



Cluster Management

Jason Allen
OSG All Hands Meeting, Boston
March 8, 2011

About Me

- Fermilab for 10 years
- Head of Fermilab Experiments Facilities (FEF) Department in the Computing Division.
- FEF manages approx 3500 physical Linux systems: farms, workstations, online servers, offline servers, and more!
- FEF is also responsible for the development and support of Scientific Linux. SL 6 now available!

For this talk


- Primarily discuss Free and Open Source Software (FOSS) tools
- Tools in production at Fermilab or ones that we have recently evaluated
- “Cluster” defined as any large group of Linux systems
- Skimming the surface of cluster management

Cluster Management Basics

- Management infrastructure
- Provisioning
- Configuration and software package management
- Monitoring


Management Infrastructure

- Remote power-cycling and serial console access
- FEF has standardized on Avocent ACS serial console servers and Avocent PM line of PDU
- Other depts use IPMI or a mix of Avocent and APC products
- All depts have scripts to control and configure remote power-cycling and serial console access




CMS Tier 1 Power Usage Plots

CMS plots power utilization by querying PDUs using SNMP. This data can be particularly useful to datacenter managers.




fcc-2-1380

```
cmssrv03 cmssrv02 cmssrv01 cmssrv95 cmsosgce  
cmsosgce2 cmsosgce3 cmsosgce4 cmssrv105 cmssrv14  
cmssrv34 cmssrv35
```



fcc-2-1381

```
cmsstor252 cmsstor253 cmsstor254 cmsstor255  
cmsstor256 cmsstor257 cmsstor258 cmsstor259  
cmsstor260 cmsdisk252 cmsdisk253 cmsdisk254  
cmsdisk255 cmsdisk256 cmsdisk257 cmsdisk258  
cmsdisk259 cmsdisk260
```



fcc-2-1382

```
cmsstor155 cmsstor154 cmsstor153 cmsstor150  
cmsstor151 cmsstor149 cmsdisk155 cmsdisk154  
cmsdisk153 cmsdisk151 cmsdisk149 cmsdisk150
```

Provisioning Tools

Preparing a system for use; OS installation and initial configuration.

Provisioning Tools

At Fermilab

- PXE/Kickstart
- Rocks



FEF's PXE/Kickstart Setup

- FEF uses custom tool built on top of MySQL and Perl DHCP server modules.
- No dhcpd restart required.
- Web front-end for specifying kickstart / node combinations.
- Very flexible
- Kickstart files are created dynamically based on selections from the web GUI
- Planning to eval Cobbler later this year.

Rocks Clusters

- Open source Linux distribution based on CentOS.
- Created in 2000
- Created for easy deployment of large clusters.
- Used by CMS Tier 1 at Fermilab for 2400 machines
- CMS can install ~500 systems in 1 hour.
- Only one OS version per Rocks server may be a deal breaker for some.

Configuration and Package Management

Tools that help manage system configuration (files, dirs, permissions, etc.) and software packages.

Configuration and Package Management Tools

Popular open source solutions:

- Cfengine
- Puppet
- Bcfg2

Cfengine



- First version released in 1993
- Still the de facto standard configuration management tool, though Puppet is quickly catching-up
- Written in C
- Fairly easy to understand syntax
- Used by D0 workstation cluster, ClueD0
- FEF used for many years
- Relatively easy to find sysadmins with experience



Puppet

- New generation of configuration management system
- Extensible, declarative language
- Understands dependencies (huge benefit)
- Better reporting than Cfengine
- Auto generation of documentation (think Javadoc)
- FEF avg is 325 actions per node every Puppet run

FEF Puppet Usage

- Management of all external mounts
- Kerberos files -- keytab files, .k5login, etc
- Package management (RPM sets grouped by cluster)
- NIC bonding configs
- Group quotas
- Grid host certs
- FEF_backup

Cfengine vs. Puppet (High Level)

