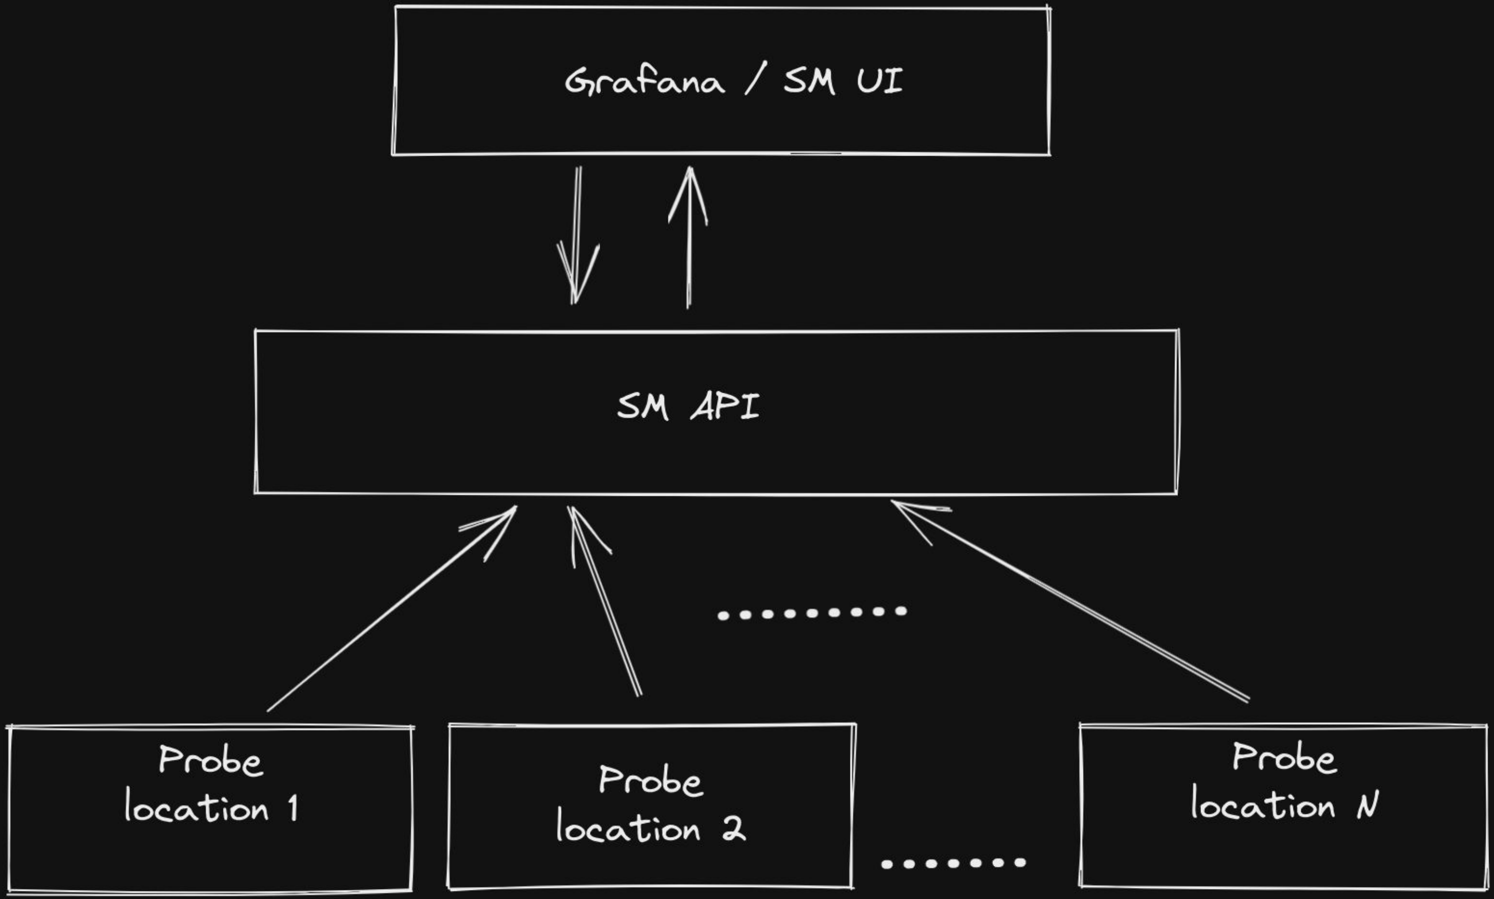

Grafana / SM UI

SM API

Probe location 1

Probe location 2

Probe location N

Grafana Labs

Blackbox Exporter

Probe targets, get prometheus metrics.

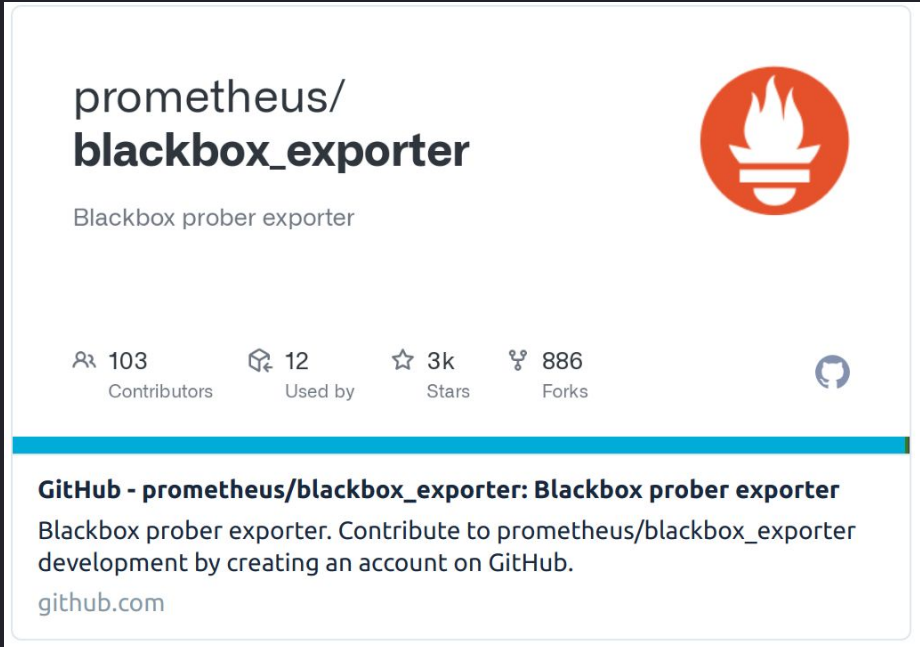

prometheus/
blackbox_exporter

Blackbox prober exporter

GitHub - prometheus/blackbox_exporter: Blackbox prober exporter

Blackbox prober exporter. Contribute to prometheus/blackbox_exporter development by creating an account on GitHub.

github.com

Grafana Labs

Why Blackbox Exporter?

