

Configuration Management at the USAtlas Tier1 Facility

Jason A. Smith
Brookhaven National Lab

Components

- Cobbler/RHEV – New system provisioning
- Puppet – Centralized config management
 - Complete service config after provisioning
 - Dashboard monitoring & change auditing
- Git – Puppet catalog repository
 - Distributed development & historical record
- GLPI – Asset mgmt. & node classification
 - Fusioninventory-agent: auto asset inventory
 - ENC uses GLPI, custom DB & dashboard

Provisioning

- Cobbler for hardware installs:
 - Powerful Cheetah templating language and config/code reuse with “Snippets”
 - Single ks template used for most systems
 - Specify OS version & arch, network (MAC, IP, etc) & template metadata to install base OS, including fusioninventory-agent & puppet
- RHEV 3.0 for virtual machines:
 - Single template image used for new systems
 - 10 node cluster with 4TB of shared fiber storage
 - Should support several hundred VMs

Why Git?

- Distributed version control system
- Faster, completely localized project copies
 - Commits and other work can be done offline
 - Local copy contains complete history
- Reduced single point of repository failure
 - Git can merge changes between many “servers”
- Simple, fast & clean branching (and merging)
 - Branches easily merged with other branches
 - All changes can be treated as branches

Why Puppet?

- Cfengine, puppet, chef, etch, bcfg2, AutomateIt
- Puppet was selected for several reasons:
 - Simple yet powerful DSL (Domain-Specific Lang) & RAL (Resource Abstraction Layer)
 - Explicitly declared dependency graphing model
 - Provides better deterministic state convergence
 - Central config catalog & dependency resolution
 - Better security, conflict resolution & logic analysis
 - Web dashboard, GraphViz config visualization
 - Long history, stable codebase, large user base
 - Free OpenSource (optional commercial support)

GLPI Node Classification

The screenshot displays the GLPI web interface for configuring a server node. The breadcrumb trail is **Central > Inventory > Computers**. The node ID is **ID 2**, and it was last updated on **2012-02-15 16:00** (Imported from OCSNG).

Node Information:

- Name: gcmaster02
- Type: -----
- Model: PowerEdge R410
- Location: BCF
- Manufacturer: Dell Inc.
- OS: Red Hat Enterprise Linux Workstation release 6.2 (Santiago)
- OS Version: 2.6.32-220.4.2.el6.x86_64
- Service Pack: #1 SMP Mon Feb 6 16:39:28 EST 2012
- OS serial: [Empty]
- OS Product ID: ID-1000049759
- Auto update OCSNG: Yes

Administrative Information:

- Contact: root
- Contact Number: 382
- User: [Nobody]
- Group: 382
- Technician in charge of the hardware: Jason Smith
- Network: 382
- Domain: rcf.bnl.gov
- Serial Number: FQ5TLM1
- Inventory number: 382
- Status: testing
- Update Source: -----

Additional Data:

- Last OCS inventory date: 2012-03-08 16:16
- Import date in GLPI: 2012-03-08 16:20
- Server localhost, Agent: FusionInventory-Agent_v2.1.14
- Comments: x86_64/00-00-26 09:27:10 Swap: 8191

Puppet Configuration:

Puppet Classes:

```
base afs lvm::fs(fs=[/var:sysvg:10G]) ganglia::node(cluster=gce_servers)
yumrepo::conf(repos=[testing:99]) git glpi mysql::server
glpi::fusioninventory-agent puppet::client(noclient=1) puppet::server
puppet::dashboard
```

Puppet Parameters:

```
iptables_allow_tcp_ports=[https,8140,130.199.6.238@3000]
httpd_www_fs_size=10G mysql_db_fs_size=20G backup_fs_size=40G
git_proxy_server=https://vproxytest02.racf.bnl.gov
git_allow_from=130.199.6.238 glpi_allow_from=130.199.6.238
ssh_root_key_list=[jd,mizuki,pryor,raot,smithj4,willsk]
```

Buttons: **Update** and **Delete** are visible at the bottom of the configuration form.

Puppet Environments

- Currently using 3 puppet environments linked to git branches:
 - Development: extensive module changes
 - Testing: small changes and wider testing
 - Changes staged for production
 - Production: main server management
 - Changes must be approved before they are merged into the production branch/environment
- Git branches are automatically sync'ed to puppet environments by push hooks.
 - Also verifies puppet syntax and other checks

Production Approval

Git/Puppet updates to production that are pending approval.

Hello Jason A. Smith, there are currently 2 changes waiting for approval:

Date	Age	User	Changes	Changelog	Approve	Reject
Fri Mar 9 15:43:12 2012	2 days	Zhenping Liu	diff	pending-zhliu-cb36590-20120309T204312UTC	merge	delete
Mon Mar 12 10:18:27 2012	1 minute	Jason A. Smith	diff	pending-smithj4-cb36590-20120312T141827UTC	merge	delete

Instructions:

- The table above lists all changes to puppet's production environment that are currently pending approval.
- The diff link in the Changes column uses the cgit interface to display the detailed changes to all files contained in that pending update.
- The branch link in the Changelog column uses the cgit interface to display the commit history of that branch since it diverged from production.
- Use the merge link in the Approve column to accept the changes and merge them into production.
- Use the delete link in the Reject column to delete the branch if you do not want it merged into production.
- Email notifications are sent after confirmation of the chosen action.
- A side effect of this approval process is that you might see a lot of these old temporary pending branches accumulate in your locally cloned repo. You can clean these up by using the "git remote prune origin" command.

