Managing build infrastructure at ALICE using Hashicorp Nomad

Computing in High Energy Physics 2023, Norfolk, VA

Timo Wilken[1] [2] Giulio Eulisse[2]

2 May 2023

[1] E-mail: timo.wilken@cern.ch

[2] for the ALICE Collaboration

What do we build?

- O^2: ALICE's online (data-taking) and offline (physics analysis, Monte-Carlo simulation) software suite
- Run 2 software still maintained for analysing old data
- pseudo-distribution of O^2 dependencies
  - designed to function on top of recent versions of CentOS, Alma, Ubuntu, MacOS, …
- nightly release builds, CI compilation checks, unit and integration tests
- 1 non-trivial CI check completed every 2 minutes, on average
  - plus lots of fast rebuilds where nothing has changed
- …all on multiple platforms (mostly) through containerization

Architecture overview

- Nomad for job scheduling
  - long-running jobs: custom continuous integration builders, Jenkins builders
  - web services: user account administration websites, tarball servers
  - scheduled/“cron” jobs: software repository maintenance and cleanup
- Consul
  - job discovery: `*.service.consul` DNS
  - Traefik auto-config for web access
  - job monitoring: simple health checks
- Vault stores secrets, using Consul as backend
- Prometheus and InfluxDB metrics of the whole cluster monitored and visualised using Grafana

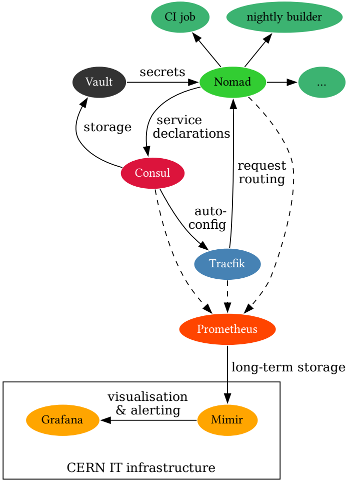

Reasons for switching away from Mesos and Aurora

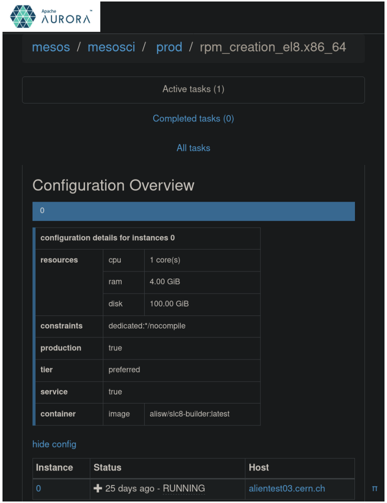

mesos / mesosci / prod / rpm_creation_el8.x86_64

Active tasks (1)

Completed tasks (0)

All tasks

Configuration Overview

0

configuration details for instances 0		
resources	cpu	1 core(s)
	ram	4.00 GiB
	disk	100.00 GiB
constraints	dedicated:*/nocompile	
production	true	
tier	preferred	
service	true	
container	image	alisw/slc8-builder:latest

hide config

Instance	Status	Host
0	✚ 25 days ago - RUNNING	alientest03.cern.ch

- previous stack: Mesos + Marathon + Apache Aurora
- Aurora not intensively developed any more
- requires Python 2 (EOL since 2020) on server and developers' machines
  - difficult to install, deploy and maintain
- some features difficult to integrate with or nonexistent
  - autoscaling (or even manual scaling without restarts of all jobs)
  - difficult to keep build caches "hot"
  - little monitoring and alerting integration

Improvements with Nomad + Consul + Vault

- simple deployment: static binary + systemd/launchd service + configuration = 3 files
- first-class support for web services: health checks, autoconfiguration
- better secrets management: Vault instead of passwords in a Git repo
- excellent monitoring & alerting support through Prometheus
  - resource use statistics (CPU, memory, disk)
  - alerts when build machines are unavailable or have problems
- …more features, for deeper future integration