Prometheus metrics, logs, flexible configuration, and many more things to love <3

Grafana Labs

Why do we need a global network?

Why can’t we just deploy it in on location?

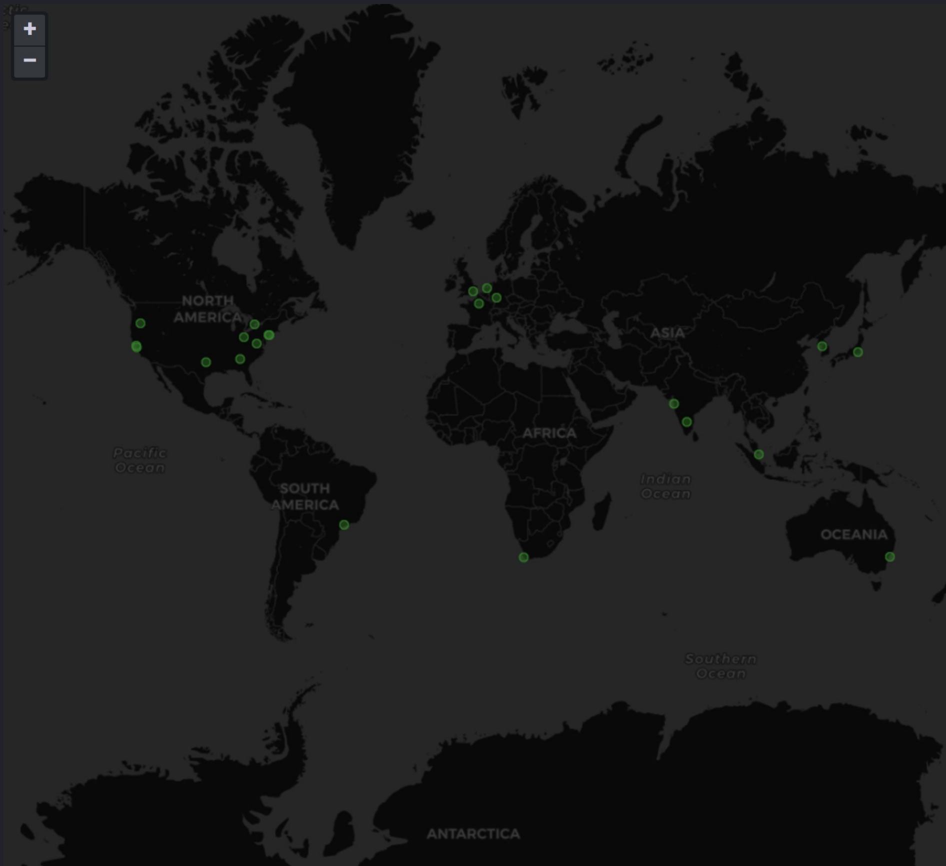

Grafana Labs

Why multiple service providers?