	Native File Editing	Dependency Management	Commercial Support	Dependency Graphs	Scalability
Cfengine	✓		✓		✓
Puppet		✓	✓	✓	✓

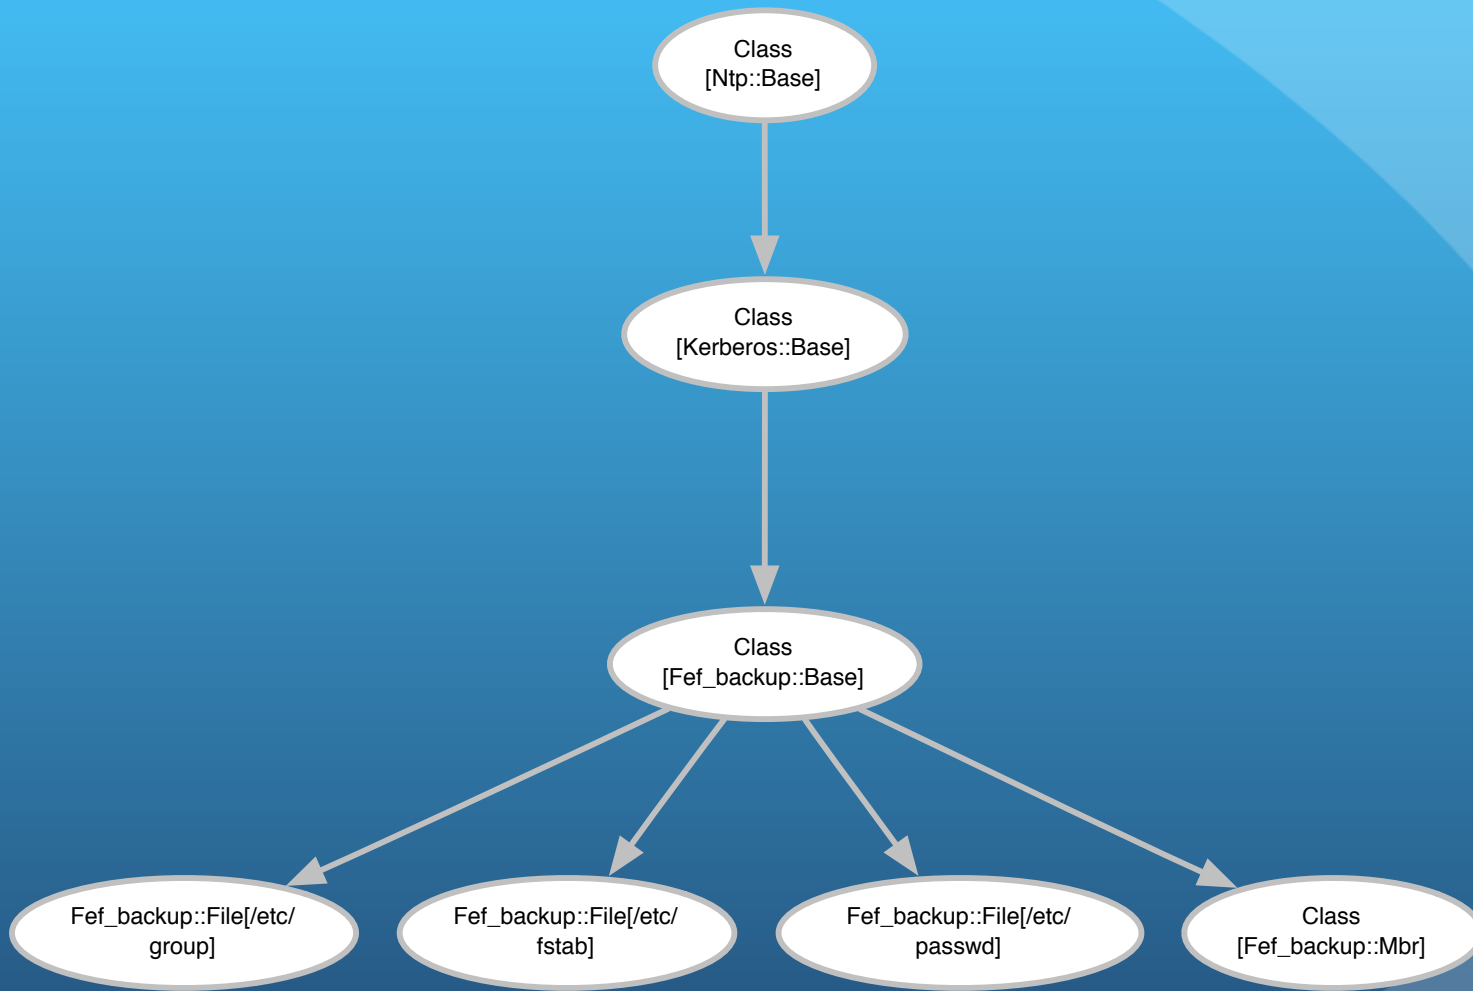
Puppet Add User Example

```
User { managehome => true,  
      ensure      => present,  
      gid         => users,  
      shell       => "/bin/bash",  
    }  
  
user { "mark":  
      uid => 1000,  
    }  
  
user { "fred":  
      uid => 1001,  
    }  
  
user { "jane":  
      uid => 1002,  
    }
```

