Consortium GARR
THE ITALIAN EDUCATION & RESEARCH NETWORK

GARR Cloud monitoring

Claudio Pisa - GARR

EAPConnect Workshop II

2019-11-22 Rome

Monitoring - day 0 (one year ago)

- Different monitoring systems:
  - Nagios
    - OpenStack services
  - Zabbix
    - Ceph
    - Hardware sensors
- Some systems not monitored

Architecture - day 0

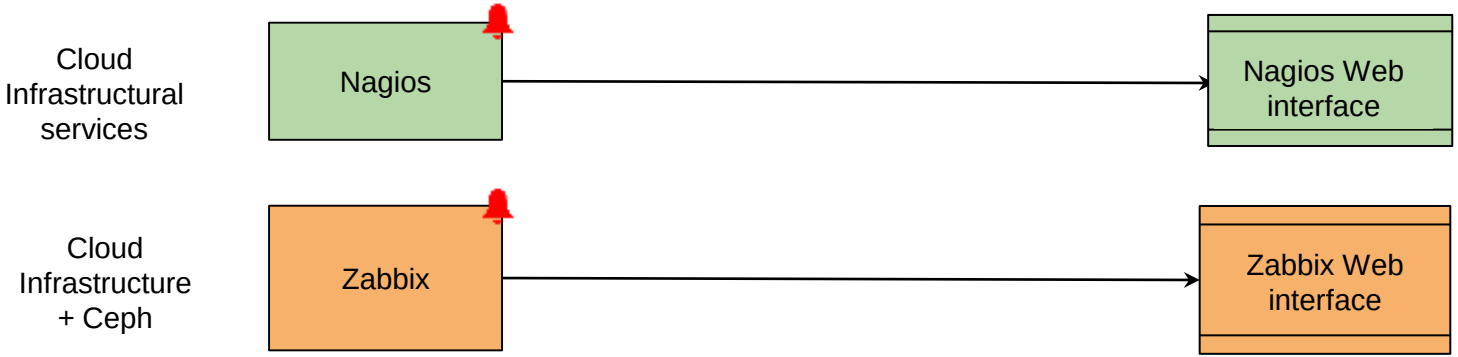

Architecture - day 0

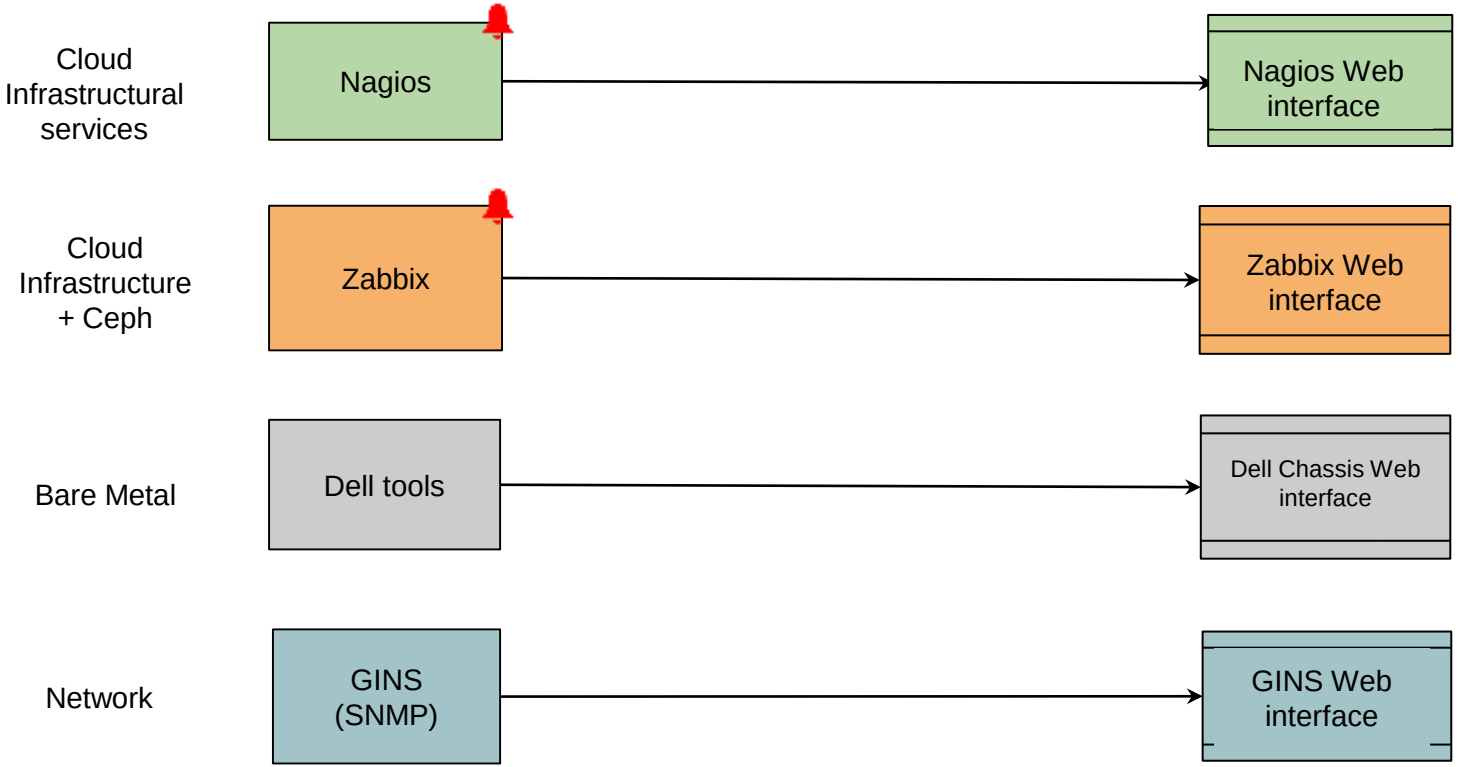

Objective

Objective:
build a unified comprehensive dashboard

State of the art:
Grafana

Considerations

- Existing monitoring systems seem to be doing their job well
  - the right tool for the right job
- What is missing is just a single viewpoint
- Grafana seems to have what we need

