

# Deploying LCG/gLite middleware with cfengine

Colin Morey (Manchester University)

# What is cfengine?

Colin Morey (Manchester University)

# What is cfengine?

- Highly scalable
- Autonomous configuration agent
- Simple configuration language
- ‘Loose’ client-server model

# What it won't do

- Install your cluster
  - But you can integrate it with Kickstart
- Replace your package manager
  - But it can augment it...
- Make you coffee/tea

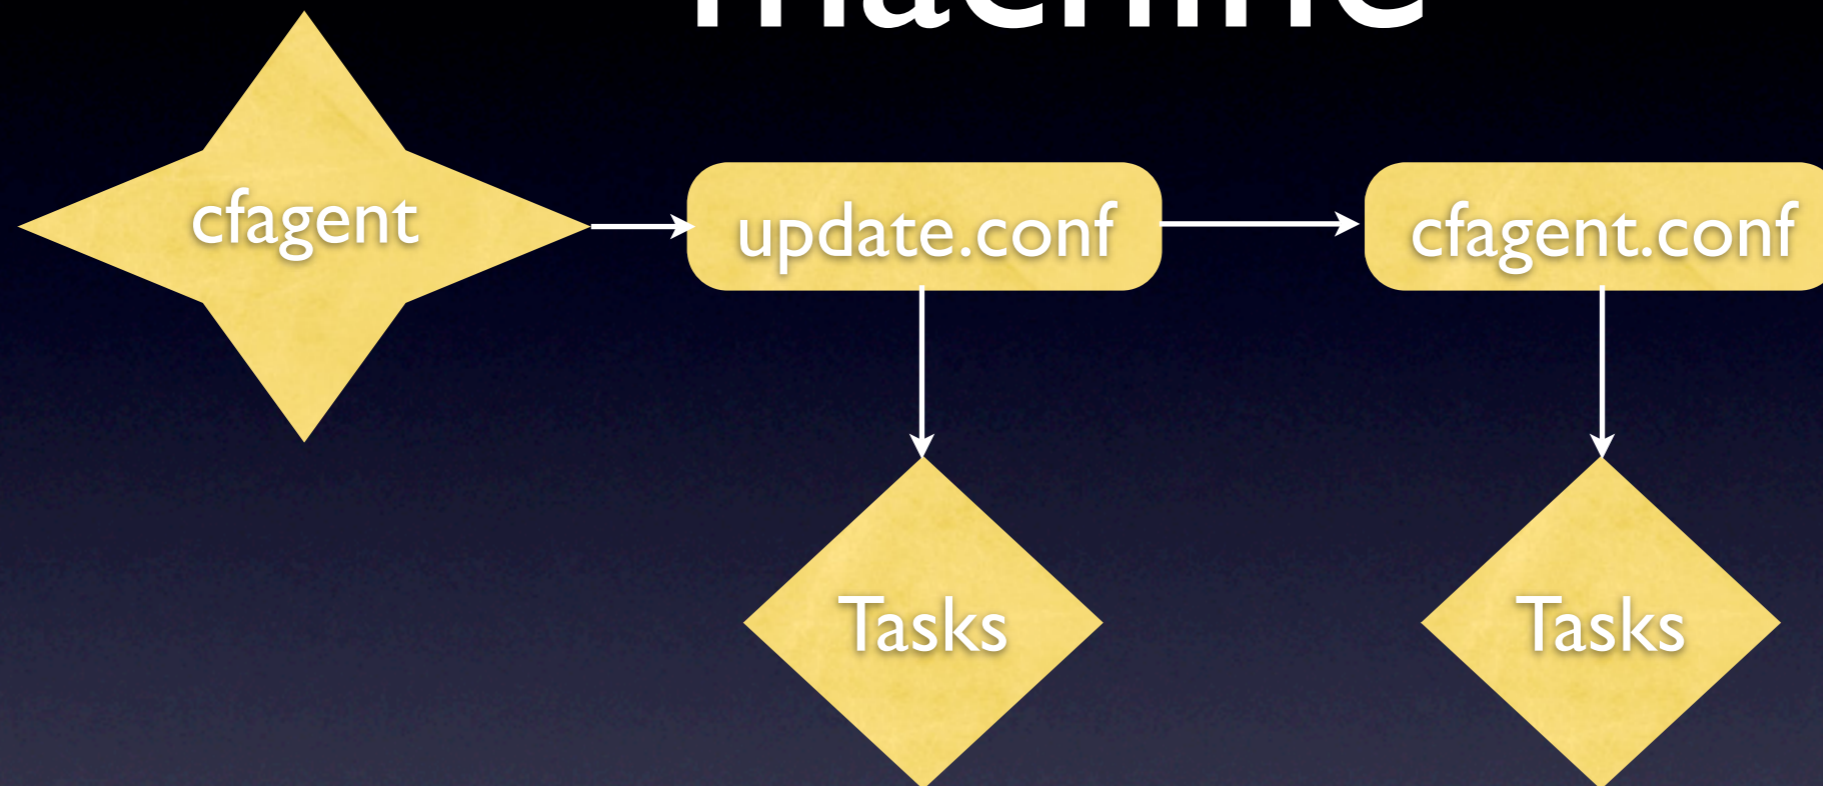
# What it will do

- Keep any configuration files how you want them
- Maintain permissions on files and directories
- Keep a track of file checksums
- And more...