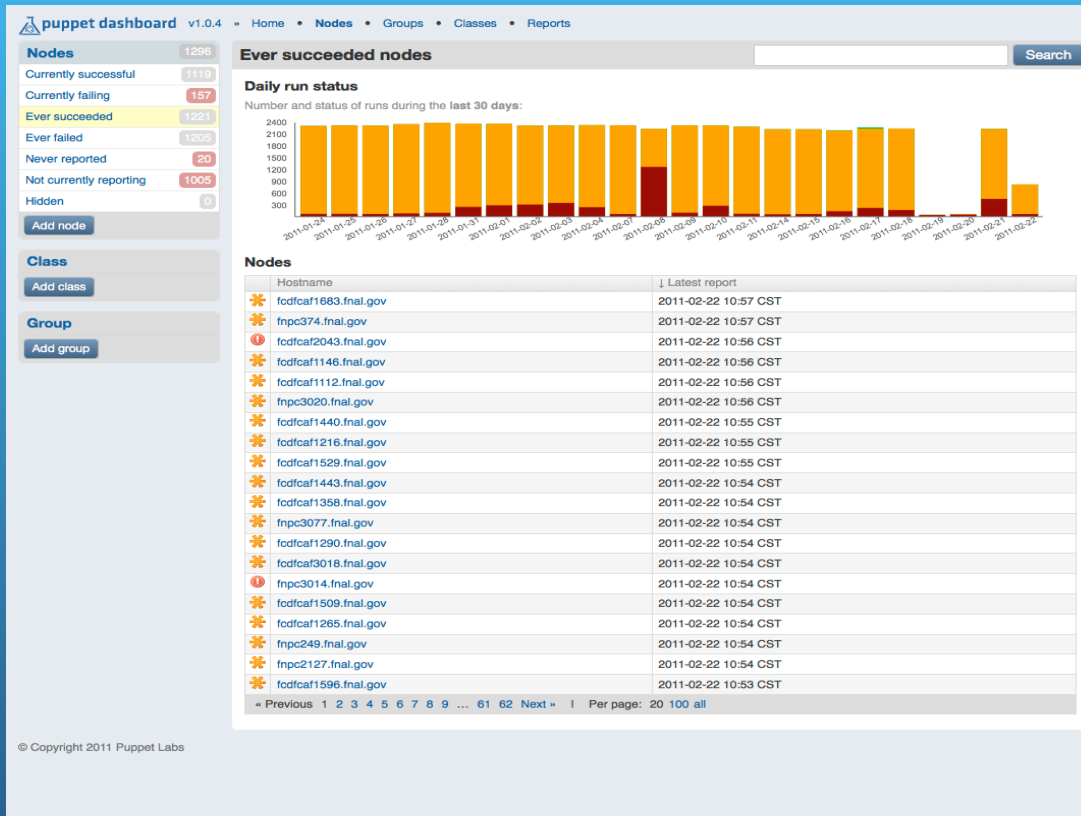
Cfengine Add User Example

```
"pw[mark]" string => "mark:x:1000:100:Mark Burgess:/home/  
mark:/bin/bash";  
  
"pw[fred]" string => "fred:x:1001:100:Right Said:/home/  
fred:/bin/bash";  
  
"pw[jane]" string => "jane:x:1002:100:Jane Doe:/home/  
jane:/bin/bash";  
  
"users" slist => getindices("pw");  
  
files:  
  "/etc/passwd"  
    edit_line => append_users_starting("addusers.pw");  
  "/etc/group"  
    edit_line => append_user_field("root","4","@  
(addusers.users)");  
  "/home/${users}/."  
    create => "true",  
    perms => mog("755","${users}","users");
```

