The evolution of monitoring system: the INFN-CNAF case study

Presenter:
D. Michelotto

Authors:
D. Michelotto
S. Bovina

diego.michelotto@cnaf.infn.it stefano.bovina@cnaf.infn.it

CHEP 2016
San Francisco, US

October 11, 2016

INFN-CNAF

CNAF

CNAF is the national center of INFN (Italian Institute for Nuclear Physics) "for the Research and Development in INFN Information and Communication Technologies". Being the central computing facility of INFN.

Mission

- Italian Tier 1 for WLCG infrastructure.
- Computing facility for 4 LHC experiments and many other astro-particle and neutrino-physics experiments.

Numbers

CNAF resources in numbers

- Core ~22000.
- Disk storage ~20PB.
- Tape storage ~34PB.
- Racks > 180.
- kHS06 ~250.

CNAF staff

- 50 persons
- 5 functional units

Targets

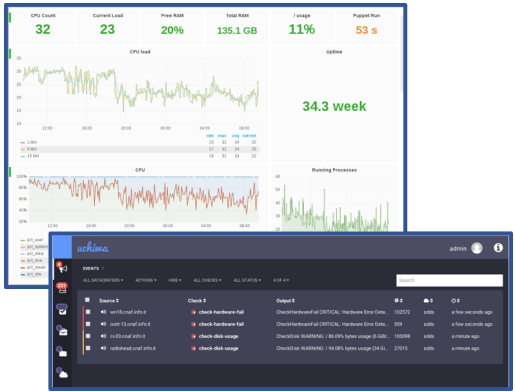

Targets

- Monitoring.
- Alerting.
- Management of monitoring through CM tools.
- Reusable data.
- Modern dashboards.
- Scalable to support CNAF numbers.
- High Available and Reliable.
- Re-usability of Nagios and Lemon scripts.

Past vs. Future

Past

- Nagios.
- Lemon.
- Home made probes.
- Home made sensors.
- Legacy UI

Future

- Central infrastructure.
- Sensu and Uchiwa.
- InfluxDB and Grafana.
- Comunity probes and sensors.
- Home made sensors and probes.

Sensu

What is Sensu

- Monitoring router.
- Schedule checks.
- Manage event actions for checks and metrics, such as send an email alert
- Use RabbitMQ as message broker.
- Use Redis to save check status.

Why Sensu

- Satisfy our requirements.
- We can reuse custum Nagios probe.
- Provide a rich API interface.
- Supported by large comunity.
- Provide modern dashboard.
- Suitable for dynamic infrastructures.
- Can scale at CNAF numbers.

InfluxDB

InfluxDB

- Time series DB born in 2013.
- Optimized for handling time series data.
- Written in Go.
- Easy to install, no external dependecies.
- Query language SQL like.
- Scalable.

CNAF numbers

- About 1500 nodes monitored.
- 5 db nodes, 1 for each functional unit.
- Installed version 1.0.0.

Performance and Optimizations

- Created 4 retention policy for data: 1 week, 1 month, 6 months, 1 year.
- Created continuous queries for data aggregation: 15 minutes, 30 minutes, 1 hour.
- Data is written into 1 week retention policy by default. Every 15 minutes/30 minutes/1 hour we downsample raw data with Continuous Query into 1 month/6 month/1 year retention policy respectively, to change the high frequency data into lower frequency data.

Dashboards

Uchiwa

- Displays real-time hosts/checks statuses.
- Written in Go and Angular JS.
- Retrive data from Sensu API.
- Used for:
  - Silence hosts/checks .
  - Filtering host and checks.
  - Remove client.

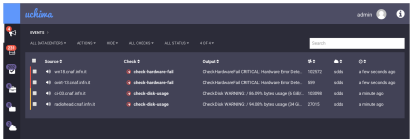

Grafana

- Draw graphs from several data sources.
- Based on Kibana 3.
- Multitenancy and read only dashboards.
- LDAP integration.