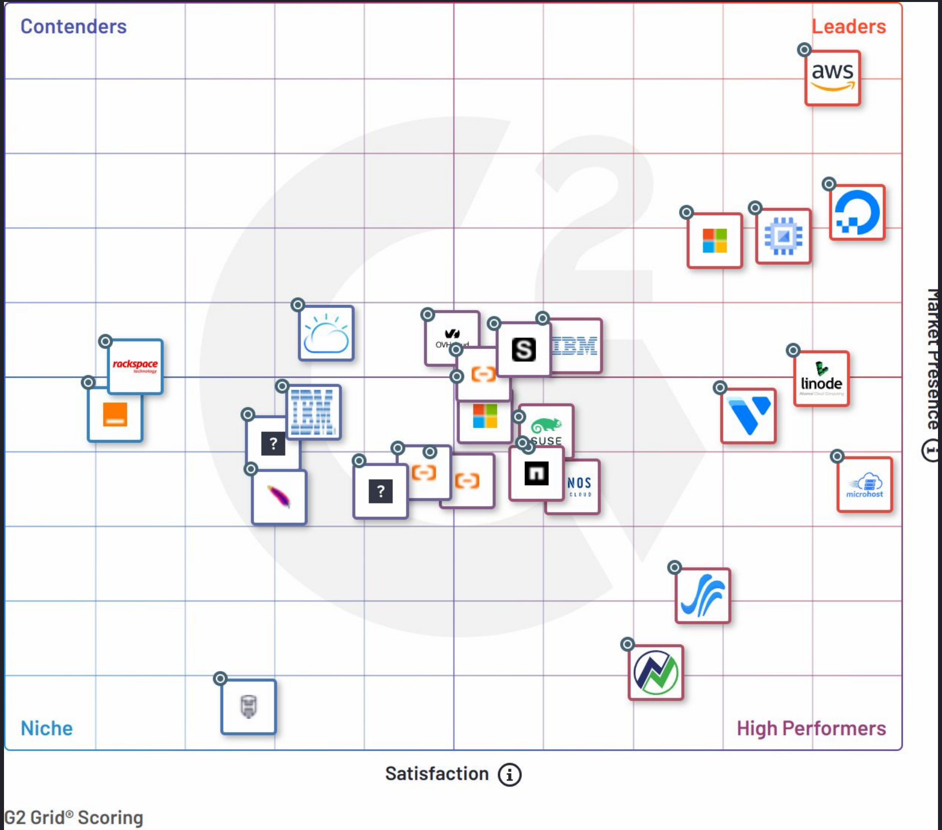

source: https://www.g2.com/categories/infrastructure-as-a-service-iaas

Grafana Labs

Selecting providers is a major challenge

Grafana Labs

IPv6 Support

No support or doesn’t work for some protocols

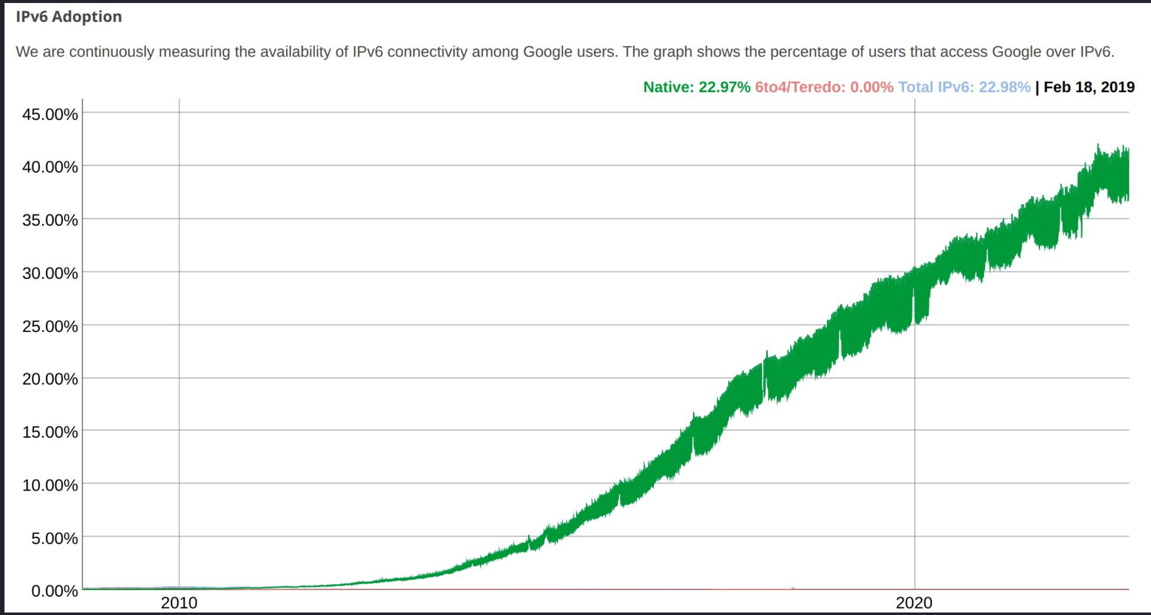

IPv6 Adoption

We are continuously measuring the availability of IPv6 connectivity among Google users. The graph shows the percentage of users that access Google over IPv6.

source: https://www.google.com/intl/en/ipv6/statistics.html#tab=ipv6-adoption

Grafana Labs

Network Stability

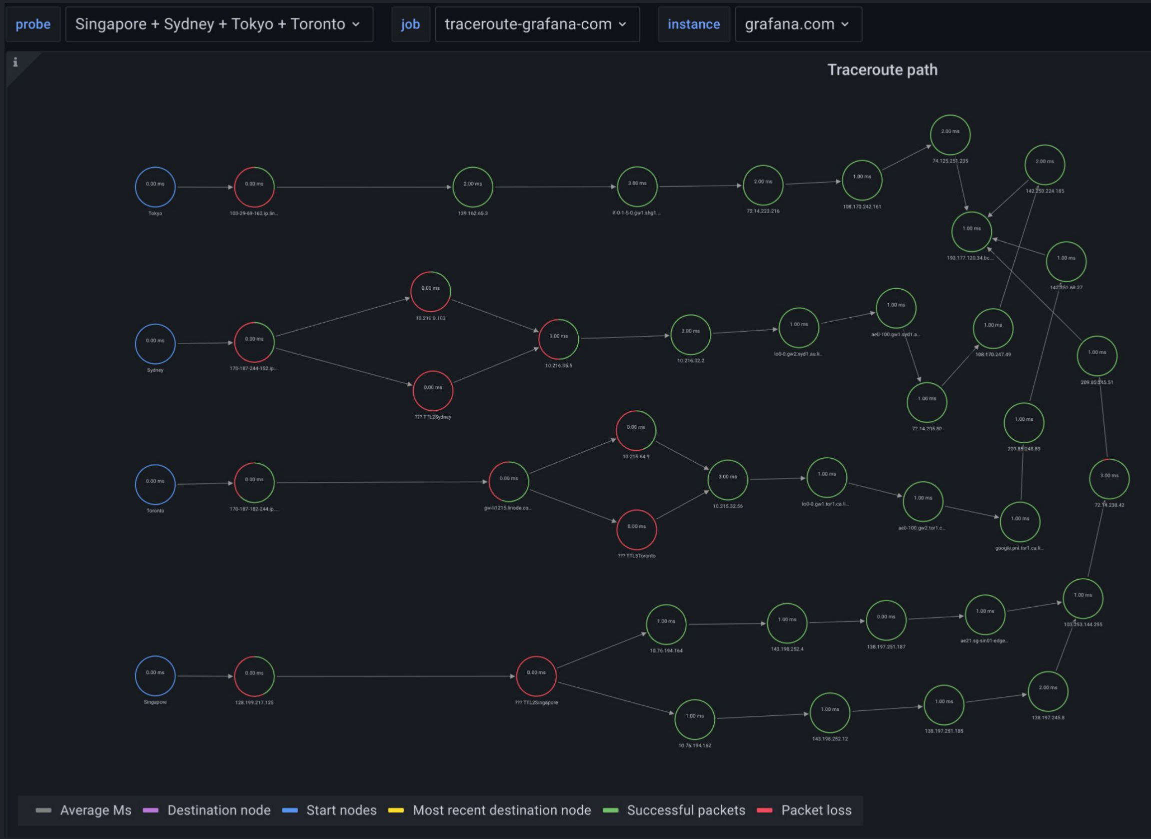

probe Singapore + Sydney + Tokyo + Toronto
job traceroute-grafana-com
instance grafana.com