Grafana

- **Grafana** is a platform for data visualization, querying and alerting
- Several pluggable data sources:
  - Zabbix
  - PNP (Nagios)
  - Prometheus
  - Gnocchi
  - Monasca
  - JSON (general purpose)
  - MySQL / PostgreSQL (general purpose)
- Data from heterogeneous sources can be mixed in the same dashboard

Proposed Architecture

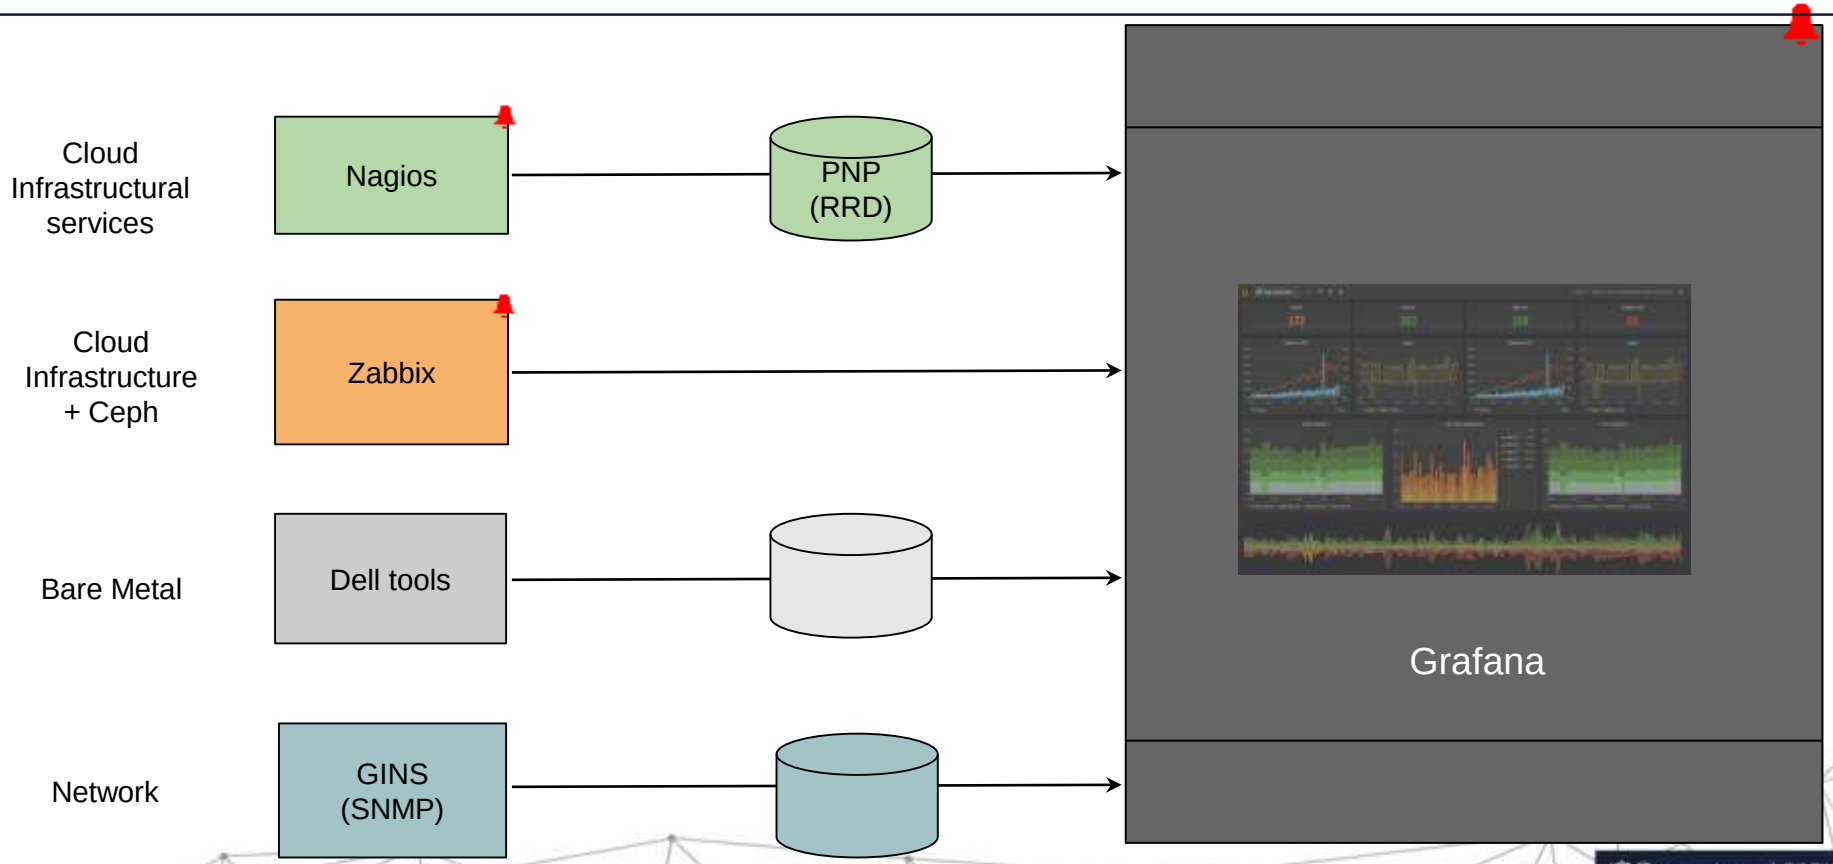

Zabbix → Grafana

ZABBIX

- Zabbix
  - “batteries included” free and open source monitoring tool
  - auto discovery
  - XML based templates
  - remote agents
  - Good Ceph integration
  - Good multiuser support
- Straightforward integration with Grafana
  - Grafana’s Zabbix datasource

