

Cfengine @ JLab

David J. Bianco
bianco@jlab.org



Table of Contents

- Introduction
- JLab's Environment
- What is cfengine?
- JLab's cfengine architecture
- JLab templates
- Summary
- Questions



JLab's Unix Environment

- ~90 central computing Unix servers (Linux, Solaris, HP)
 - General computing resources, web, email, etc.
- ~50 CAD nodes (HP)
- ~185 compute farm nodes (Linux)
- A large number of user-managed Unix workstations (mostly Linux)



JLab's Unix Environment

- The lab's Unix admin staff is just 6 people.
- Changes are made to these machines all the time
- As with any environment, proper communication & documentation can be a problem
- Once a problem is fixed... will it remain fixed?
- Several recent incidents have underscored the need for proper configuration management
- In January 2002, JLab started looking into cfengine to help solve these problems



What is cfengine?

- Stands for “Configuration Engine”
- Policy driven configuration management for a network of machines
- Open source
- Unix & NT/2000
 - Mostly Unix, though



What is cfengine?

- Developed by Mark Burgess @ Oslo University College in 1993
- Used on an estimated 100,000 nodes worldwide
- Currently in version 2.0

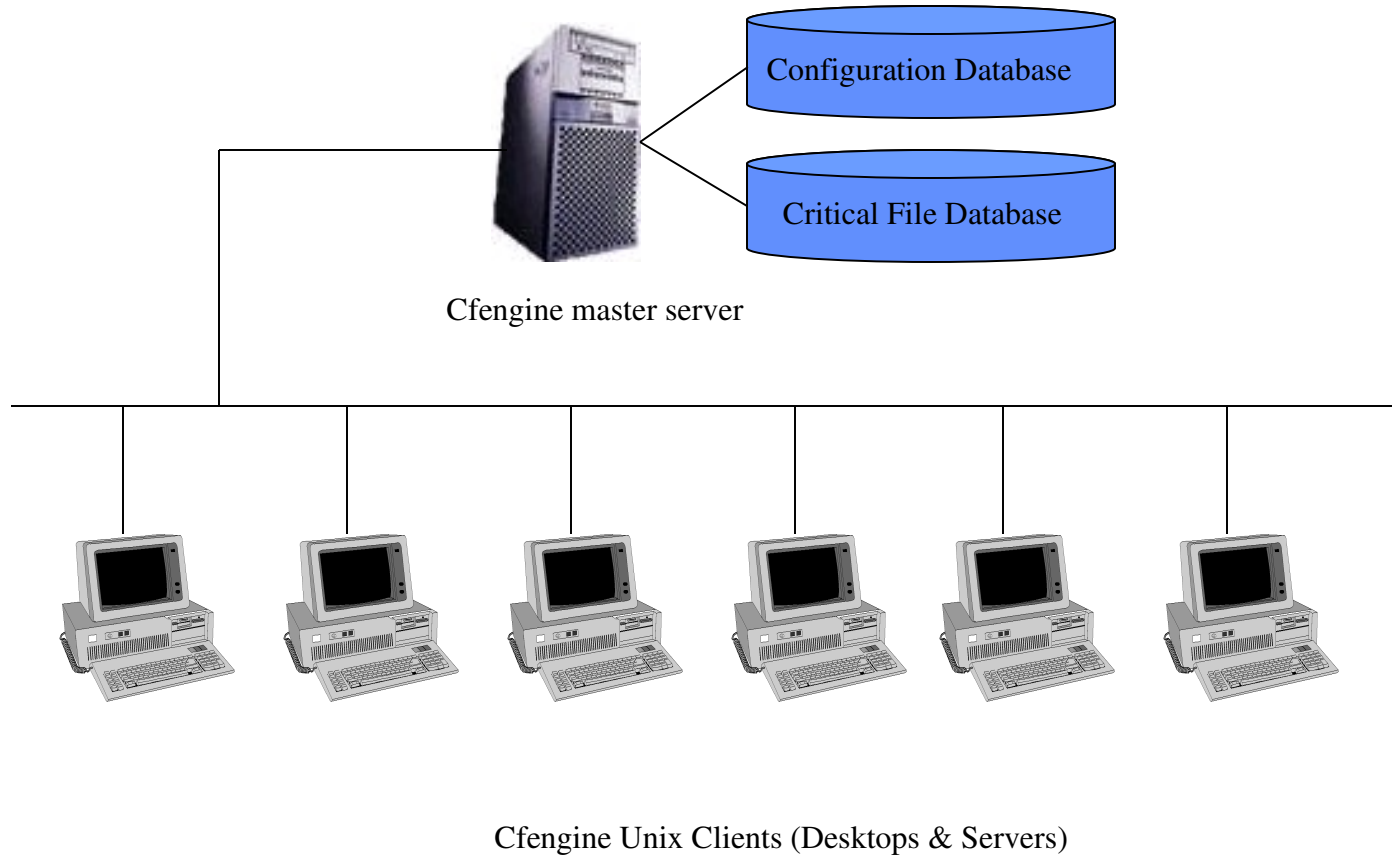


What is cfengine?

- Three main parts
 - cfagent
 - Network services
 - Declarative configuration templates
- Optional anomaly detection service



JLab's cfengine Architecture



JLab's cfengine Architecture

- Cfengine master server contains
 - Cfengine binaries for all platforms
 - 'All configuration templates
 - Master copies of critical system/software configuration files
- Cfengine clients contain
 - Local copies of their own binaries
 - A complete copy of the configuration templates



JLab's cfengine Architecture

- Clients use crontab to run “*cfexecd -F*” every 30 minutes
 - Wrapper to run cfagent and email any output to the system administrator
 - “splay time” keeps all client from overloading the master at once
 - Cfagent automatically copies updated binaries and config templates from master
 - Most configuration checks are performed during each run
 - Expensive checks (file sweeps) performed only during the midnight run



JLab's cfengine Architecture

- Administrator can also run cfengine manually
 - Local root user (on a single client): *cfagent* (local root user)
 - Cfengine admin (remotely from the master): *cfrun*



Installing cfengine on a host

- Run `/local/cfengine/bin/cfinstall <hostname>` as root
- Log is `/tmp/cfinstall-<hostname>`

```
Starting cfengine installation for sysdevs1 @ Tue Mar  5 09:08:42 EST 2002
Installation host is: SunOS
Generating keypair... DONE
Exchanging keypairs...
Running cfagent for the first time...
cfengine:sysdevs1: Update of image /home/janed/.ssh/authorized_keys from master
/local/cfengine/REPOSITORY/common/home/janed/.ssh/authorized_keys on cfm.jlab.org
[Additional config output]
```



Summary

- JLab uses cfengine 2.0 to manage configuration on a network of hundreds of Unix hosts
- The configuration master contains full copies of all configuration binaries, templates and important system files
- All network connections are encrypted and mutually authenticated
- The template files are modular, enabling us to pick and choose among the pieces we run for a particular host



Questions?

David J. Bianco
<bianco@jlab.org>