# How does cfengine work?

Colin Morey (Manchester University)

# The Basics - Local machine



# Adding the file daemon, cfservd



Colin Morey (Manchester University)



# Anatomy of (basic) cfagent.conf

classes:

```
bohr_nodes = HostRange(bohr,0-9999)
```

import:

```
any::
```

```
bohr_nodes::  
    cf.bohr
```

control:

```
any::
```

```
    actionsequence = ( files links )
```

```
    policyhost = ( ks1.tier2.hep.manchester.ac.uk )
```

```
doitnow::
```

```
    ifelapsed = ( 0 )
```

files:

```
/etc/passwd mode=644 owner=root action=fixall
```

links:

```
/exp_software -> /nfs
```

# Classes - the Backbone

Default cfengine classes look something like this:-

```
Defined Classes = ( 195_194_109 195_194_109_1 32_bit Day24 Hr20 Hr20_Q1 January Min05_10 Min09 Q1 Wednesday Yr2007 ac_uk any bohr01
bohr01_test_tier2_hep_manchester_ac_uk bohr_nodes cfengine_2 cfengine_2_1 cfengine_2_1_21 compiled_on_linux_gnu dteam_home_dir
entropy_cfengine_in_low entropy_cfengine_out_low entropy_dns_in_low entropy_dns_out_low entropy_ftp_in_low entropy_ftp_out_low
entropy_icmp_in_low entropy_icmp_out_low entropy_irc_in_low entropy_irc_out_low entropy_misc_in_low entropy_misc_out_low
entropy_netbiosdgm_in_low entropy_netbiosdgm_out_low entropy_netbiosns_in_low entropy_netbiosns_out_low entropy_netbiossn_in_low
entropy_netbiossn_out_low entropy_nfsd_in_low entropy_nfsd_out_low entropy_smtp_in_low entropy_smtp_out_low entropy_ssh_out_low
entropy_tcpack_in_low entropy_tcpack_out_low entropy_tcpfin_in_low entropy_tcpfin_out_low entropy_tcpsyn_in_low entropy_tcpsyn_out_low
entropy_udp_in_low entropy_udp_out_low entropy_www_in_low entropy_www_out_low entropy_wwws_in_low entropy_wwws_out_low
hep_manchester_ac_uk i686 ipv4_195 ipv4_195_194 ipv4_195_194_109 ipv4_195_194_109_1 linux linux_2_4_21_40_ELSmp linux_i686
linux_i686_2_4_21_40_ELSmp linux_i686_2_4_21_40_ELSmp__1 SMP_Wed_Mar_15_12_28_15_CST_2006 loadavg_high_ldt manchester_ac_uk
net_iface_eth0 net_iface_lo pid_file_exists redhat rootprocs_high_normal scientific scientific_sl scientific_sl_3 scientific_sl_3_0
test_tier2_hep_manchester_ac_uk tier2_hep_manchester_ac_uk uk update_cert update_ssh use_ksl )
```

# Hard Classes

- OS - Linux/IRIX/Solaris
- OS Release - eg. scientific scientific\_sl scientific\_sl\_3 scientific\_sl\_3\_0
- Domain name
- IP Address

# Soft Classes - envd

- System Load
- Network sockets
- Memory Usage
- Abnormal number of sockets open to a daemon

# Soft Classes - User

- Bohr
- Update\_cert
- Cluster\_01
- doitnow

# Some Examples

Colin Morey (Manchester University)

# Monitoring and Restarting pbs\_mom

processes:

bohr\_nodes::

```
"/usr/sbin/pbs_mom" restart "/etc/init.d/pbs_mom"
```

# Building a Centrally Distributed grid-mapfile

copy:

```
/etc/grid-security/grid-mapfile      dest=/var/cfengine/masterfiles/inputs/all-nodes/grid-  
mapfile
```

```
mode=0440 o=root g=root define=new_gridmapfile
```

shellcommands:

```
ks | .Min10_15::
```

```
"/opt/edg/sbin/edg-mkgridmap --output=/etc/grid-security/grid-mapfile --safe >> /var/  
log/edg-mkgridmap.log 2>&1"
```

```
ks | .new_gridmapfile::
```

```
"/bin/sh -c 'cd /var/cfengine/masterfiles/inputs/all-nodes; CVS_RSH=/usr/bin/ssh /usr/  
bin/cvs commit -m\"New autogenerated grid-mapfile\" grid-mapfile' "
```

alerts:

```
ks | .new_gridmapfile::
```

```
"A new grid-mapfile has been copied into /var/cfengine/masterfiles/inputs/all-nodes"
```



