Upgrade and integration of the Configuration and Monitoring tools for the ATLAS Online farm

ATLAS TDAQ System Administration Team

S. Ballestrero, F. Brasolin, G.-L. Darlea, I. Dumitru, D.A. Scannicchio, M.S. Twomey, M.L. Valsan, A. Zaytsev

atlas-tdaq-sysadmins@cern.ch

Introduction

The ATLAS Online farm is a non-homogeneous cluster of nearly 3000 PCs which run the data acquisition, trigger and control of the ATLAS detector. The systems are configured and monitored by a combination of open-source tools, and tools developed in-house for specific requirements.

Centralized Configuration Management Tools (CMT) allow our team to quickly provision new systems or replace faulty hardware. The in house developed ConfDB v1 has been replaced by the much improved ConfDB v2. The well-proven Quattor CMT is now complemented by Puppet, allowing us more comprehensive automation even for complex configurations.

The well-performing but ageing monitoring system is being upgraded, to take advantage of more modern tools now available.

NetBooted Nodes

- ~2350 nodes boot the Scientific Linux CERN 5 OS via PXE
- ~80 Local File Server (LFS) hosts provide DHCP, PXE, TFTP for booting, `/usr` read-only directory via NFS.
- Configuration of DHCP, PXE and boot parameters provided by ConfDB, our CMT for NetBoot nodes which is described in a separate poster
  (Centralized configuration system for a large scale farm of network booted computers)

Boot With Me tool
- Generates PXE boot images (kernel + RAMdisk root) and `/usr`
- Uses a reference SLC5 VM image as source

BWM post-boot script system
- Hierarchy of shell scripts, configures services, disk and NFS mounts etc
- Uses the standardised hostname to decide which sequence of scripts
- Stored on central Network-Attached Storage, executed by the client

BWM puppet
- Start to introduce Puppet profiles to replace BWM scripts
- Improve consistency and maintainability
- Serverless configuration, for scalability, using the NAS as storage

A Subversion repository is used to track changes of the BWM image creation configuration, of the post-boot scripts and Puppet profiles.

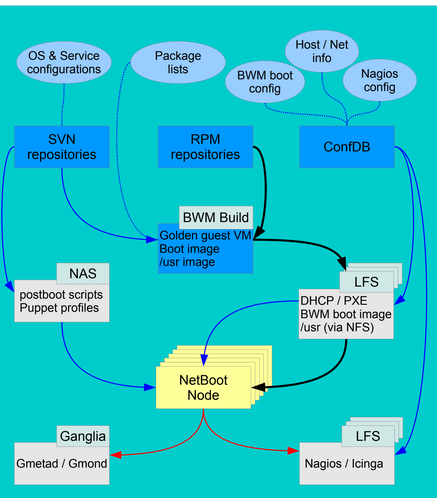

A netbooted node receives OS and configuration information from various sources (blue arrows). Its boot and operating system are provided by BWM via its LFS (**black arrows**). Monitoring information is collected by Nagios server (usually the boot LFS); some hosts also send more detailed performance information to the dedicated Ganglia server (red arrows).

Monitoring system - alerting and performance

Existing system based on Nagios 2.5
- Scalability: each LFS runs an independent Nagios server instance with a limited number of clients
- Configuration of each Nagios instance generated by ConfDB
- Overall view generated by a web GUI developed in-house, displaying data from Nagios and ConfDB
- Detailed monitoring of hardware health status via IPMI
- Good coverage of alerting for critical OS issues

Ganglia, introduced in 2011
- Detailed performance monitoring, excellent and easy scalability
- Flexible Web interface with advanced visualisation functionality
- Used for ~300 special purpose nodes
- Per-host configuration is client-side only, managed by Puppet.
- Limited alerting functionality in latest GangliaWeb versions, not sufficient for us
- Testing integration with Nagios/Icinga, the functionality is promising but reliability and scalability need to be proven

An overall upgrade of the monitoring system, based on Icinga and currently in a testing phase, is described in a separate poster
(Tools and strategies to monitor the ATLAS online computing farm)

Virtualization

Virtual Machines have been introduced for special purpose services, not for computing or data flow:
- low-load but critical services, e.g Quattor, Puppet, TDAQ OKSCVS
- complement existing systems and increase redundancy e.g. Nagios for core monitoring
- misc test/validation servers, web development etc.

VMs increase the overall system flexibility, allow easier testing and better user of HW resources.

Our new management tools allow us to handle effectively the many special purpose configurations.

Standardizing on KVM hypervisor on SLC6 physical hosts

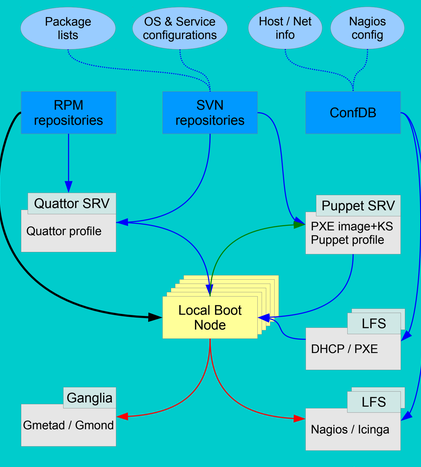

A node with locally installed OS receives configuration information from various sources (blue arrows); the Puppet server uses node status information (green arrow) to compile the profile. Its OS installation and updates are provided directly by the repository (**black arrow**). Monitoring information is collected by one of the Nagios servers and more detailed performance information is sent to the dedicated Ganglia server (red arrows).

Local Boot nodes

Provisioning by PXE + KickStart
- DHCP+PXE provided by an LFS, from ConfDB info
- Kickstart files generated by template-based system

Quattor
- CERN standard Configuration Management Tool
- Production system, managing 237 hosts in the Online Farm
- Tight control on installed packages
- Lack of flexibility for complex configuration/service dependencies
- Multiple languages for implementing modules

Puppet
- Widespread industry adoption, active development
- Full features, high flexibility
- Gentler learning curve
- Focus on consistency and idempotence
- In production, manages exclusively 25 complex servers and complements Quattor on the remaining 237
- Planned to completely replace Quattor on SLC6

For both Quattor and Puppet the configuration code is maintained in a Revision Control System (Subversion).

Conclusions

We are introducing in the ATLAS Online Farm new strategies for configuration management and monitoring. They are performing very well and are providing the expected improvement in maintainability, extending the coverage and flexibility. In addition, they give better control of the system, and deeper insight into its status.

All text/code based configurations are now covered by Subversion, allowing for safe change management, and together with the multiple data sources integrated in ConfDB, this provides a clear structure and a well identified information flow.

The introduction of these new technologies in the production system has been gradual and fully transparent, and will continue during 2012, to be completed in the LHC Long Shutdown.

Corresponding author: S. Ballestrero <sergio.ballestrero@cern.ch>