Grafana Labs

Capacity Issues

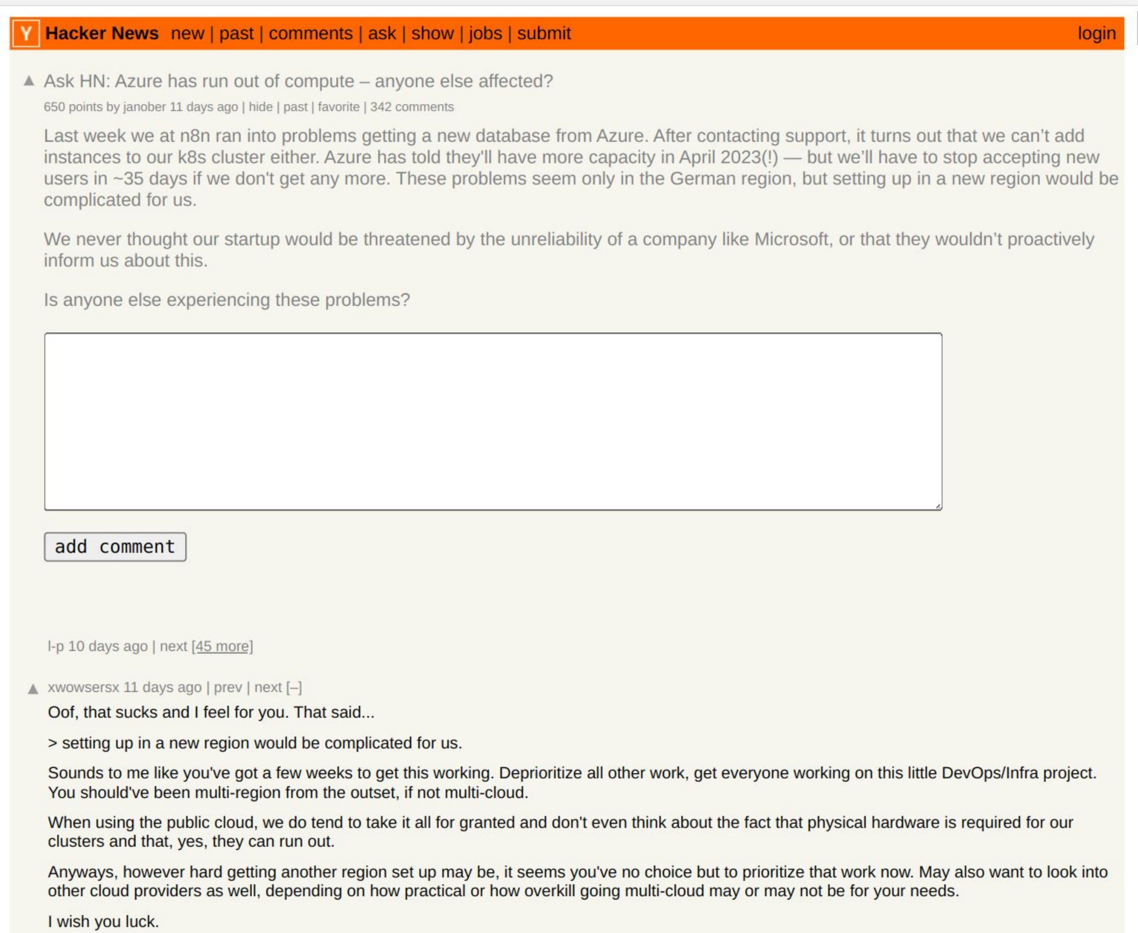

Y Hacker News new | past | comments | ask | show | jobs | submit
login

▲ Ask HN: Azure has run out of compute – anyone else affected?

650 points by janober 11 days ago | hide | past | favorite | 342 comments

Last week we at n8n ran into problems getting a new database from Azure. After contacting support, it turns out that we can't add instances to our k8s cluster either. Azure has told they'll have more capacity in April 2023(!) — but we'll have to stop accepting new users in ~35 days if we don't get any more. These problems seem only in the German region, but setting up in a new region would be complicated for us.

We never thought our startup would be threatened by the unreliability of a company like Microsoft, or that they wouldn't proactively inform us about this.

Is anyone else experiencing these problems?

l-p 10 days ago | next [45 more]

▲ xwowsersx 11 days ago | prev | next [–]

Oof, that sucks and I feel for you. That said...

> setting up in a new region would be complicated for us.

Sounds to me like you've got a few weeks to get this working. Deprioritize all other work, get everyone working on this little DevOps/Infra project. You should've been multi-region from the outset, if not multi-cloud.

When using the public cloud, we do tend to take it all for granted and don't even think about the fact that physical hardware is required for our clusters and that, yes, they can run out.

Anyways, however hard getting another region set up may be, it seems you've no choice but to prioritize that work now. May also want to look into other cloud providers as well, depending on how practical or how overkill going multi-cloud may or may not be for your needs.

I wish you luck.

Tracking Health and Reliability of probes

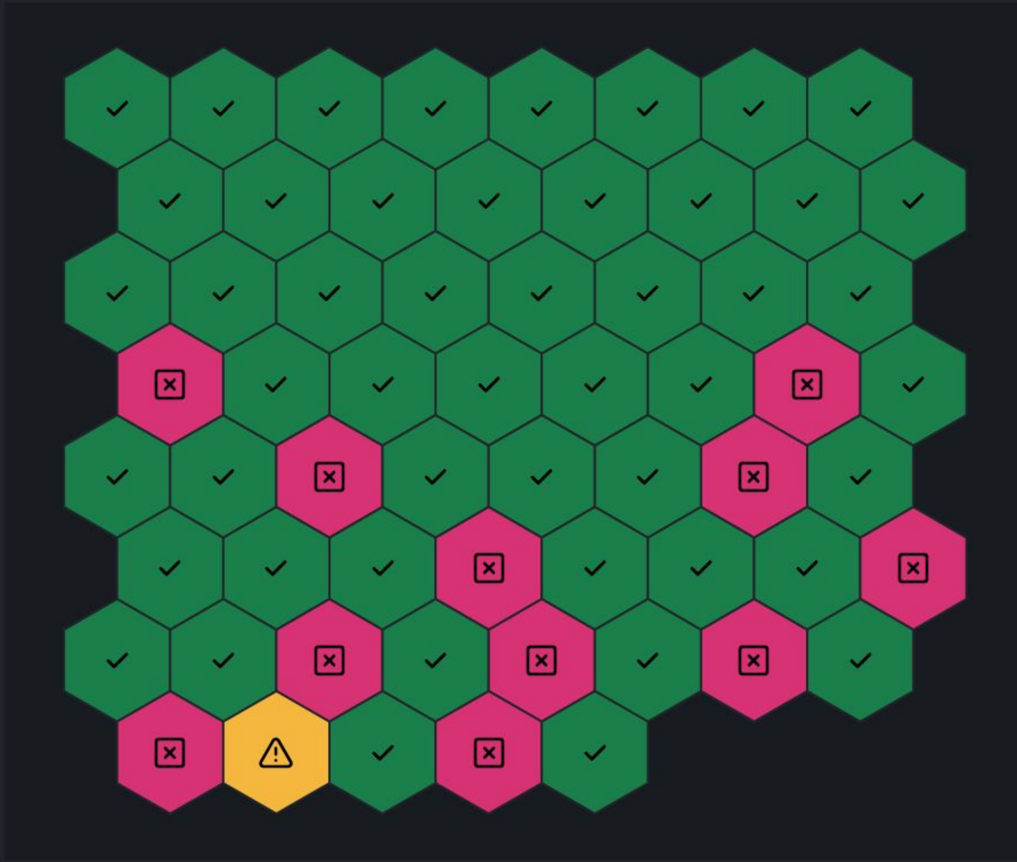

Grafana Labs

Is Probe Healthy & Reliable?