Ceph

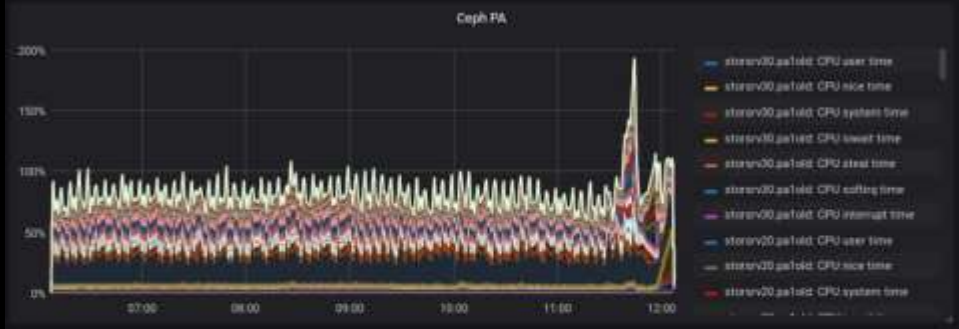

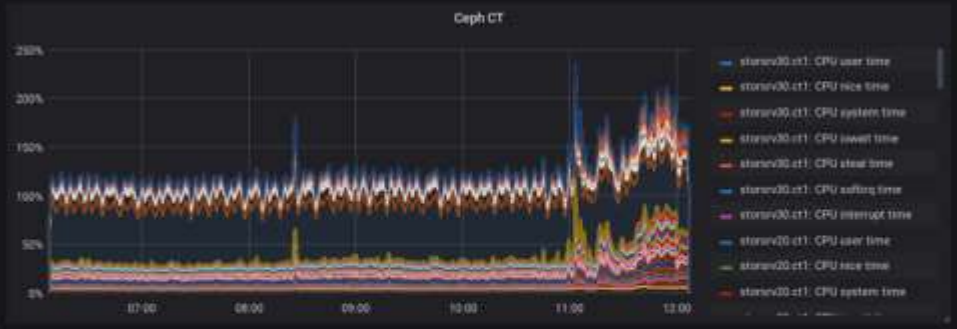

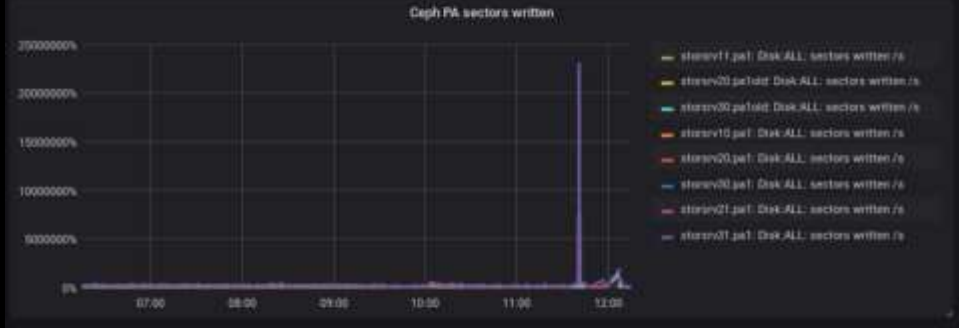

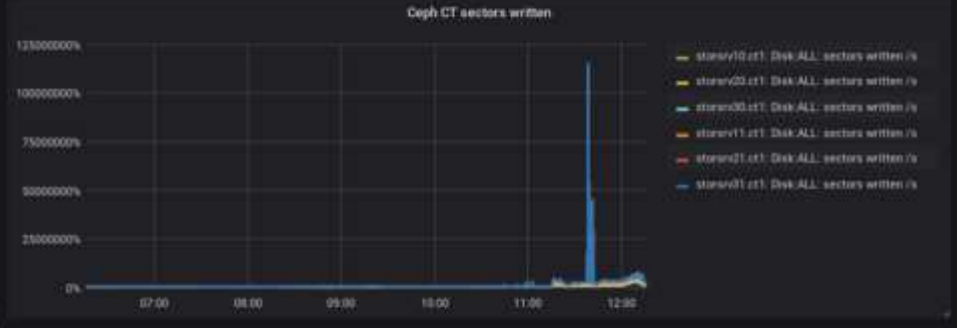

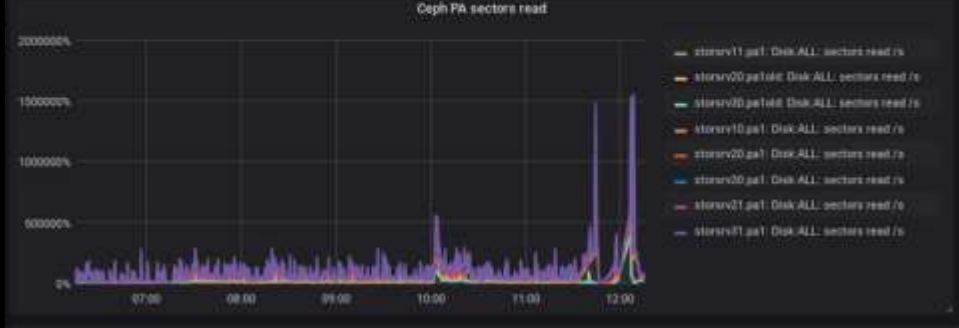

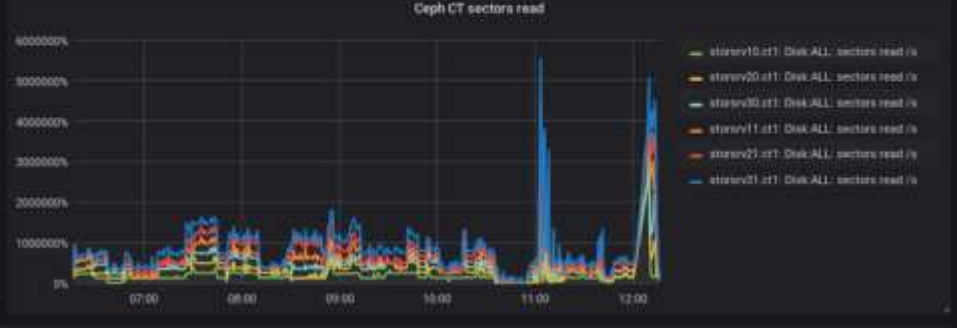

Zabbix Server (7 panels)

Nagios

- **Nagios** is a free and open source monitoring tool
  - plugins - scripts to check and report
  - Nagios remote plugin executor (NRPE) - for remote hosts
  - known especially for alerting
  - FAQ: how do you pronounce Nagios?
    - the author pronounces it as “nah-ghee-ose”
    - but “you can pronounce it however the heck you'd like”
    - even “nachos”

Nagios

- Nagios is very well integrated in Juju
  - Nagios charm
  - NRPE charm
  - Canonical OpenStack charms come with handy Nagios configuration options

PNP my Nagios

- PNP is an addon to Nagios which analyzes performance data provided by Nagios plugins and stores them automatically into Round Robin Databases (RRD)
  - and there is a PNP Grafana datasource

OpenStack → Nagios

- The Nagios Charm monitors the OpenStack services
  - but not what is happening inside OpenStack
- Simple idea: **write a set of Nagios plugins which collect metrics from the OpenStack API**
  - number of projects
  - number of servers
  - floating IP address usage
  - volume usage
  - OpenStack APIs reachability
  - ...

CT1 →

Allocated VCPUs CT1

Allocated RAM CT1

Taken Public IPs CT1

Active instances CT1

551

Volume Usage CT1

340 TB

Active compute hosts CT1

24

PA1 →

Allocated VCPUs PA1

Allocated RAM PA1

Taken Public IPs PA1

Active instances PA1

552

Volume Usage PA1

183 TB

Active compute hosts PA1

16

NA2 →

Allocated VCPUs NA2

Allocated RAM NA2

Taken Public IPs NA2

Active NA2 instances

34

Volume Usage NA2

1.925 TB

Active compute hosts NA2

3

Active users

614

Active Projects

935

Demo projects

602

vDCs

47

Active images

474

Floating IPs

Active Instances

Active volumes

Total RAM Usage (GiB)

Total volume usage (GiB)