Example Puppet Dependency Graph



Puppet Dashboard



Puppet dashboard is a web interface for quickly viewing puppet run status and the state of individual system configurations.

Bcfg2



- BCFG2 is an xml-based configuration management system
- Developed by Argonne National Lab
- Being used by CMS Tier 1 at Fermilab for 2 years to manage a limited number of configuration items. Not RPMS
- Developers are very responsive; provide support via mailing list and IRC
- Complex file manipulation can be tricky
- Does simple pre/post dependencies

Bcfg2 Reporting System

Bcfg2 Reporting System

- Home
- Clients
 - Detailed List
- Displays
 - System
 - Summary
 - Timing
- Config Items
 - Bad
 - Modified

Detailed Client List

Enter date or use calendar popup: @

Node	State	Good	Bad	Modified	Extra	Last Run	Server
cms-sleepgw.fnal.gov	clean	7	0	0	798	2011-03-03 06:55	cmssrv01.fnal.gov
cms-xen20.fnal.gov	dirty	4	1	2	821	2011-02-28 06:30	cmssrv01.fnal.gov
cms-xen39.fnal.gov	clean	9	0	0	664	2011-03-02 08:35	cmssrv01.fnal.gov
cms-xen40.fnal.gov	clean	9	0	0	638	2011-03-03 05:03	cmssrv01.fnal.gov
cms-xen41.fnal.gov	clean	9	0	0	552	2011-03-03 05:06	cmssrv01.fnal.gov
cmscbn.fnal.gov	clean	7	0	0	803	2010-12-14 04:23	cmssrv01.fnal.gov
cmscode01.fnal.gov	clean	7	0	0	871	2011-03-03 04:36	cmssrv01.fnal.gov
cmscode02.fnal.gov	clean	7	0	0	826	2011-03-03 04:26	cmssrv01.fnal.gov
cmscode03.fnal.gov	clean	7	0	0	824	2011-03-02 05:35	cmssrv01.fnal.gov
cmsdcd1.fnal.gov	clean	10	0	0	886	2011-03-03 04:02	cmssrv01.fnal.gov
cmsdcd2.fnal.gov	clean	10	0	0	888	2011-03-02 04:32	cmssrv01.fnal.gov
cmsdcm1.fnal.gov	clean	12	0	0	898	2011-03-02 06:50	cmssrv01.fnal.gov
cmsfilemover.fnal.gov	clean	7	0	0	957	2011-03-03 07:22	cmssrv01.fnal.gov
cmsfnal01.fnal.gov	clean	11	0	0	936	2011-03-03 06:48	cmssrv01.fnal.gov
cmsfnal02.fnal.gov	clean	11	0	0	936	2011-03-03 05:06	cmssrv01.fnal.gov
cmsfnal03.fnal.gov	clean	11	0	0	977	2011-03-03 06:23	cmssrv01.fnal.gov
cmsfts1.fnal.gov	dirty	2	5	0	1023	2010-12-14 05:11	cmssrv01.fnal.gov
cmsfts2.fnal.gov	clean	7	0	0	1018	2010-12-11 04:32	cmssrv01.fnal.gov
cmsfts3.fnal.gov	clean	7	0	0	792	2011-03-01 04:57	cmssrv01.fnal.gov
cmslmdc.fnal.gov	clean	7	0	0	872	2011-03-03 06:00	cmssrv01.fnal.gov