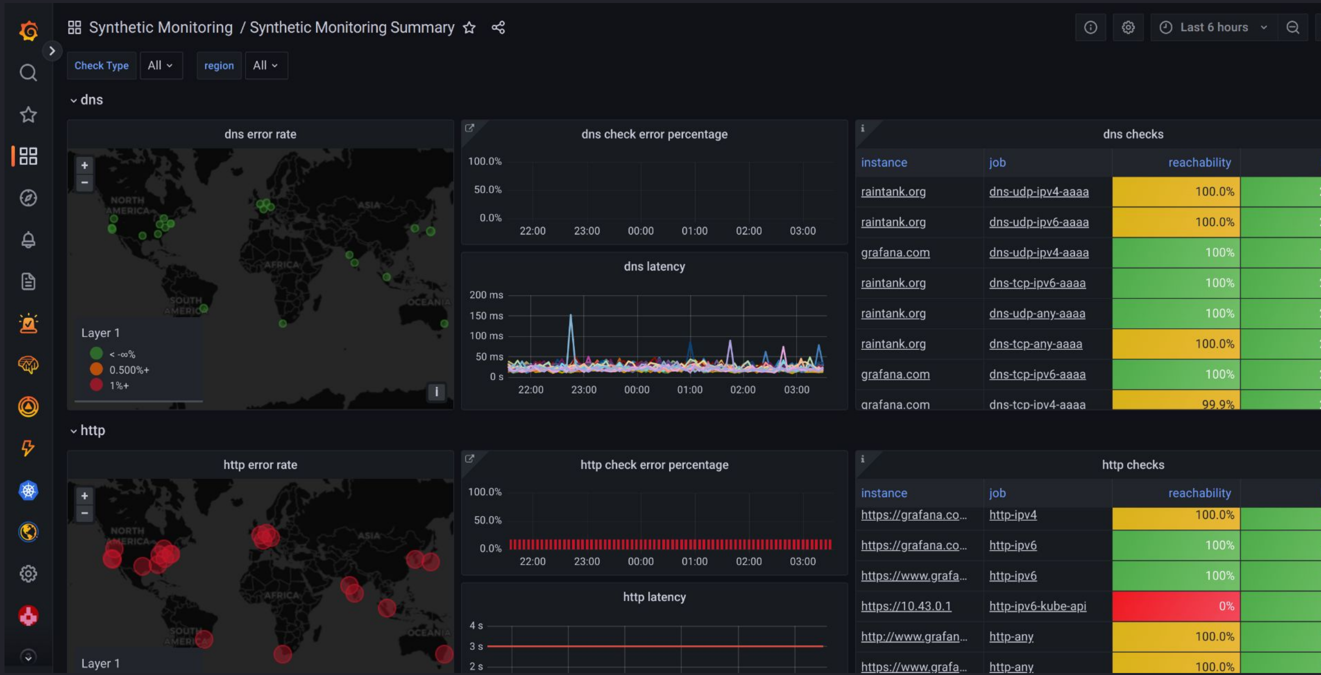
Synthetic Monitoring / Synthetic Monitoring Summary

Check Type: All · region: All · Last 6 hours

dns

dns checks

instance	job	reachability
raintank.org	dns-udp-ipv4-aaaa	100.0%
raintank.org	dns-udp-ipv6-aaaa	100.0%
grafana.com	dns-udp-ipv4-aaaa	100%
raintank.org	dns-tcp-ipv6-aaaa	100%
raintank.org	dns-udp-any-aaaa	100%
raintank.org	dns-tcp-any-aaaa	100.0%
grafana.com	dns-tcp-ipv6-aaaa	100%
grafana.com	dns-tcp-ipv4-aaaa	99.9%

http

http checks

instance	job	reachability
https://grafana.co...	http-ipv4	100.0%
https://grafana.co...	http-ipv6	100%
https://www.grafa...	http-ipv6	100%
https://10.43.0.1	http-ipv6-kube-api	0%
http://www.grafan...	http-any	100.0%
https://www.grafa...	http-any	100.0%

Grafana Labs

Evolution of Probe Deployments

Grafana Labs

VMs: simple and easy

but not ideal

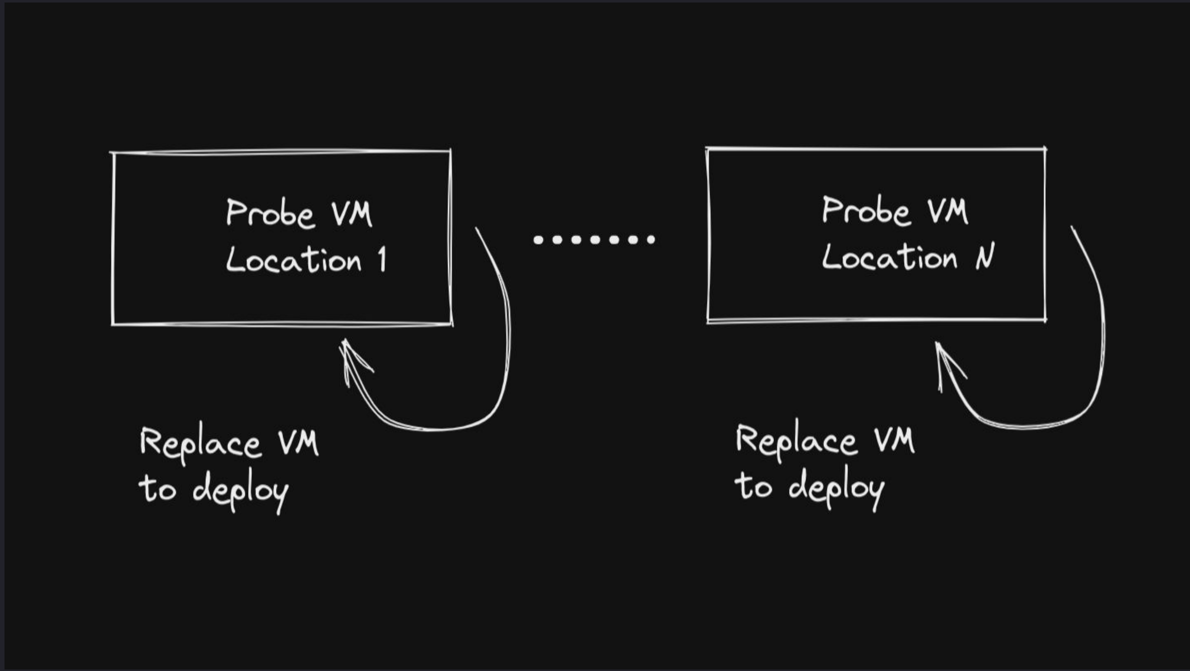

Locations * Regions = ERR_TOO_MANY_VMS

22 * 9 = 198 VMs, for current scale

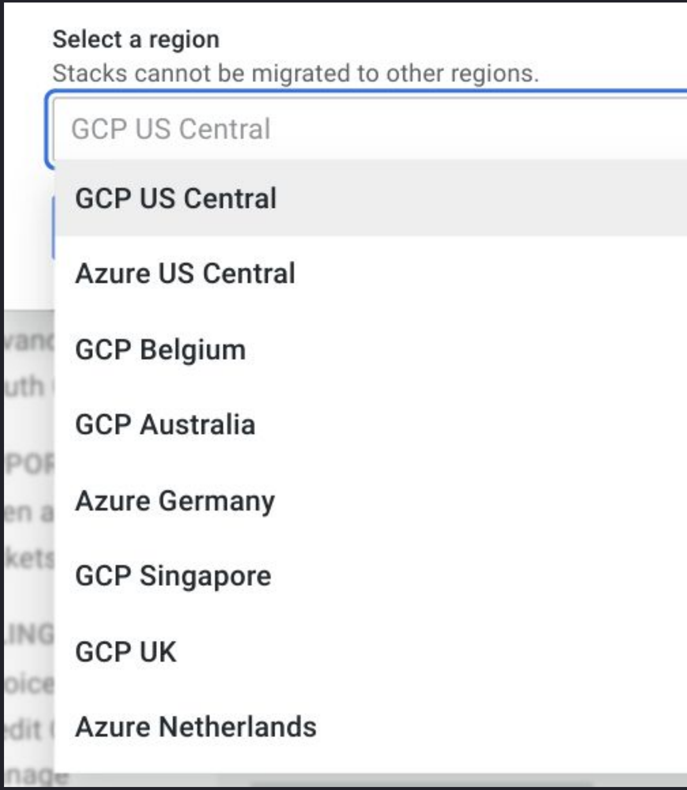

Select a region

Stacks cannot be migrated to other regions.