Flavor Distribution CT1

Flavor Distribution PA1

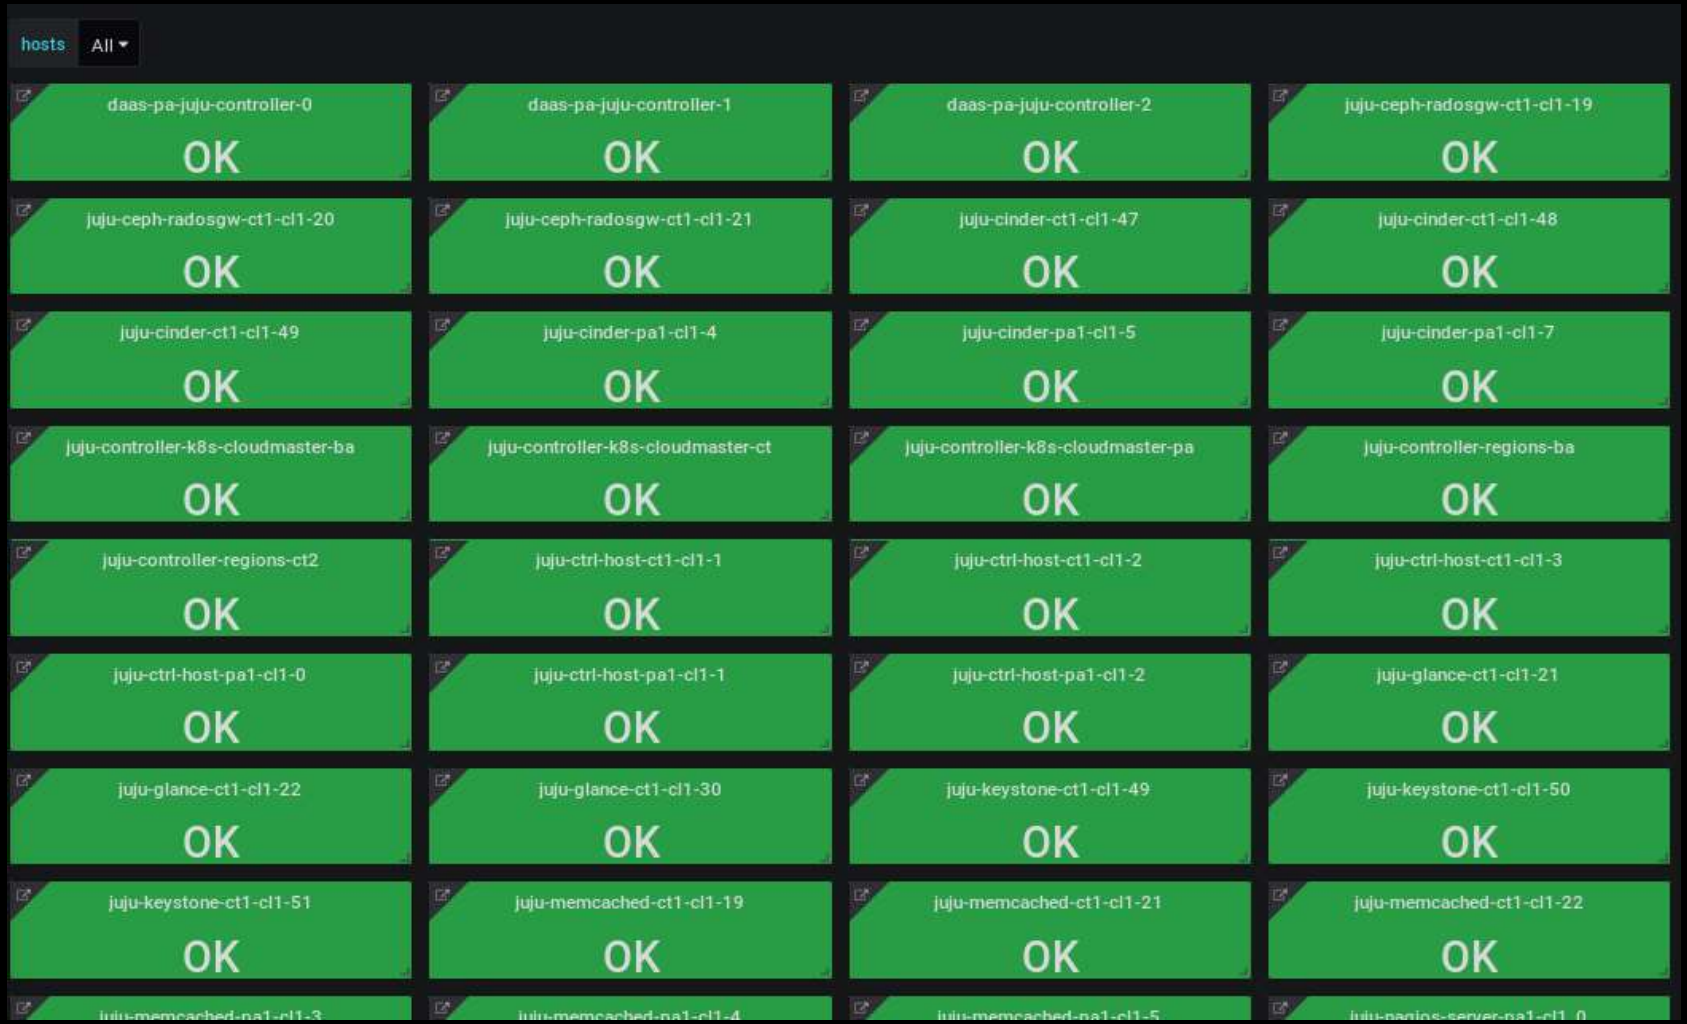
hosts
All
daas-pa-juju-controller-0
OK
daas-pa-juju-controller-1
OK
daas-pa-juju-controller-2
OK
juju-ceph-radosgw-ct1-cl1-19
OK
juju-ceph-radosgw-ct1-cl1-20
OK
juju-ceph-radosgw-ct1-cl1-21
OK
juju-cinder-ct1-cl1-47
OK
juju-cinder-ct1-cl1-48
OK
juju-cinder-ct1-cl1-49
OK
juju-cinder-pa1-cl1-4
OK
juju-cinder-pa1-cl1-5
OK
juju-cinder-pa1-cl1-7
OK
juju-controller-k8s-cloudmaster-ba
OK
juju-controller-k8s-cloudmaster-ct
OK
juju-controller-k8s-cloudmaster-pa
OK
juju-controller-regions-ba
OK
juju-controller-regions-ct2
OK
juju-ctrl-host-ct1-cl1-1
OK
juju-ctrl-host-ct1-cl1-2
OK
juju-ctrl-host-ct1-cl1-3
OK
juju-ctrl-host-pa1-cl1-0
OK
juju-ctrl-host-pa1-cl1-1
OK
juju-ctrl-host-pa1-cl1-2
OK
juju-glance-ct1-cl1-21
OK
juju-glance-ct1-cl1-22
OK
juju-glance-ct1-cl1-30
OK
juju-keystone-ct1-cl1-49
OK
juju-keystone-ct1-cl1-50
OK
juju-keystone-ct1-cl1-51
OK
juju-memcached-ct1-cl1-19
OK
juju-memcached-ct1-cl1-21
OK
juju-memcached-ct1-cl1-22
OK

Kubernetes → Prometheus → Grafana

- A new kid on the block: Kubernetes
  - container platform
  - inspired by **Google Borg**
- Prometheus
  - open source monitoring tool
  - inspired by the **Google Borg Monitor**
  - powerful query language (PromQL)
  - alerting
  - white-box monitoring
- Kubernetes supports Prometheus natively
- Grafana supports Prometheus natively

