



HARRY

Aggregate hardware usage metrics to optimise procurement of computing resources

Hervé Rousseau — CERN IT/CF-FPP

Agenda

Introduction

Technical architecture

Examples

Goals

- Assist CERN's IT/CF teams for adequate sizing of hardware and infrastructure
- Have a long term view (≥ 5 years) of hardware resources usage
- Ultimate goal: better resource acquisition and allocation process

Data needed (1)

Data

- CPU utilisation (all modes)
- Memory utilisation
- Disk I/O & storage (lower resolution)
- Network utilisation

Data needed (2)

Tags/Attributes

- Hostgroup (\approx Cluster)
- Location (Room and rack)
- IP Service (\approx Network “domain”)
- Purchase Order

Interlude

Wait a minute... we already have all this data !

- Yes, the raw data is here...
- but grouping by tags/attributes is not possible
- and displaying it is not visually attractive.

Interlude

Hey, but we already have all this data !




Figure: CPU Idle time on 7934 servers

Agenda

Introduction

Technical architecture

Examples

HARRY: Software used

On servers


- `hw_exporter`: exposes metrics and (optionally) tags

Data collection, aggregation and archival

- Tool to query PuppetDB (CERN specific)^a
- Prometheus
- Thanos

^aOnly custom development on this project

HARRY: Architecture

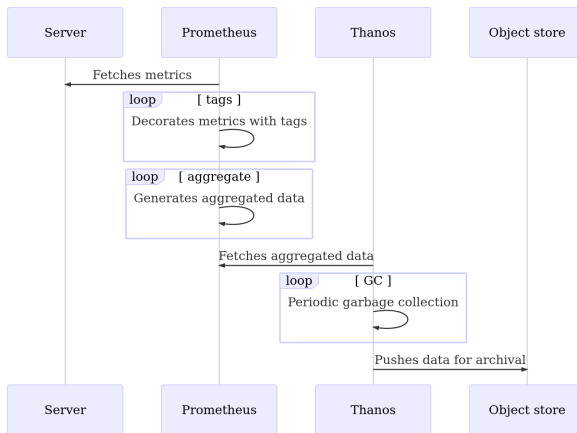


HARRY: On servers

hw_exporter instantiation:

```
class hardware::include::hw_exporter (  
  # Values set in module-level hiera data  
  Array[Stdlib::Host] $collectors = [],  
  Optional[Integer[1, 65535]] $port = 4242,  
  Optional[Boolean] $enabled = !$::facts['is_virtual'],  
  Optional[Boolean] $open_firewall = str2bool($::writefirewall),  
)
```

HARRY: Data flow



Each block is 2 hours long, "hot" data is kept for ~ 12 hours

Agenda


Introduction

Technical architecture

Examples


HARRY: Examples

CPU Idle time (on 7934 servers)




HARRY: Examples

CPU Idle time (on 7934 servers)




HARRY: Examples

Cluster with high network throughput requirements



HARRY: Examples

Request for high-IOPS NVMe storage



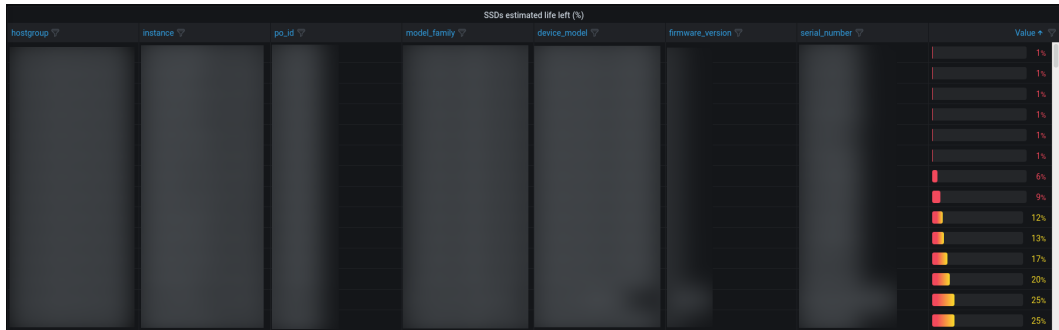
HARRY: Examples

SSDs estimated life left per purchase order




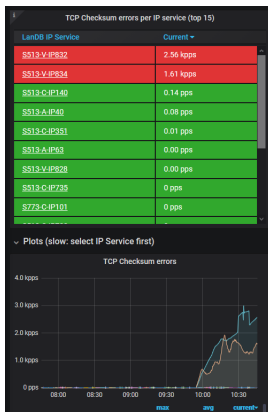
HARRY: Examples

SSDs estimated life left drill-down



HARRY: Troubleshooting tool


Datacentre-wide network outage



Next steps

- Delegate long-term storage to Central Monitoring
- Deprecate old home-made tools in favor of HARRY

Thank you





www.cern.ch

Links

- PuppetDB Query API v4
- Node Exporter
- Prometheus
- Thanos
- MONIT Documentation

HARRY Resources

Computing resources

Quantity	Usage	Specs
3 ^a	Collector	64 GB of RAM, 16 vCPU VMs
1	Aggregation	32 GB of RAM, 16 vCPU VMs

^aOne per availability zone

Storage resources

- Storage space usage: 1.6 TB ^a

^aAssuming Gorilla-style double delta encoding. Raw space usage: 4.1 TB