Node status and last run times are viewable from the Bcfg2 web interface.

Monitoring

Tools to help monitor system state and performance.

Monitoring

In use at Fermilab:

- Zabbix
- Nagios
- Ganglia
- Many custom solutions using MRTG, RRDtool, etc

Zabbix

- Being used by CMS Tier 1 to monitor approx 2.4K nodes; performs 100K checks.
- Does status and performance monitoring.
- Relatively new compared to Nagios.
- Most configuration is done via the web interface.
- Easy to add custom checks and alerts.

Zabbix Dashboard

Zabbix provides a polished web interface that displays finely grained status and performance information.

The screenshot displays the Zabbix web interface. At the top, there is a navigation menu with options like Monitoring, Inventory, Reports, Configuration, and Administration. Below this, a breadcrumb trail shows the current path: History > Search > Hosts > Status of triggers > QUEUE > Status of triggers. The main content area is divided into several sections:

- PERSONAL DASHBOARD:** Includes links for Favourite Graphs, Favourite Screens, and Favourite Maps.
- Status of Zabbix:** A table showing key parameters and their values.
- System status:** A table showing the status of various host groups.
- Host status:** A table showing the status of individual hosts.

Parameter	Value	Details
Zabbix server is running	Yes	-
Number of hosts (monitored/not monitored/templates)	636	612 / 6 / 18
Number of items (monitored/disabled/not supported)	52192	41731 / 7745 / 2716
Number of triggers (enabled/disabled)[true/unknown/false]	26048	23655 / 2393 [23 / 6892 / 16739]
Number of users (online)	17	4
Required server performance, new values per second	473.0381	-

Host group	Disaster	High	Average	Warning	Information	Not classified
dcache_all	0	10	0	0	1	0
Disk arrays	0	0	0	3	0	0
external	0	0	0	0	0	0
Linux servers	0	9	0	1	0	0
LPC interactive	0	0	0	0	0	0
Misc	0	0	0	0	0	0
Templates	0	0	0	0	0	0
ZABBIX Servers	0	1	0	0	4	0

Host group	Without problems	With problems	Total
dcache_all	243	6	249
Disk arrays	181	2	183
external	2	0	2
Linux servers	93	5	98
LPC interactive	25	0	25
Misc	23	0	23
Templates	1	0	1
ZABBIX Servers	0	1	1

Nagios

- Used by FEF; 3.5K nodes, approx 30K checks on one server
- Around for many years.
- Create a new check by dropping shell script on node (check_mk plugin)
- Nagios support built-in to Puppet.
- Web interface can be slow and feels dated.

Nagios Web Interface

The Nagios web interface is functional but feels dated. Performance is an issue when monitoring many hosts.

The screenshot displays the Nagios web interface. On the left is a navigation menu with sections: General, Current Status, Reports, and System. The main content area is titled 'Current Network Status' and includes a summary table for Host Status Totals and Service Status Totals. Below these are links for service status details and a table titled 'Host Status Details For All Host Groups' listing individual hosts with their status, last check time, duration, and status information.

