Next Generation Monitoring

Rob Fay (fay@hep.ph.liv.ac.uk), John Bland (jbland@hep.ph.liv.ac.uk),
Stephen Jones (sjones@hep.ph.liv.ac.uk)

Overview

Monitoring IT infrastructure is essential to maximize availability and minimize disruption, enabling rapid intervention by detecting failures and developing issues.

The HEP group at Liverpool have been working on a project to modernize local monitoring infrastructure (previously provided using Nagios) for both HEP systems and Liverpool Tier-2, with the goal of streamlining configuration and maintenance, increasing coverage, and improving visualization and analytical capabilities. Here we present progress to date, with the developed monitoring infrastructure and the tools used to build it.

The system is configured with Puppet. Basic system checks are configured in Puppet using Hiera, and managed by Sensu. Centralised logging is managed with Elasticsearch, together with Logstash and Filebeat. Kibana provides an interface for interactive analysis, including visualization and dashboards. Metric collection is also configured in Puppet, managed primarily by Sensu, and sent to Graphite, with Grafana providing a visualization and dashboard tool.

The Uchiwa dashboard for Sensu provides a web interface for viewing infrastructure status. Alert capabilities are provided via external handlers. Liverpool are developing a custom handler to provide an easily configurable, extensible and maintainable alert facility.

Integrated Configuration of Monitoring Framework with *Puppet*

The open source configuration management software Puppet is used to manage configuration of Liverpool's HEP systems. With Puppet, in combination with Sensu, Graphite, and ELK, monitoring checks, metric collection and log collection can be defined in the same place the subject services are configured. This makes it much easier to define and maintain the monitoring framework, particularly as nodes and services are added over time. Stale or incomplete monitoring definitions are a common problem with traditional statically configured systems such as Nagios.

In addition, Hiera is used with Puppet to provide a hierarchical configuration, enabling, for example, all CentOS 7 systems to easily share a base configuration of installed packages and monitoring checks, all defined together in one place:

Hiera Node level definition contains:

- node-specific classes
- node-specific class configuration
- node-specific metric collection
- node-specific monitoring checks

Hiera OS level definition contains:

- OS classes
- OS class configuration
- OS metric collection
- OS monitoring checks

Monitor with *Sensu*

Sensu is a monitoring platform designed to replace Nagios, featuring monitoring plugin compatibility accordingly, allowing existing Nagios checks to be easily ported to Sensu.

Two types of check are defined in the Liverpool configuration:

- 'Standalone' checks configured on the clients, which report to the server.

- Checks run from the server, used where standalone checks are not applicable (e.g. systems where Sensu cannot be run, such as network switches, where checks are made from an external system via SNMP).

The Uchiwa dashboard is used to see status.

Alerts are sent using sensu handlers.

Limitations: Systems and checks don't appear on the dashboard until they've run at least once and reported.

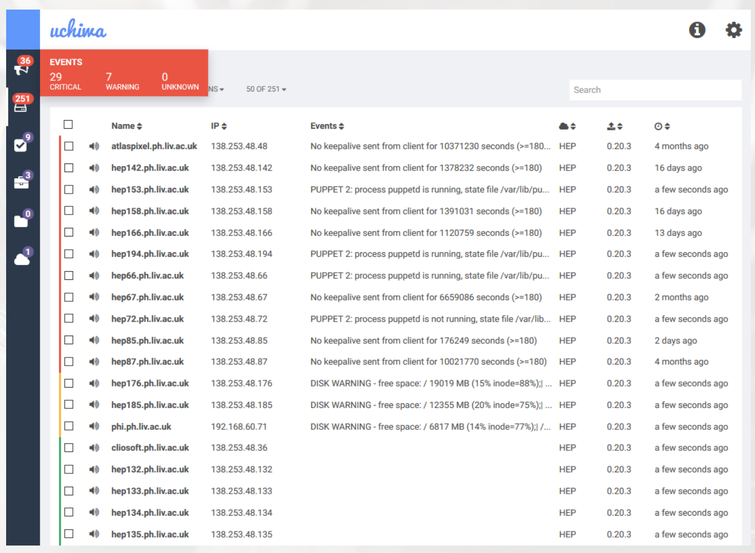

uchiwa dashboard showing the status of HEP desktops

Collect Metrics with *collectd, Graphite & Grafana*

Metrics are collected with collectd, configured on nodes by Puppet.

The metrics are then sent and stored by Carbon running on the Graphite server.

Grafana provides advanced visualisation.

Sensu checks can then be configured to run on the Graphite server, running against metrics, e.g. using rolling averages.

Compared to Ganglia, collectd with Graphite & Grafana requires significantly more installation and configuration for collection of basic site metrics.

Collectd with Graphite & Grafana is easier to expand and considerably more powerful than Ganglia.