GCP US Central

- GCP US Central
- Azure US Central
- GCP Belgium
- GCP Australia
- Azure Germany
- GCP Singapore
- GCP UK
- Azure Netherlands

We started looking for a solution

and made a list

Things we need

- Tolerate Node failures
- IPv6 Support
- Security

Nice to have

- Declarative Deployment Model
- Rolling Updates
- Good Tooling & Developer Experience

Things we don’t need

- Persistence / Disks

Grafana Labs

We looked At:

Nomad, Ansible, Google Anthos, Provider Managed Kubernetes, Self Managed Kubernetes

Self Managed Kubernetes

- Fits our needs and wants
- K8s is widely used across company
- In-house expertise and tooling

Grafana Labs

Why Self Managed Kubernetes

- Can configure as per our needs
- Stateless workload makes management easy
- Uniform setup across providers (just bring VMs)

Grafana Labs

Why Self Managed Kubernetes

- Native IPv6 support with Dual Stack
- Network Policies for Network Security
- Gradual rollouts with Existing tooling

Cluster Setup

One Cluster per location.

Total 22 clusters

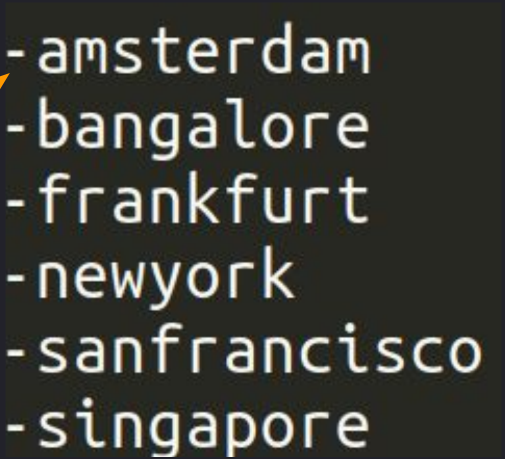

```
-amsterdam
-bangalore
-frankfurt
-newyork
-sanfrancisco
-singapore
```

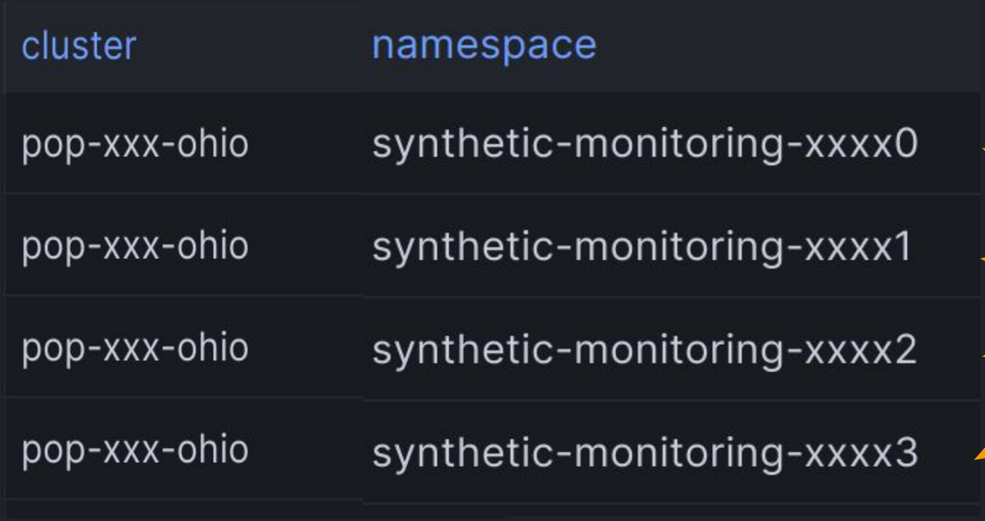

cluster	namespace
pop-xxx-ohio	synthetic-monitoring-xxxx0
pop-xxx-ohio	synthetic-monitoring-xxxx1
pop-xxx-ohio	synthetic-monitoring-xxxx2
pop-xxx-ohio	synthetic-monitoring-xxxx3

One namespace per Grafana Cloud Region in each cluster

Grafana Labs

Cluster Setup

- One PoP cluster per location
- Used RKE with Calico CNI and Terraform for provisioning
- One namespace for each region in every cluster

Grafana Labs

Security concerns

- Exposing Internal network and cluster components
- NET_RAW capability in Agent Pods

Grafana Labs

Secure Pop Clusters

- Calico Network Policy to block private network
- Public DNS in SM Agent pods

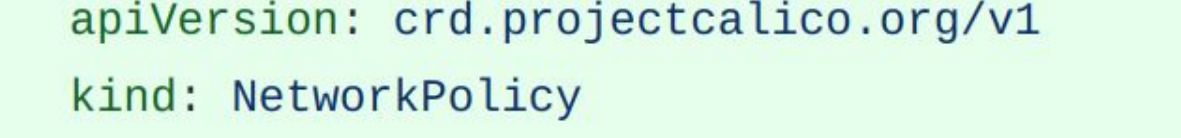

```yaml
apiVersion: crd.projectcalico.org/v1
kind: NetworkPolicy
```