Current Network Status
Last Updated: Tue Feb 22 14:34:36 CST 2011
Updated every 300 seconds
Nagios® Core™ 3.2.0 - www.nagios.org
Logged in as Jason M. Allen

Host Status Totals

Up	Down	Unreachable	Pending
2563	14	0	0

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
33671	0	413	44	0

Host Status Details For All Host Groups

Host	Status	Last Check	Duration	Status Information
131.225.15.13	UP	02-22-2011 14:31:44	91d 5h 23m 21s	PING OK - Packet loss = 0%, RTA = 0.42 ms
131.225.23.46	UP	02-22-2011 14:31:44	91d 5h 23m 21s	PING OK - Packet loss = 0%, RTA = 0.51 ms
dfvmadmin01	UP	02-22-2011 14:31:24	91d 5h 23m 20s	PING OK - Packet loss = 0%, RTA = 0.18 ms
dfvmv01	UP	02-22-2011 14:31:44	91d 5h 20m 49s	PING OK - Packet loss = 0%, RTA = 1.25 ms
dfvmv02	UP	02-22-2011 14:31:44	91d 5h 23m 21s	PING OK - Packet loss = 0%, RTA = 0.84 ms
dfvmv03	UP	02-22-2011 14:31:24	47d 1h 9m 24s	PING OK - Packet loss = 0%, RTA = 1.18 ms
dfvmv04	UP	02-22-2011 14:31:44	47d 0h 19m 54s	PING OK - Packet loss = 0%, RTA = 0.15 ms
dfcabam1	UP	02-22-2011 14:32:14	13d 17h 49m 58s	PING OK - Packet loss = 0%, RTA = 0.36 ms
dfcabav1	UP	02-22-2011 14:32:14	13d 17h 49m 29s	PING OK - Packet loss = 0%, RTA = 0.50 ms
dfcabav2	UP	02-22-2011 14:32:14	13d 17h 49m 31s	PING OK - Packet loss = 0%, RTA = 0.49 ms
dfca0841	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 10.12 ms
dfca0842	UP	02-22-2011 14:31:24	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 10.93 ms
dfca0843	UP	02-22-2011 14:32:05	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 7.71 ms
dfca0844	UP	02-22-2011 14:29:44	14d 1h 42m 3s	PING OK - Packet loss = 0%, RTA = 10.37 ms
dfca0845	UP	02-22-2011 14:31:24	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 10.12 ms
dfca0846	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 10.40 ms
dfca0847	UP	02-22-2011 14:31:24	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 10.22 ms
dfca0848	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 9.55 ms
dfca0849	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 9.67 ms
dfca0850	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 9.30 ms
dfca0851	UP	02-22-2011 14:31:24	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 10.16 ms
dfca0852	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 9.24 ms
dfca0853	UP	02-22-2011 14:31:24	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 10.24 ms
dfca0854	UP	02-22-2011 14:31:24	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 9.65 ms
dfca0855	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 9.71 ms
dfca0856	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 9.49 ms
dfca0857	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 9.17 ms
dfca0858	UP	02-22-2011 14:31:44	14d 1h 40m 43s	PING OK - Packet loss = 0%, RTA = 10.07 ms
dfca0859	UP	02-22-2011 14:31:24	14d 1h 42m 3s	PING OK - Packet loss = 0%, RTA = 9.95 ms