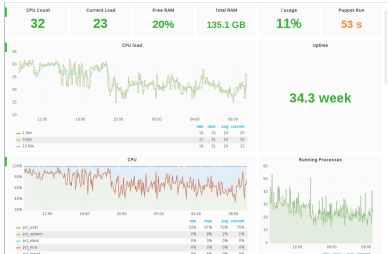

Architecture

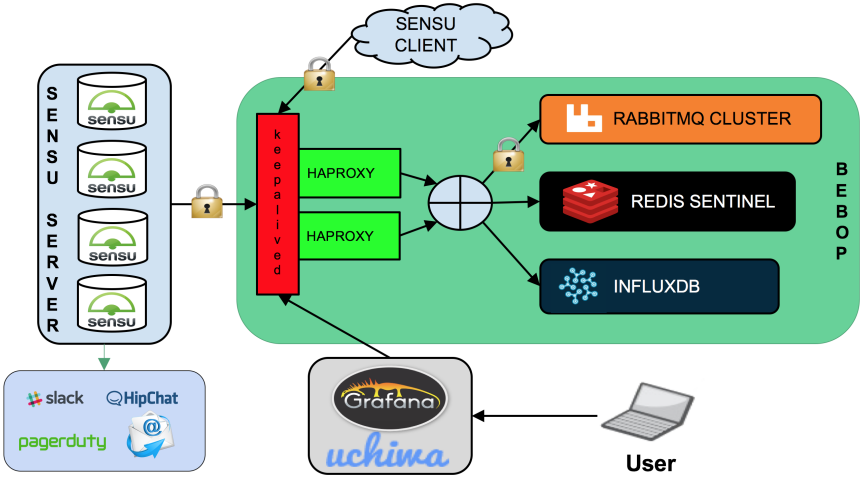

Architecture

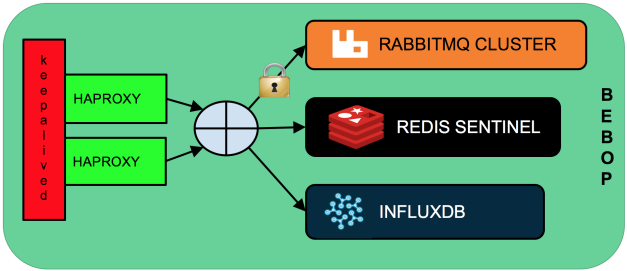

Architecture

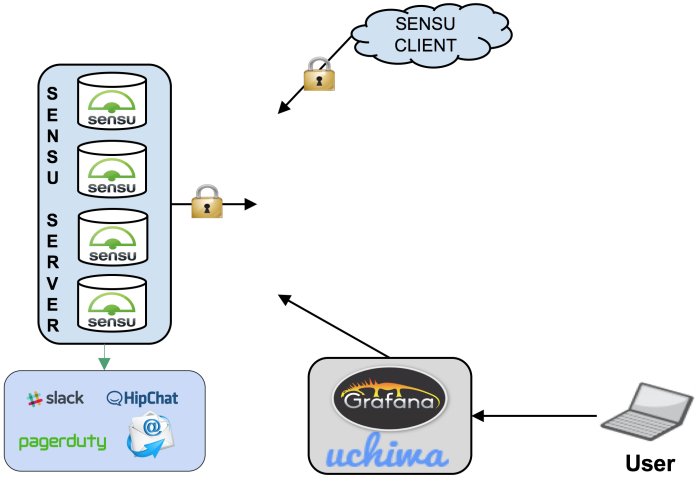

Architecture

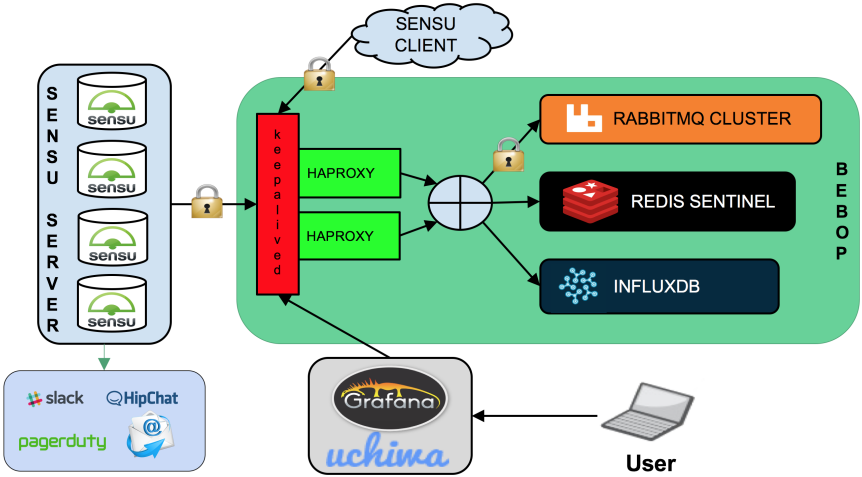

Future works

Future works

- Finish porting of probes and sensors.
- Monitoring data center networks.
- Optimizations.
- Scaling components if necessary.
- Decommission of Nagios and Lemon.

Conclusion

Conclusion

- Setup ready for production.
- About 1500 servers monitored.
- All the infrastructure managed by Puppet.
- Separated environment for each functional unit but unique infrastructure.

References

- https://sensuapp.org
- https://influxdata.com
- http://grafana.org/
- https://uchiwa.io/
- https://puppet.com

Thanks


THANK YOU
GRACIAS
BÏYAN
SHUKRIA
TINGKI
DANKSCHEEN
YAQHANYELAY
TASHAKKUR ATU
SUKSAMA
EKHMET
ARIGATO
SHUKURIA
GOZAIMASHITA
EFCHARISTO
KOMAPSUMNIDA
MAAKE
GRAZIE
MEHRBANI
PALDIES
JUSPAXAR
BOLZÏN
MERCI