Cgit Diff View

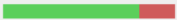
git index : puppet/catalog development switch
Git Repository for the RACF Puppet Catalog. RACF Puppet Master

summary refs log tree commit **diff** stats log msg search

path: root/gce/glpi/manifests/fusioninventory-agent.pp

[Side-by-side diff](#)

Diffstat (limited to 'gce/glpi/manifests/fusioninventory-agent.pp') ([more/less](#) context) ([ignore](#) whitespace changes)

-rw-r--r-- gce/glpi/manifests/fusioninventory-agent.pp 24 

1 files changed, 19 insertions, 5 deletions

```
diff --git a/gce/glpi/manifests/fusioninventory-agent.pp b/gce/glpi/manifests/fusioninventory-agent.pp
index 93ed6fb..dfd2f40 100644
--- a/gce/glpi/manifests/fusioninventory-agent.pp
+++ b/gce/glpi/manifests/fusioninventory-agent.pp
@@ -1,8 +1,22 @@
-class glpi::fusioninventory-agent {
+class glpi::fusioninventory-agent ( $server=undef ) {
+  # Check for parameterized class invocation or global parameter:
+  if ( $server ) {
+    # Called as a parameterized class, use that server name:
+    $server_name = $server
+  } elsif ( $glpi_server ) {
+    # Global parameter is set, use that server name:
+    $server_name = $glpi_server
+  } else {
+    # Default puppet server name:
+    $server_name = 'puppet.racf.bnl.gov'
+  }
+
+  # Install the fusioninventory-agent package:
+  package { 'fusioninventory-agent':
-    ensure => installed,
-    require => Yumrepo[ 'base' ],
+    ensure => latest,
+    schedule => daily,
+    require => Yumrepo[ 'base' ],
+    notify => Service[ 'fusioninventory-agent' ],
+  }
}
```

Puppet Dashboard

puppet dashboard v1.1.1 » Home • Nodes • **Groups** • Classes • Reports

Nodes 2579

- Currently successful 1829
- Currently failing 10
- Ever succeeded 1833
- Ever failed 175
- Never reported 740
- Not currently reporting 1691
- Hidden 0

File Search

Custom query

Add node

Class

- afs 0
- base 41
- bnl_banner 1
- cloudtestbed_vm 0
- cvms 4
- desktop 21
- dns 0
- epel 0
- frontier 4
- httd_host 1

Add class

Group

- Atlas dCache** 20
- Desktops 21
- Phenix dCache 24
- Web servers 9

Add group

Group: Atlas dCache Edit Destroy

Parameters
— No parameters —

Groups
— No groups —

Classes
— No classes —

Derived groups
— No child groups —

Daily run status

Number and status of runs during the last 30 days:

Date	Runs
2012-03-08	350
2012-03-09	350
2012-03-10	350
2012-03-11	320
2012-03-12	210
2012-03-13	380
2012-03-14	350
2012-03-15	330
2012-03-16	250

Nodes for this group

Hostname	Source	Latest report
✓ dcsrm02.usatlas.bnl.gov	dcsrm02.usatlas.bnl.gov	2012-03-16 16:40 EDT
✓ ddcap03.usatlas.bnl.gov	ddcap03.usatlas.bnl.gov	2012-03-16 16:40 EDT
✓ dcsrddb01.usatlas.bnl.gov	dcsrddb01.usatlas.bnl.gov	2012-03-16 16:40 EDT
✓ dcdoor11.usatlas.bnl.gov	dcdoor11.usatlas.bnl.gov	2012-03-16 16:31 EDT
✓ dcdoor12.usatlas.bnl.gov	dcdoor12.usatlas.bnl.gov	2012-03-16 16:22 EDT
✓ dcdoor08.usatlas.bnl.gov	dcdoor08.usatlas.bnl.gov	2012-03-16 16:22 EDT
✓ chimera01.usatlas.bnl.gov	chimera01.usatlas.bnl.gov	2012-03-16 16:21 EDT
✓ dcdoor13.usatlas.bnl.gov	dcdoor13.usatlas.bnl.gov	2012-03-16 16:21 EDT
✓ dcdoor06.usatlas.bnl.gov	dcdoor06.usatlas.bnl.gov	2012-03-16 16:20 EDT
✓ dcdoor14.usatlas.bnl.gov	dcdoor14.usatlas.bnl.gov	2012-03-16 16:20 EDT

Puppet Config & Scalability

- Still using 2.6.14 on RHEL5 with ruby 1.8.5
 - testing upgrade to 2.7 on RHEL6 with ruby 1.8.7
- Apache with Phusion Passenger (mod_rails)
- Queue daemon with activemq for fast DB updates of storeconfigs
- Over 2k agents currently using puppet
- Noticed MySQL errors with inventory service enabled with a rate of about 1 client/second
- Will look at Tomcat/JRuby later since early test show promise, but it still has problems

Future Plans

- Change Management
 - Policy & procedures used to control changes made to production systems (ITIL, DevOps).
 - Changes made only during official windows.
 - Absolutely no unauthorized changes, no “cowboy” type behavior tolerated.
 - Use testbed environment to test changes before putting them into production.
 - Create replica of prod using VMs for auto-tests
 - Tools like Puppet, Git & GLPI can help make changes and keep a historical change record.

Why do it?

- Uncontrolled change can work sometimes, but often cause self inflicted problems and future firefighting episodes & upgrade nightmares.
- Stop duplicating work and effort, standardize.
- Stop making time consuming manual changes.
- Without it, servers become like snowflakes: they may all start out identical, but over time, config drift eventually makes each one unique.

Benefits

- Shift staff time from perpetual reactive firefighting mode, that often only addresses the symptoms, to more proactive work, that addresses the root causes of problems (fire prevention).
- Repeatable and standard build & config process means it is often faster and easier to rebuild problematic servers, rather than waste hours or days troubleshooting problems.