

Evolution of monitoring, accounting and alerting services at INFN-CNAF Tier1




Dal Pra Stefano^a, Falabella Antonio^a, Fattibene Enrico^a, Cincinelli Gianluca^b, Magnani Matteo^b, De Cristofaro Tiziano^b, Ruini Martin^b

a. INFN-CNAF, Bologna, Italy

b. Università di Bologna, Sede di Cesena, Italy
antonio.falabella@cnafe.infn.it

INFN Tier1 data center

- The INFN Tier1 at CNAF provides resource to experiments of LHC and 30 other in which INFN is involved
- >25000 CPU cores
- ~30PB of disk storage
- ~47PB of tape storage
- ~ 30TFlops of HPC Resources
- Elastic data center extension
 - ~15% Resources installed at INFN Bari-ReCaS (Level 2 VPN)
 - ~50% of CPU Resources at Cineca datacenter



Old monitoring infrastructure ¹

- In past years, the teams at Tier-1 sub-units were using several software packages:
 - Lemon
 - developed at CERN and customized to suit the needs of a datacenter managing scientific data
 - Nagios
 - mainly used for alerting purposes
 - a system based on Graphite database
 - a number of in-house developed services for web reporting



Nagios



New monitoring infrastructure ²

- Recently the adoption of off-the-shelf solution has been chosen
- The new CNAF monitoring infrastructure uses Sensu as alternative to Lemon and Nagios
 - Sensu is able to issue checks as well as metrics measurements (easy port of nagios probes)
 - A Sensu server for each CNAF Unit
 - Communication between client and server through dedicated RabbitMQ message queue
 - Configuration → Puppet, Foreman
 - Time series persistence → InfluxDB
 - Dashboard composition → Grafana



Migration

- Server monitoring and standard metrics not ported
 - provided as Sensu community plugins
- Specific probes and script for batch system, GPFS servers, TSM server, building operation has been adapted to Sensu



Historical Data Migration

- Migration only of historical measurements to InfluxDB from:
 - Graphite
 - Accounting and specific service monitoring
 - Easy migration between time series Databases
 - Lemon
 - Standard metrics
 - Developed a layer between backend Lemon Database (Oracle) and InfluxDB
 - Schema is not in a human-readable format
 - Not every metrics has been migrated
- Resampled measurements to match InfluxDB retention policies

Pledges Add-on

- Web app developed in-house to handle pledges
 - php, javascript
 - Backend PostgreSQL DB
- Flask-based server to exposed the PostgreSQL data to Grafana

Experiment	20160101	20160201	20160301	20160401	20160501	20160601	20160701
ALICE	10000	10000	10000	10000	10000	10000	10000
ATLAS	20000	20000	20000	20000	20000	20000	20000
LHCb	30000	30000	30000	30000	30000	30000	30000
COMPASS	40000	40000	40000	40000	40000	40000	40000
FAIR	50000	50000	50000	50000	50000	50000	50000
SPS	60000	60000	60000	60000	60000	60000	60000
TRIUMF	70000	70000	70000	70000	70000	70000	70000
INFN-CNAF	80000	80000	80000	80000	80000	80000	80000
INFN-BARI	90000	90000	90000	90000	90000	90000	90000
INFN-TRIESTE	100000	100000	100000	100000	100000	100000	100000



References

[1] Antonelli S. et. al. "INFN-CNAF Monitor and Control System", 2010, JPCS IOP Publishing 331 042032 2011

[2] Bovina S. et. al. "The evolution of monitoring system: the INFN-CNAF case study", 2017, Journal of Physics: Conf. Series 898