Web services: health checks & Traefik autoconfiguration

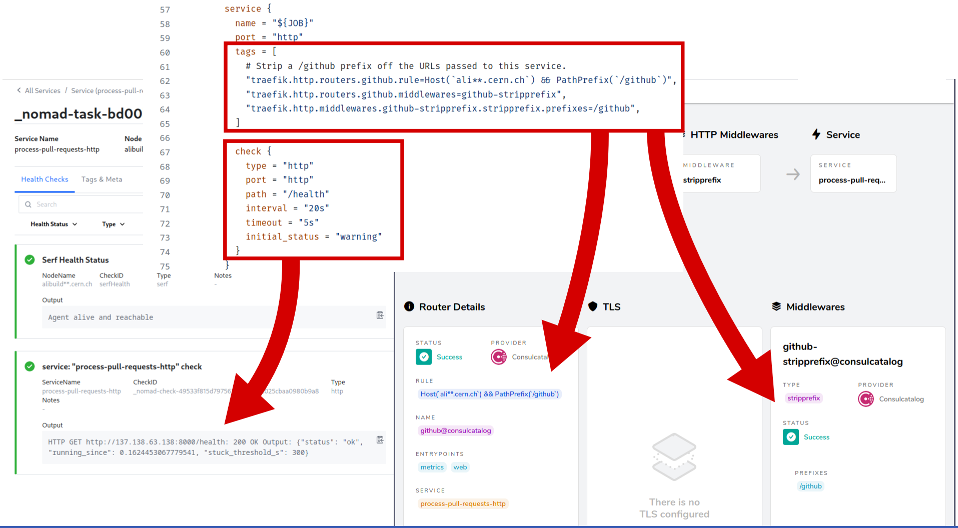

```
service {
  name = "${JOB}"
  port = "http"
  tags = [
    # Strip a /github prefix off the URLs passed to this service.
    "traefik.http.routers.github.rule=Host(`ali**.cern.ch`) && PathPrefix(`/github`)",
    "traefik.http.routers.github.middlewares=github-stripprefix",
    "traefik.http.middlewares.github-stripprefix.stripprefix.prefixes=/github",
  ]

  check {
    type = "http"
    port = "http"
    path = "/health"
    interval = "20s"
    timeout = "5s"
    initial_status = "warning"
  }
}
```

< All Services / Service (process-pull-r

_nomad-task-bd00

Service Name process-pull-requests-http

Node alibuil

Serf Health Status

NodeName: alibuild**.cern.ch
CheckID: serfHealth
Type: serf
Notes: -

Output

```
Agent alive and reachable
```

service: "process-pull-requests-http" check

ServiceName: process-pull-requests-http
CheckID: _nomad-check-49533f815d79756 025cbaa0980b9a8
Type: http
Notes: -

Output

```
HTTP GET http://137.138.63.138:8000/health: 200 OK Output: {"status": "ok",
"running_since": 0.1624453067779541, "stuck_threshold_s": 300}
```

HTTP Middlewares

MIDDLEWARE: stripprefix → **Service** SERVICE: process-pull-req...

Router Details

STATUS: Success
PROVIDER: Consulcatal
RULE: Host(`ali**.cern.ch`) && PathPrefix(`/github`)
NAME: github@consulcatalog
ENTRYPOINTS: metrics web
SERVICE: process-pull-requests-http