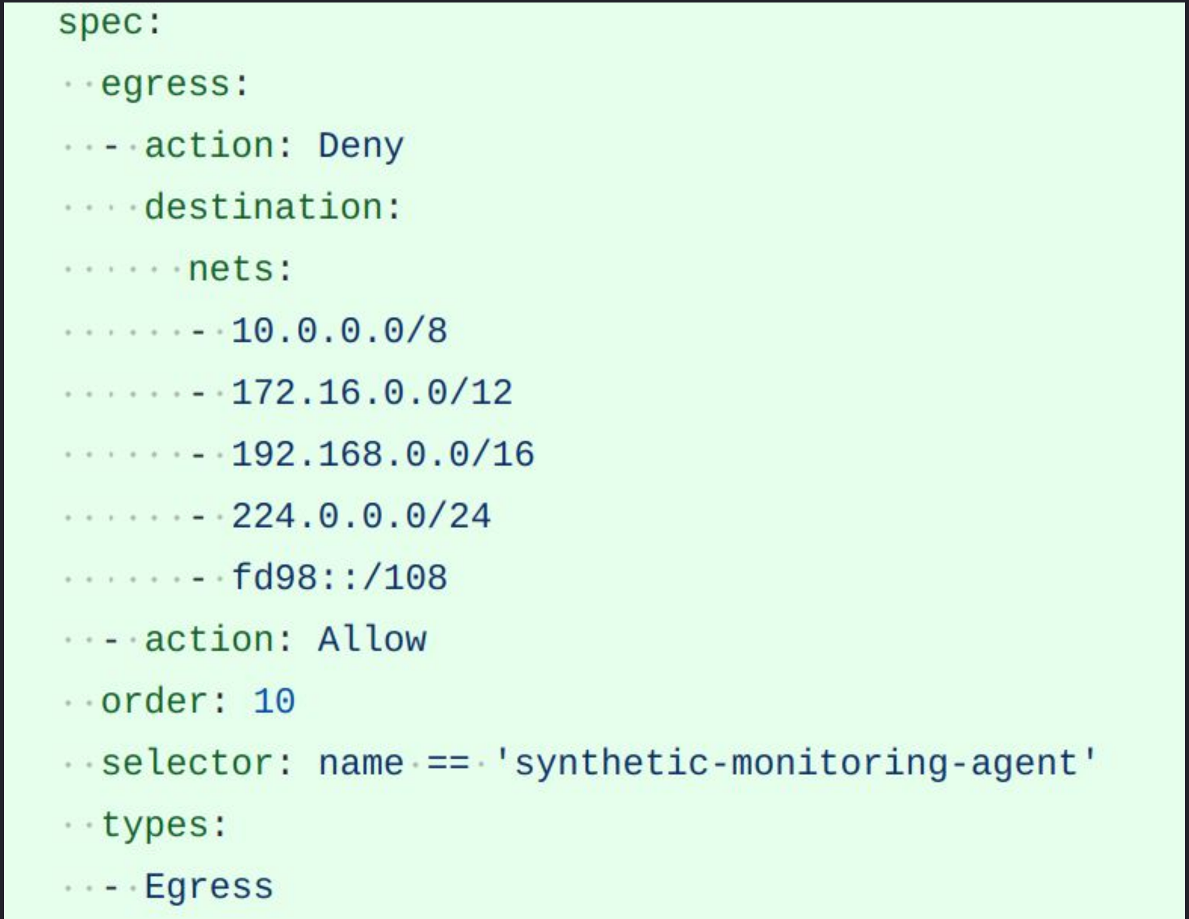

```yaml
spec:
  egress:
  - action: Deny
    destination:
      nets:
      - 10.0.0.0/8
      - 172.16.0.0/12
      - 192.168.0.0/16
      - 224.0.0.0/24
      - fd98::/108
  - action: Allow
  order: 10
  selector: name == 'synthetic-monitoring-agent'
  types:
  - Egress
```

Grafana Labs

Agent release rollouts

GitHub → Drone → Argo Workflows → Flux

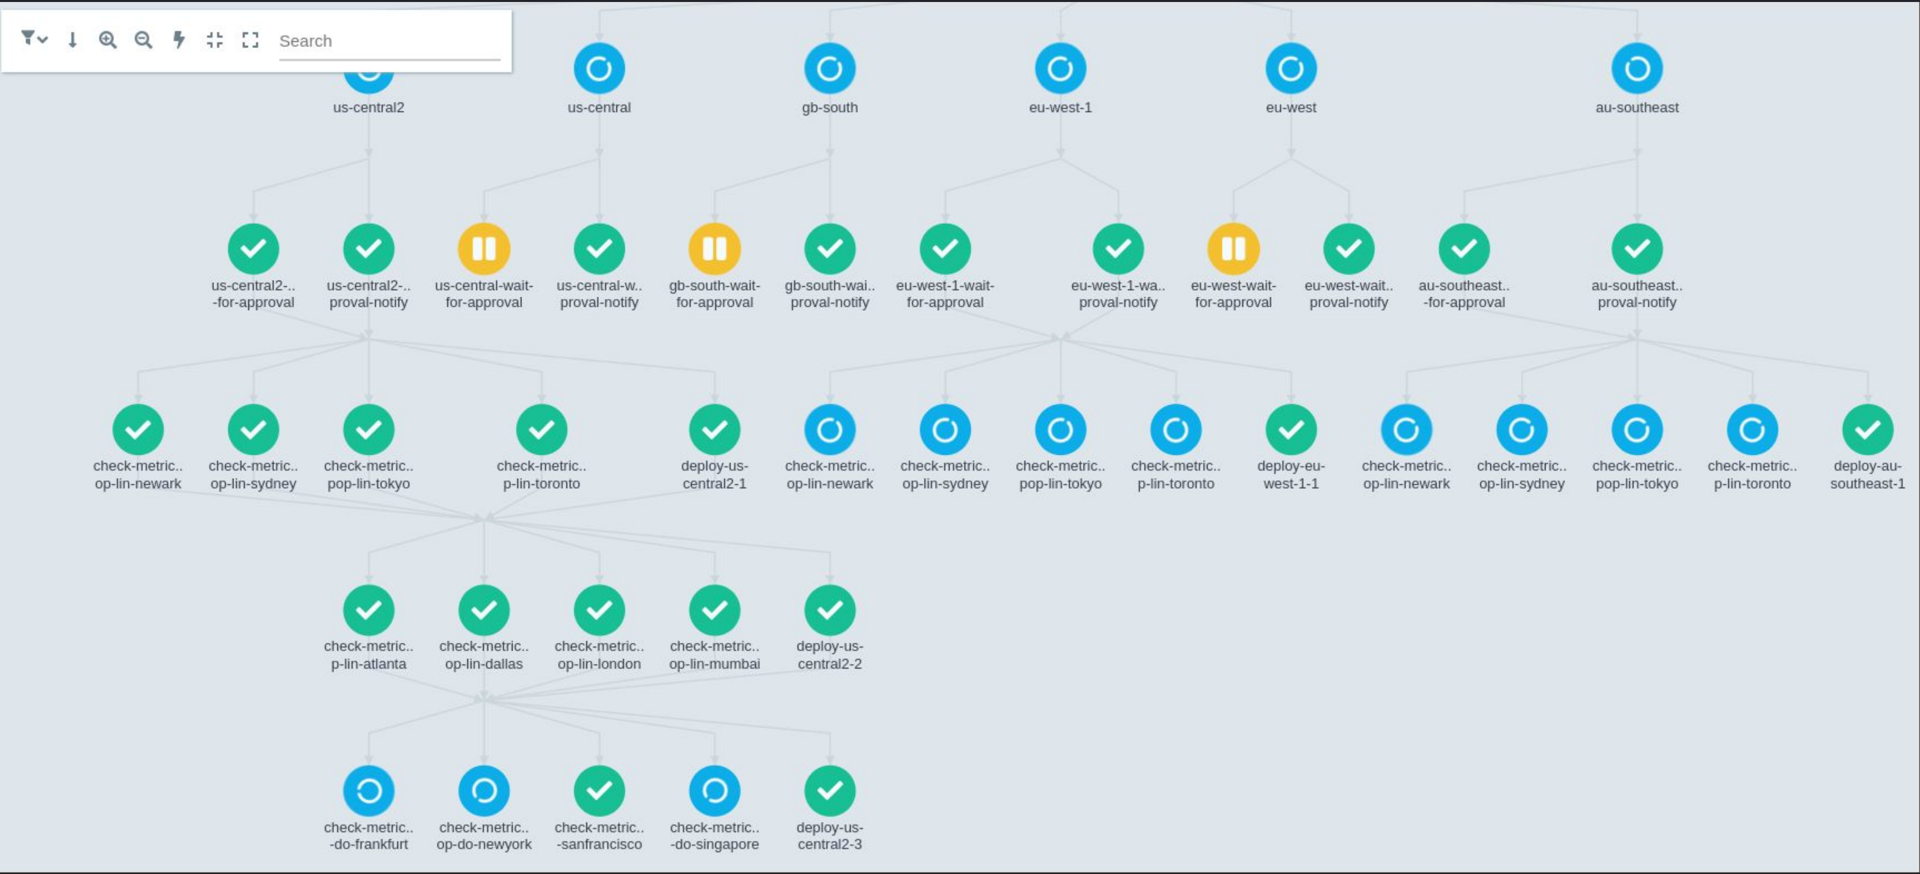
Search
us-central2
us-central
gb-south
eu-west-1
eu-west
au-southeast
us-central2-..
-for-approval
us-central2-..
proval-notify
us-central-wait-
for-approval
us-central-w..
proval-notify
gb-south-wait-
for-approval
gb-south-wai..
proval-notify
eu-west-1-wait-
for-approval
eu-west-1-wa..
proval-notify
eu-west-wait-
for-approval
eu-west-wait..
proval-notify
au-southeast..
-for-approval
au-southeast..
proval-notify
check-metric..
op-lin-newark
check-metric..
op-lin-sydney
check-metric..
pop-lin-tokyo
check-metric..
p-lin-toronto
deploy-us-
central2-1
check-metric..
op-lin-newark
check-metric..
op-lin-sydney
check-metric..
pop-lin-tokyo
check-metric..
p-lin-toronto
deploy-eu-
west-1-1
check-metric..
op-lin-newark
check-metric..
op-lin-sydney
check-metric..
pop-lin-tokyo
check-metric..
p-lin-toronto
deploy-au-
southeast-1
check-metric..
p-lin-atlanta
check-metric..
op-lin-dallas
check-metric..
op-lin-london
check-metric..
op-lin-mumbai
deploy-us-
central2-2
check-metric..
-do-frankfurt
check-metric..
op-do-newyork
check-metric..
-sanfrancisco
check-metric..
-do-singapore
deploy-us-
central2-3