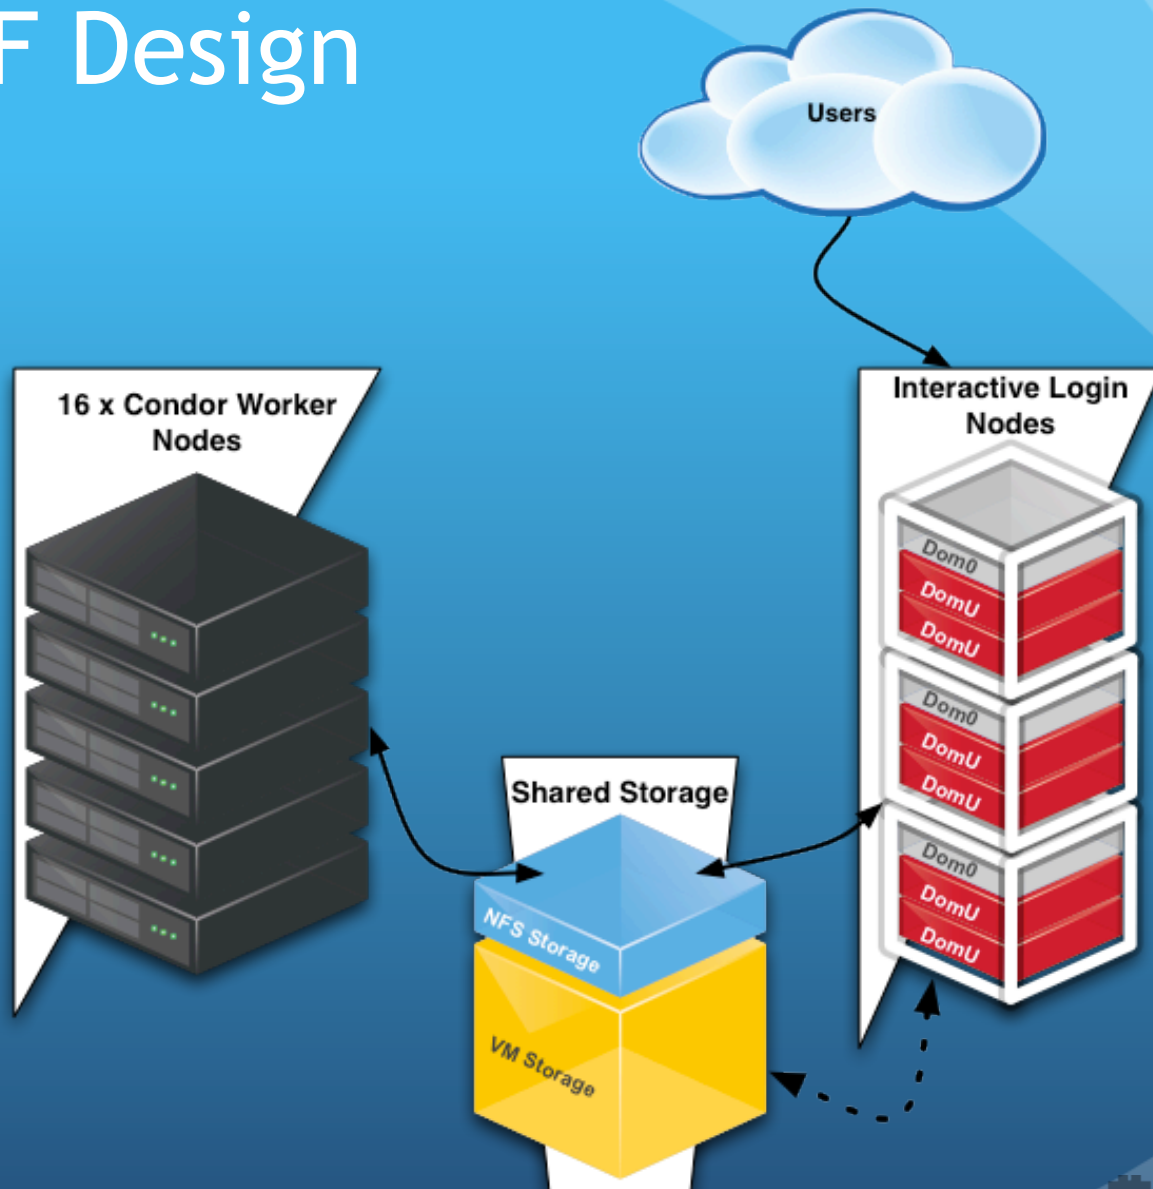
Misc Cluster Management Tools and Techniques

Other things you might find interesting

GPCF

- General Purpose Computing Facility
- Oracle VM cluster + batch cluster
- Quick provisioning for smaller or new experiments.
- Computing resources in days instead of weeks/months.
- Small batch cluster serves as training wheels for the Grid.