Proposed Architecture - updated

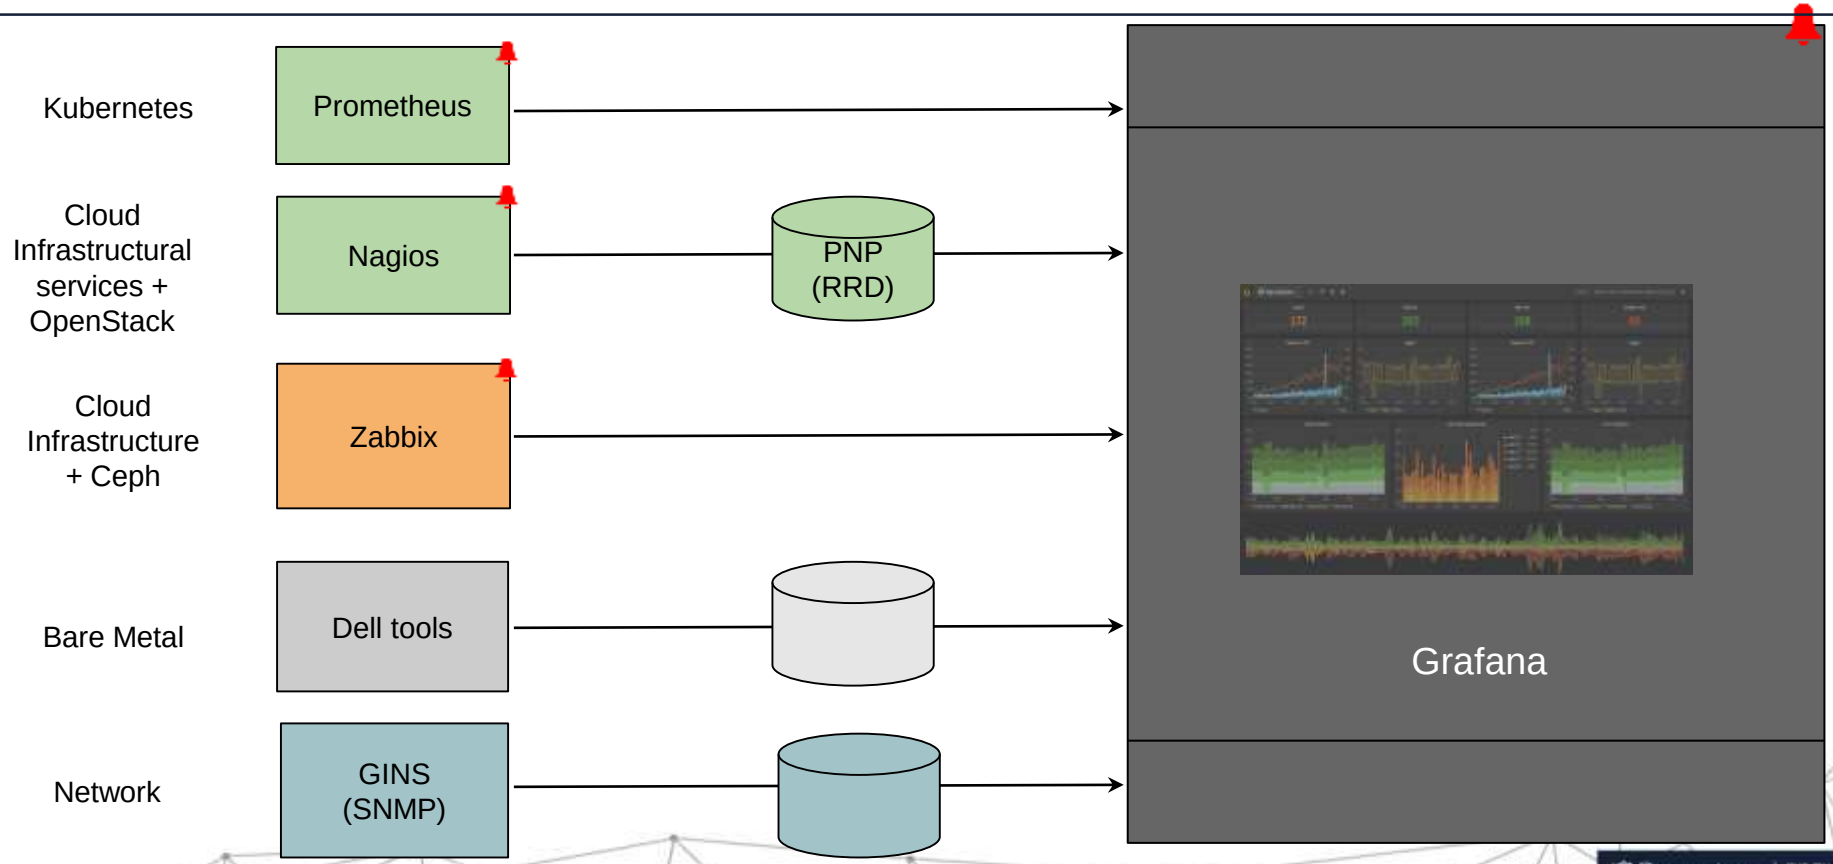

Instance All Namespace kube-system

all pods

Memory Working Set

Used	Total
49.17 GiB	1.60 TiB

Cpu Usage

Used	Total
3.29 cores	224.00 cores

Filesystem Usage

Used	Total
529.09 GiB	2.39 TiB

Network

each pod

Cpu Usage

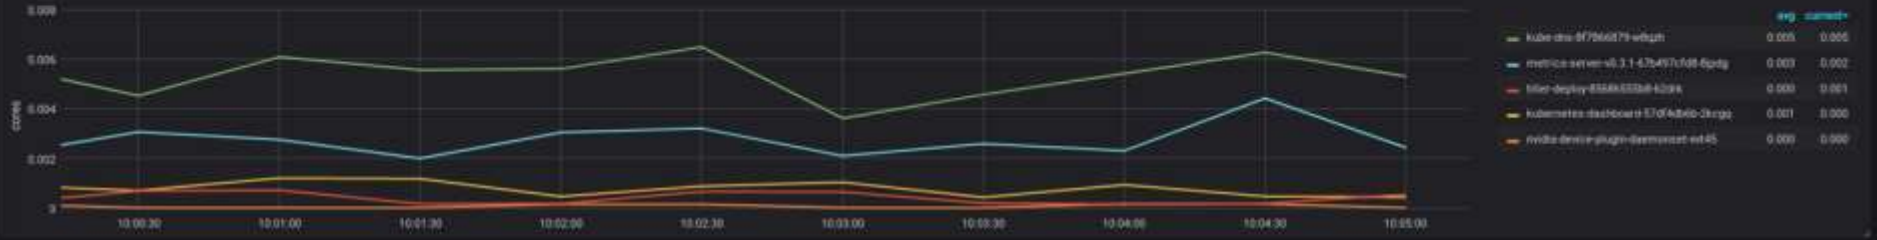

OpenStack Virtual Networks Monitoring

- The OpenStack tenants can create virtual networks, connected to virtual routers
- How to monitor if their network/virtual router is working?

OpenStack Virtual Networks Monitoring

1. Create a new project, with a VM inside
2. Add all the networks from all the tenants to the project
3. For each network, add a virtual NIC to the VM
4. From inside the VM, try to ping from each NIC a public address
5. Report the results to Nagios

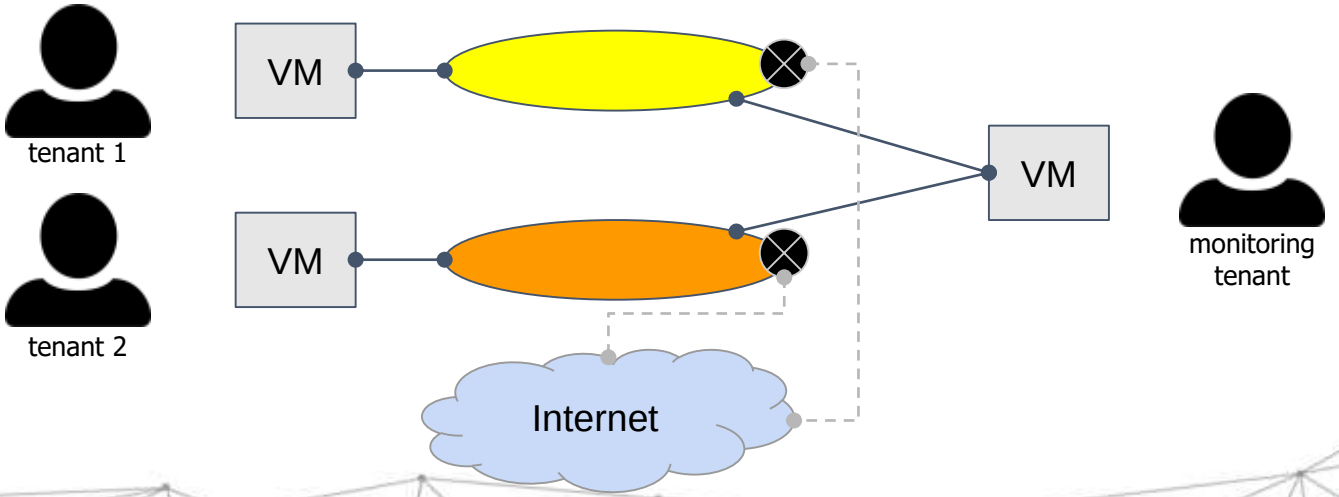

OpenStack Virtual Networks Monitoring

- Caveats
  - OpenStack VMs support a limited number of virtual NICs
    - A KVM limitation
    - In general we will need more than one VM
      - and an algorithm to distribute the networks among the VMs
  - How to ping choosing a different interface and routing without changing the IP address assignment mechanism?
    - policy routing
    - network namespaces **<------**

Service Overview For Host Group 'network-monitor-ct1'

network-monitor-ct1 (network-monitor-ct1)

Host	Status	Services	Actions
network-monitor-ct1-0	UP	1 OK	
network-monitor-ct1-1	UP	1 OK	
network-monitor-ct1-2	UP	1 OK	
network-monitor-ct1-3	UP	1 OK	

Host	Service	Status	Last Check	Duration	Attempt	Status Information
network-monitor-ct1-0	networks	OK	2019-11-22 09:32:36	1d 20h 17m 13s	1/4	networks OK: 12 checked