Gradual rollout of probes across regions & locations

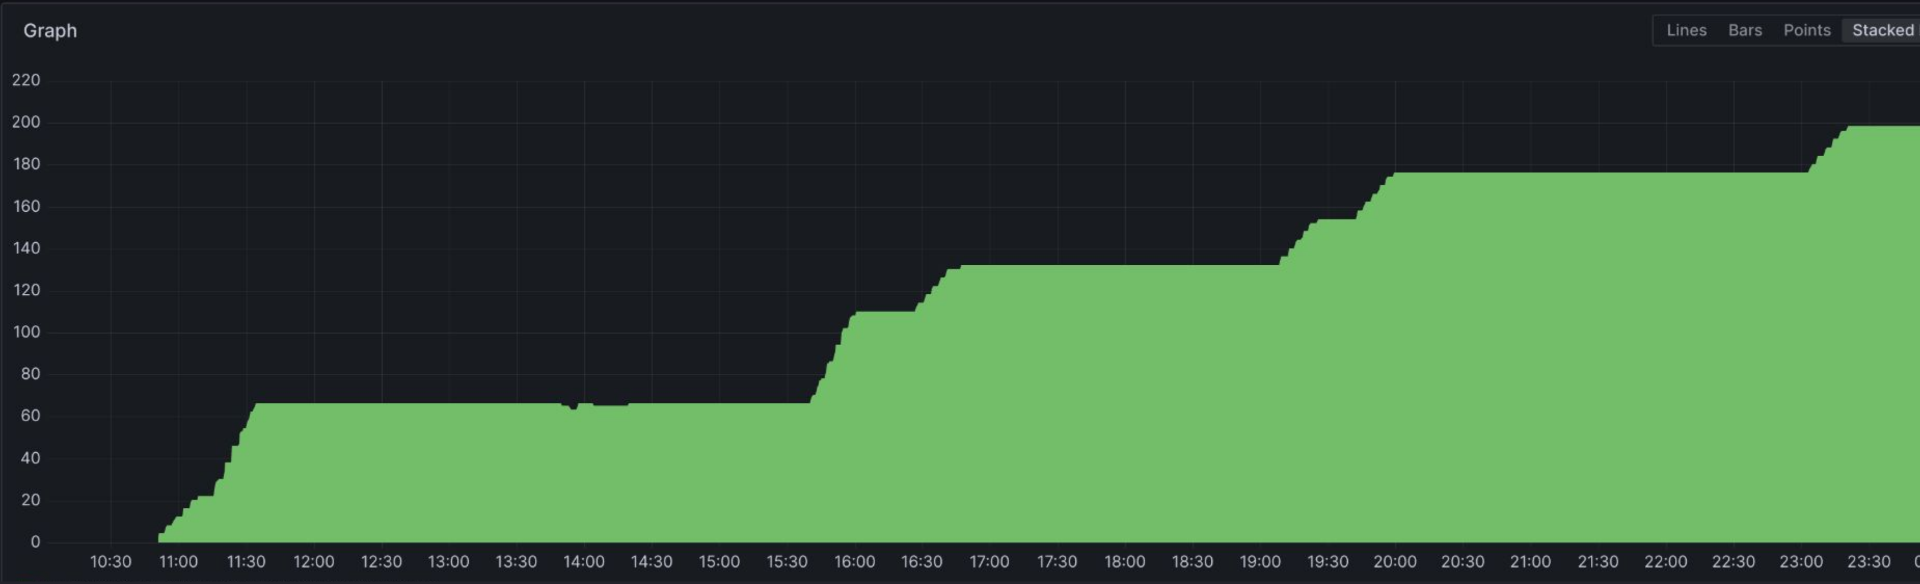

Grafana Labs

Lessons Learned

Parts of Internet will always be unreliable

Internet is a big distributed system, and
is built to route around unreliable parts.

Add Redundancy where you can

External single point of failures will burn your error budget.

Grafana Labs

Maintenance Windows

Communicate service provider’s maintenance up to your users.

Grafana Labs

Good Defaults & User Education

Nudge users to make good decisions

Grafana Labs

Limits and Abuse

Limits on resources a single users can provision, i.e. number of checks

Bad actors will try to use you for recon work.

Be Verbose

Otherwise it will be assumed that you are an bad actor, and end up in blocklist

Implicit Dependencies

Service providers are not happy when overloaded their DNS

Grafana Labs

Thank You

Slides: https://suraj.dev/talks

Twitter: @electron0zero