GPCF Design

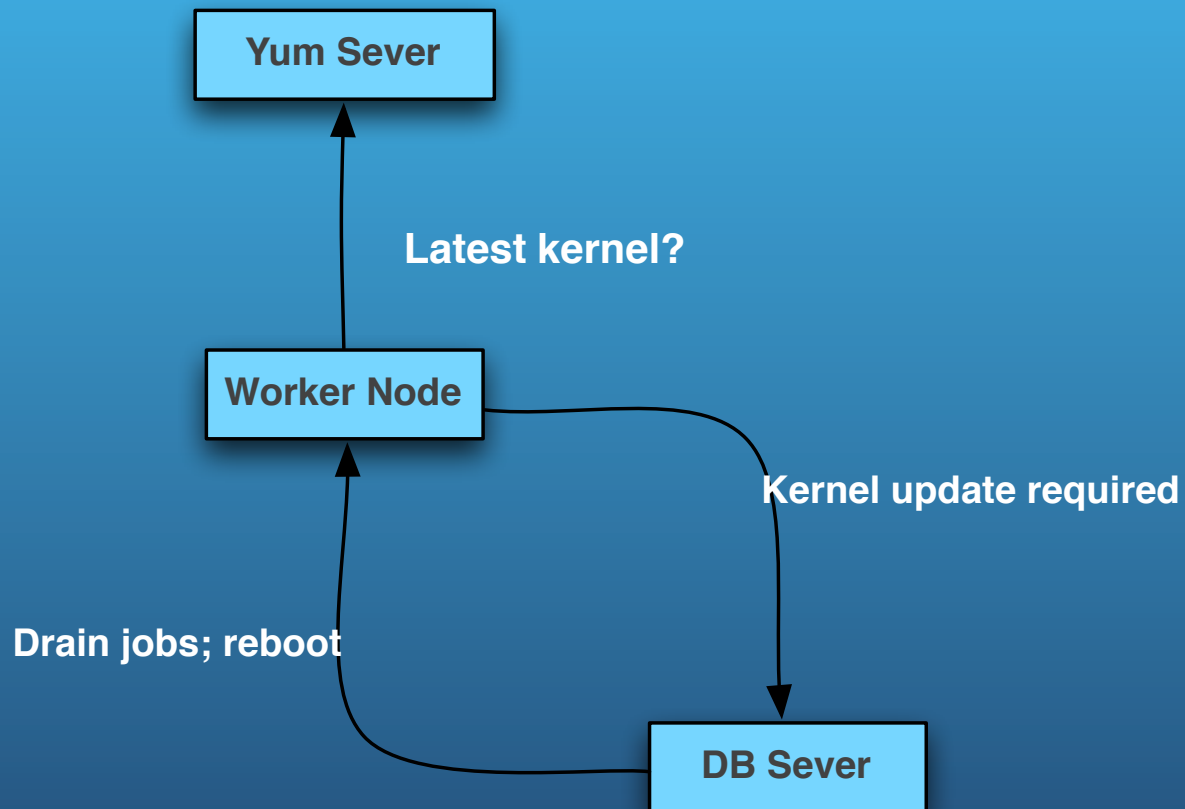


CDF Grid Rolling Upgrades

- Fermi security policy states that software patches to be applied within 60 days of release.
- Systems must be every rebooted every 2-3 months to load new Linux kernel.
- Tool automatically drains and reboots 50 nodes at a time.
- 1K node cluster is updated in approx 2 weeks.

Rolling Upgrades

Conceptual Overview



Resources

Comparison of Puppet/Cfengine/Bcfg2

<https://cd-docdb.fnal.gov:440/cgi-bin/ShowDocument?docid=3967>

Evaluation and feature comparison of the Nagios and Zabbix monitoring systems

<http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=3277>

Download Scientific Linux 6

<http://scientificlinux.org/>