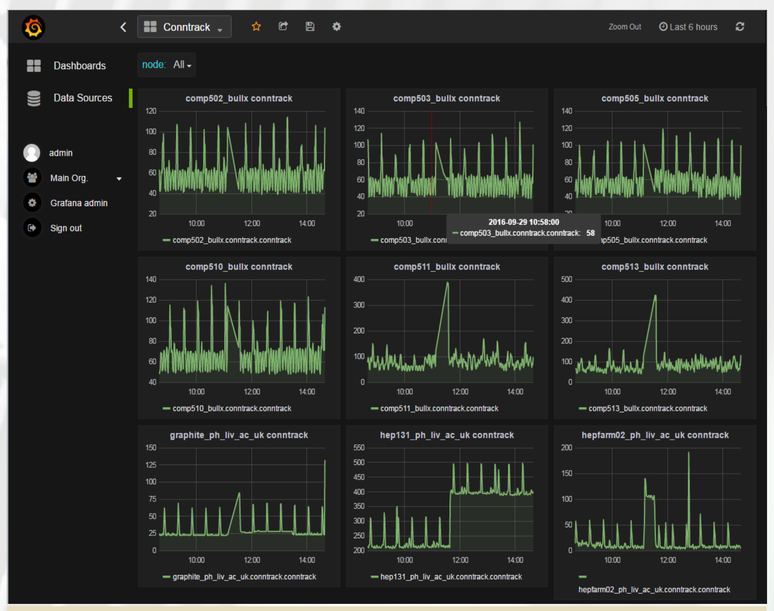

Grafana visualisation showing connection tracking statistics over the last six hours

Analyse Logs with *ELK* (Elasticsearch, Logstash, Kibana)

Logs are collected with filebeat, processed with logstash, and stored in elasticsearch, all configured with Puppet.

Kibana provides an interface for searching and visualisation of logs.

While elasticsearch and logstash modules from the Puppet Forge help, installation and configuration of the stack is potentially complex.

For logs where pre-written grok filters are either not available or unsuitable, defining efficient and correct grok filters is a non-trivial task requiring a degree of expertise.

Once configured, ELK provides analytical and visualisation capabilities that are hard to match.

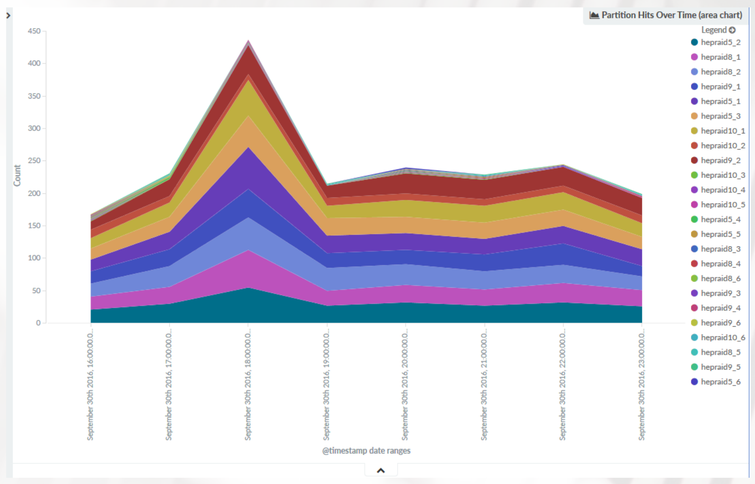

Kibana visualisation showing DPM filestore partition hits over time

Handle Events with *'Lerts*

'Lerts is an event handler in development to process Sensu events and replace Nagios notifications.

We found that with a sufficiently large infrastructure, even one that's in principle largely static, an essentially static configuration of event handling is difficult to accurately and comprehensively define, resulting in excessive alerts (which can be overlooked due to volume) or inadequate alerting. Software that does not allow simple tweaking, instead requiring relatively complex static configuration defined in text files to be edited, tested, and services restarted, can become unwieldy, resulting in a potential degradation of the service over time.

'lerts is an event handler, designed to provide dynamic adjustable handling of events, configurable on both a system and individual user level, allowing alert handling to be adjusted easily over time to reflect evolving requirements and circumstances.

Sensu sends events to 'lerts, 'lerts converts them to an internal format, filters are then applied and output plugins invoked accordingly.

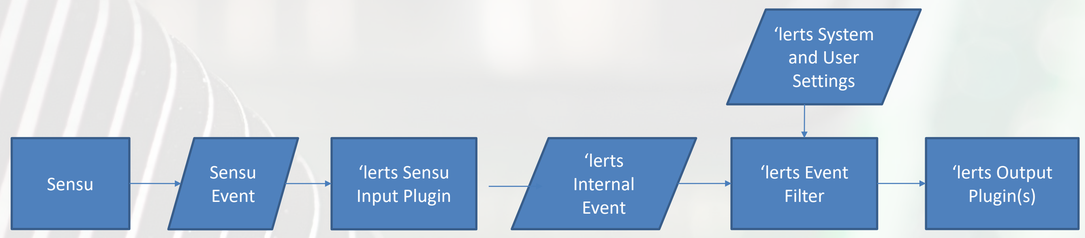

Conclusions and Recommendations

- A solution using configuration management with standalone checks defined on systems enables services and checks to be defined in one place, simplifying configuration and reducing the likelihood of omissions.
- Graphite and Grafana are initially more demanding to set up for a basic configuration than Ganglia, but can be more easily extended and offer significantly greater visualisation capabilities.
- Centralised logging with ELK enables faster and more comprehensive debugging and analysis, but can be demanding in resources, and complex to set up initially.
- The assessed combination of software overall offers great potential for monitoring and analysis, but at a potentially high cost in terms of manpower and resources which may not be available at all Tier 2 sites; for effective widespread adoption, it is likely that collaborative, maintained effort would be necessary to reduce overheads through development and provision of common installation instructions and configurations.

UNIVERSITY OF LIVERPOOL

GridPP
UK Computing for Particle Physics