

Monitoring at LHCb: Migrating to Icinga2, Puppet, Hierapro and Foreman Stack for Monitoring.



Hristo Mohamed
CERN - LHCb Experiment

Infrastructure intro

- The LHCb Online ~ 2600+ hosts currently monitored, majority of them on SLC6
- ~1700 in the High Level Trigger all NFS booted
- ~400 Embedded systems
- 300+ VMs
- 200+ Switches
- 65+ FreeNAS NFS Boot Servers
- Various site critical servers
 - Hypervisors
 - Various Storage solutions
 - DNS,DHCP,Puppet Master
 - And more and more...
- Webservers
- and even a few Desktops(Consoles)



The monitoring challenge

- Changing services
- Changing requirements
- Different environments
- Different ways of monitoring a service
- Massive deployment of changes



The choice for the stack

- We already had experience with Icinga
- Wanted to make Icinga 2 a bit more fancy



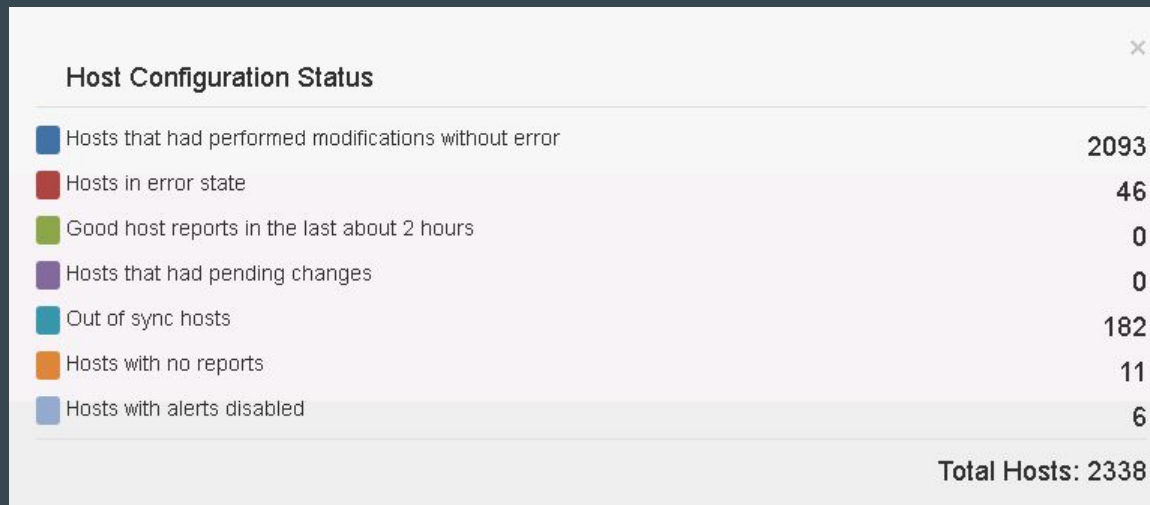


First add some Puppet

- Gathers facts and reports them to Foreman
- When a change is made, Foreman knows
- I promise a more interesting slide about Puppet later!

Foreman

- Host life cycle management tool
- Provides us with complete reports on Puppet runs
- Provides and keeps fact information about hosts



Host Configuration Status

Hosts that had performed modifications without error	2093
Hosts in error state	46
Good host reports in the last about 2 hours	0
Hosts that had pending changes	0
Out of sync hosts	182
Hosts with no reports	11
Hosts with alerts disabled	6

Total Hosts: 2338

Hiera

- Sadly not everyone host can run puppet :((Various switches, FreeNAS hosts, DDN)
- Defining “static” hosts - providing hardcoded facts about them
- Used to annoy everyone who’s editor is not set to 4 spaces

```
hosts:
  sxnas:
    facts:
      i2_type:
        - nfs_server
        - freeNas server
```



Icinga2 - auto generating everything!

- Hosts changing from development machines to staging/production and vice versa
- Various new hosts in various roles
- Various services have to be monitored, some in different way
- Do we hire an intern to do all the boring data entry in the config files?
- Autodiscovery - ugh, so many cases...

Tying it all together with little magical ruby

- Very simple to debug workflow
- Initiate a puppet run
- Simulate what would a real person do
 - Extract everything from Foreman and Hiera
 - Has anything changed in the config of already existing data?
 - Do we have new data?
 - Do we have phased out data?
 - If any change - restart Icinga and reread the config
- That's it!



Icinga some statistics

- Fully regenerating all the config files ~5-6 minutes
- Restart of the Icinga service ~ 5 minutes
- At the moment 2648 Monitored Hosts
- 32 000+ services



2498 / 0 UP

150 / 0 / 0 DOWN

0 / 0 / 0 UNREACHABLE

0 PENDING

150 / 2648 IN TOTAL



29262 / 0 OK

1108 / 0 / 1 WARNING

493 / 0 / 1092 CRITICAL

746 / 0 / 284 UNKNOWN

69 PENDING

3793 / 33055 IN TOTAL

Icinga some statistics

- Machine... a bit loaded, still able to support fully our needs
- Intel(R) Xeon(R) CPU E5420 @ 2.50GHz
- 2x10Gbit interfaces in active-backup mode
- Currently running on one server only, performing OK 400+ days of uptime
- Could be clusterized in the future

```
1  [||||||||||||||||||||||||||||||||| 84.5%]    5  [||||||||||||||||||||||||||||||||| 84.0%]
2  [||||||||||||||||||||||||||||||| 82.5%]    6  [||||||||||||||||||||||||||||||| 83.1%]
3  [||||||||||||||||||||||||||||||| 83.4%]    7  [||||||||||||||||||||||||||||||| 84.0%]
4  [||||||||||||||||||||||||||||||| 87.6%]    8  [||||||||||||||||||||||||||||||| 83.1%]
Mem [||||||||||||||||||||||||||||||| 9166/11910MB]  Tasks: 654, 174 thr; 4 running
Swp [||||||||||||||||||||||||||| 0/0MB]           Load average: 16.47 14.60 14.98
                                           Uptime: 405 days(!), 19:14:19
```

Adding and phasing out a host

- Adding host connected to puppet = automatic
- Adding a host that cannot run puppet = still done by hand :(
- Working on jailing Puppet for FreeNAS
- Working on running Puppet on Windows machines
- Deleting a host = deleting it from Foreman/Removing it from Hiera

FreeNas 9.3 - “We finally stopped spamming”

- No longer emails with “Nothing new, everything is okay, carry on!”
- No need for custom scripts to get basic reporting of ZFS volume status, Disk Status, System Crashes and Security reports
- Still cannot update root’s email address with their API :(API not mature enough
- Editing /etc/aliases with root: WHATEVER@WHATEVER.COM for mass deployment without using the GUI



Case study of a difficult case: FreeNAS

- Reporting to Icinga2
- No need to install anything on the FreeNAS
- Using NSCA passive checks, running NSCA daemon on Icinga 2 machine
- Running some simple Perl and Cron on the FreeNas boxes

Monitoring FreeNas: NSCA strikes back

- NSCA “works”, before “randomly” crashing
- EPIPE (Broken pipe) hits first
- -1 ENXIO Joined the party
- A few /dev/urandom - EAGAIN (Resource temporarily unavailable)”

Questions time!