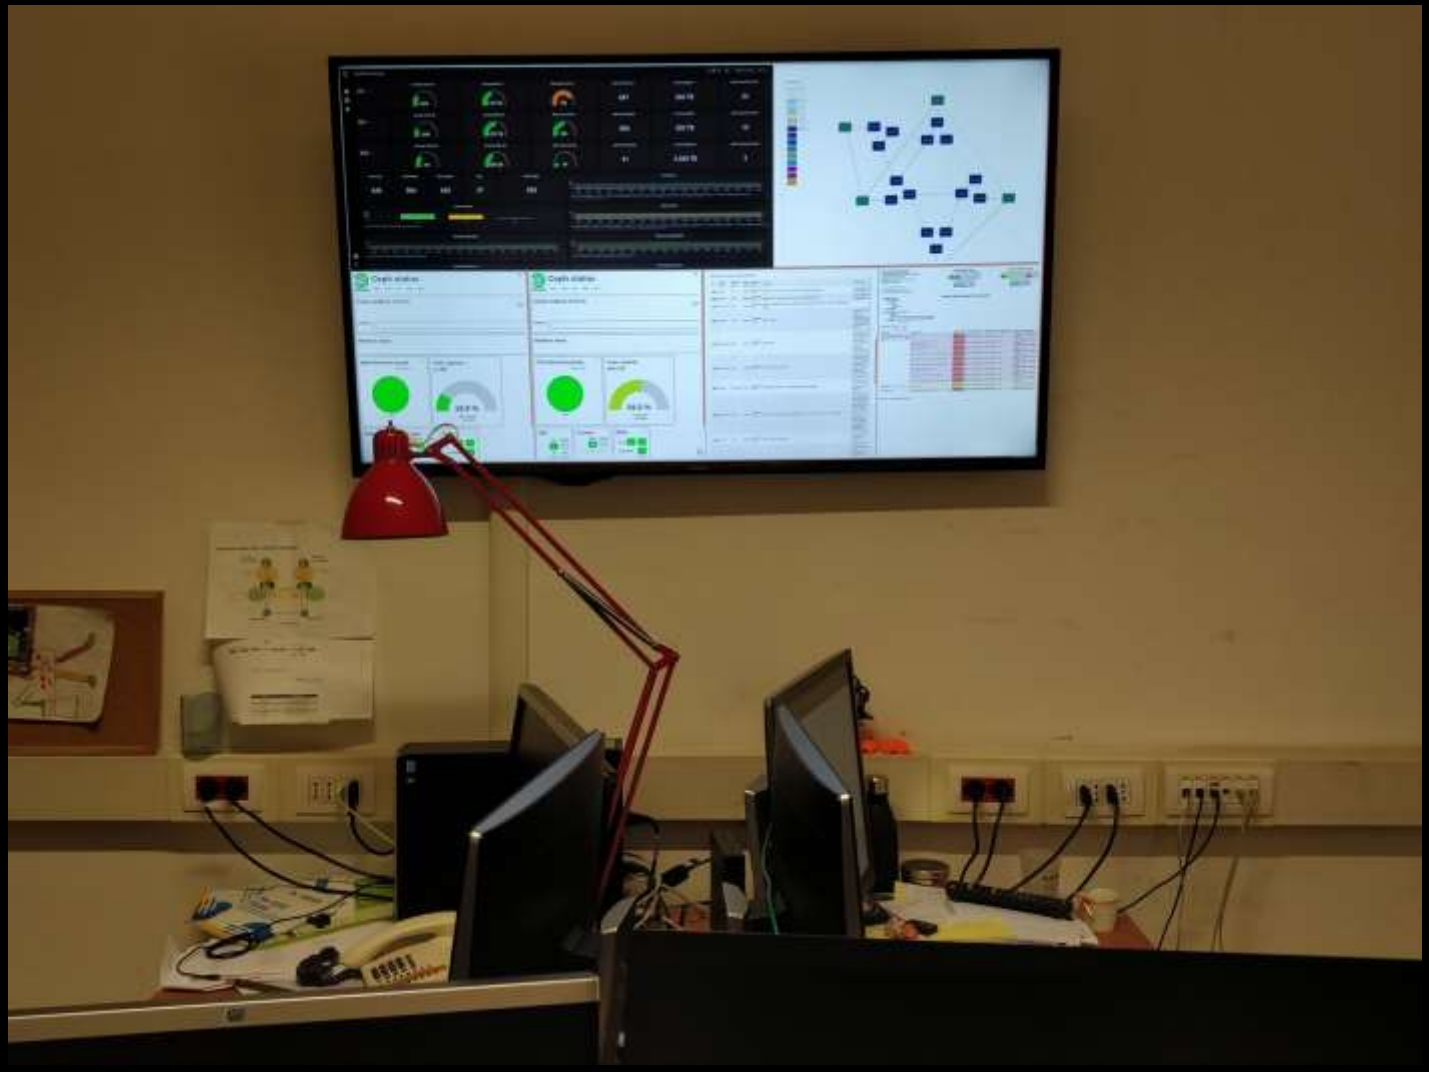

Considerations / Takeaways

- Nagios (vs. Prometheus)
  - Nagios has no query language
    - but it is very easy to develop new plugins
  - Nagios (+PNP) uses RRD based storage
    - not suitable for highly dynamic environments
      - e.g. cloud-native applications
    - but suitable for infrastructure monitoring
- Grafana performance (time to render graphs)
  - very good with Nagios+PNP
  - OK with Zabbix
  - can be slow with Prometheus

Future Work

- Self healing
  - react to well known bugs
  - with well known recipes
  - with some hysteresis / guard time
- AIOps - Artificial Intelligence for Operations
  - collect many many metrics
  - annotate incidents/events
  - train an AI using the collected data
  - profit!

Consortium GARR
THE ITALIAN EDUCATION & RESEARCH NETWORK

Thank you

Claudio Pisa - GARR - Distributed Computing and Storage Department
claudio.pisa@garr.it