TLS

There is no TLS configured

Middlewares

github-stripprefix@consulcatalog

TYPE: stripprefix
PROVIDER: Consulcatalog
STATUS: Success
PREFIXES: /github

Monitoring example: nightly build performance

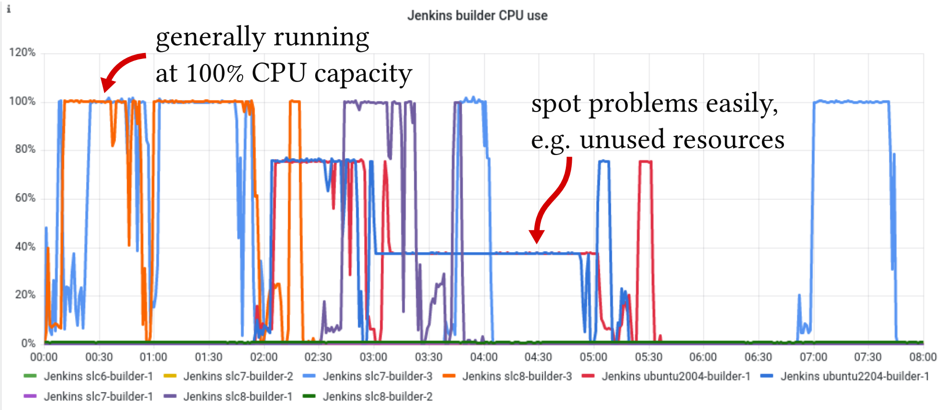

generally running at 100% CPU capacity

spot problems easily, e.g. unused resources

Figure: CPU use of a sequence of nightly builds as a fraction of total allocated CPU resources (usually the entire VM).

Monitoring example: nightly build performance

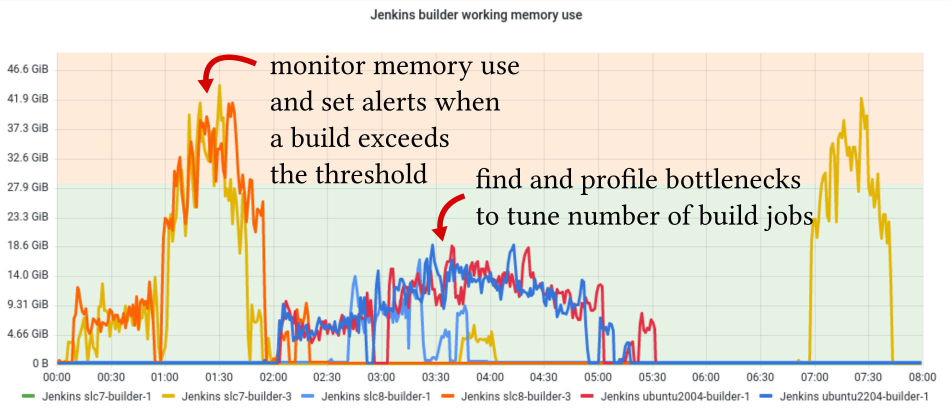

monitor memory use and set alerts when a build exceeds the threshold

find and profile bottlenecks to tune number of build jobs

Figure: Working memory (RSS) use of a sequence of nightly builds. Total available memory on a typical build VM in green.

Rough edges

1. Nomad's handling of disk space allocation
    - restarting daemon with non-empty disk confuses Nomad's accounting
    - can cause scheduling issues much further down the line
    - must manually clean up the node and restart the Nomad agent process
2. integration with CERN single sign-on
    - rely on Nomad/Consul/Vault tokens for authentication
    - could integrate SSO with Vault, which would then issue Nomad/Consul tokens
    - client certificate authentication is supported, so we use that in addition to tokens

Future work integrating build infra with Nomad

- “true” autoscaling, based on real-time demand
  - manual scaling already much smoother than previously: build caches are kept most of the time, existing builders uninterrupted
  - remaining challenge is cache invalidation: scaling often invalidates multiple gigabytes of cached builds
- temporary configuration (e.g. for testing software deployment) through Consul instead of text files
- get build secrets from Vault only when needed, instead of storing them in env variables and relying on sanitisation during build

Questions?

Useful links

- `aliBuild`: alisw.github.io/alibuild
- CI & ALICE software documentation: alisw.github.io
- ALICE O2 user guides: aliceo2group.github.io

Contact details

- timo.wilken@cern.ch