# Monitor CE<=>WN ssh connectivity

control:

```
ssh_to_ce01_test = ( ExecResult("/usr/bin/su - NONROOTUSER -c {quote}/usr/bin/ssh  
CEHOSTNAME /bin/hostname{quote}") )
```

classes:

```
ssh_to_ce01 = ( Regcmp(".*CEHOSTNAME.*",{ssh_to_ce01_test} ) )
```

alerts:

```
!ssh_to_ce01::
```

```
"Unable to ssh to CEHOSTNAME! Error was :- $(ssh_to_ce01_test)"
```

# Checking local host certificate

control:

```
AddInstallable = ( cert_date_error )
```

copy:

```
{master_scripts}/check-cert-date.sh dest=/usr/local/bin/check-cert-date.sh  
mode=0550 o=root g=root server={policyhost} backup=true
```

shellcommands:

```
"/usr/local/bin/check-cert-date.sh" umask=022 elsedefine=cert_date_error
```

alerts:

```
cert_date_error::  
"Problem with Host Certificate"
```

# Watching disk usage

control:

```
threshold = ( 20 )  
keep = ( 7 )  
AddInstallable = ( scratch_needs_cleaning scratch_tidied )  
actionsequence = ( disks tidy )
```

disks:

```
/scratch freespace=${threshold}%  
inform=true force=true define=scratch_needs_cleaning
```

tidy:

```
scratch_needs_cleaning::  
/scratch pattern=* recurse=inf age=${keep} rmdirs=true define=scratch_tidied
```

alerts:

```
scratch_tidied::  
"Files Deleted from /scratch"
```

# Simple Tripwire - File Checksums

files:

Hr04::

/usr

owner=root,daemon # all files must be owned by root or daemon

checksum=md5

recurse=inf

ignore=tmp

ignore=nagios

action=fixall

# Large Scale cfengine

strategies:

```
{ Simple_Alternate
  use_ks1: "100"
  use_ks2: "100"
}
```

control:

any:

```
master_cfinput = ( /var/cfengine/masterfiles/inputs )
```

use\_ks1::

```
policyhost = ( ks1.tier2.hep.manchester.ac.uk )
```

use\_ks2::

```
policyhost = ( ks2.tier2.hep.manchester.ac.uk )
```

copy:

```
${master_cfinput}/all-nodes/at.allow dest=/etc/at.allow
```

```
mode=0644 o=root g=root server=${policyhost}
```

# dcache - upgrade from 1.6.6 to 1.7.0

```
#####
# Everything needed to upgrade a dcache install from 1.6 to 1.7
#####
# Note:At Manchester, the dcache class refers to the head nodes(s)

control:
  do_dcache_upgrade::
    actionsequence = ( editfiles.dcache_upgrade packages.dcache_upgrade shellcommands.stop_dcache copy.dcache_upgrade shellcommands.upgrade_dcache_database shellcommands.update_pkg
    copy.dcache_upgrade shellcommands.dcache_upgrade )

editfiles:
  dcache_upgrade::
  {
    /etc/yum.conf
      BeginGroupIfNoLineMatching "[local-dcache]"
      Append "[local-dcache]"
      Append "name=Local Mirror of dcache.org Stable"
      Append "baseurl=http://ks.tier2.hep.manchester.ac.uk/linux-install/yum/dcache/"
      EndGroup
    }
  {
    /opt/glite/yaim/etc/site-info.def
      DeleteLinesContaining "RESET_DCACHE_CONFIGURATION=no"
      AppendIfNoSuchLine "RESET_DCACHE_CONFIGURATION=yes"
    }
}

copy:
  dcache_upgrade.dcache::
    ${master_scripts}/upgrade-dcache-database.sh dest=/usr/local/bin/upgrade-dcache-database.sh
    mode=0550 o=root g=root server=${policyhost} backup=true

packages:
  dcache_upgrade::
    dcache-server action=install version=1.7.0-34
    dcache-client action=install version=1.7.0-34

shellcommands:
  stop_dcache::
    "/etc/init.d/dcache-pool stop"
    "/etc/init.d/dcache-core stop"
  stop_dcache.dcache::
    "/etc/init.d/pnfs stop"
    "/bin/umount /pnfs/fs"
  stop_dcache.!dcache::
    "/bin/umount /pnfs/test.tier2.hep.manchester.ac.uk"
  upgrade_dcache_database::
    "/usr/local/bin/upgrade-dcache-database.sh"
  update_pkg::
    "/bin/rpm -ve --nodeps dcache-server dcache-client"
    "/bin/rm /opt/d-cache/config/*.batch"
    "/bin/rpm -Uvh http://ks.tier2.hep.manchester.ac.uk/linux-install/yum/dcache/dcache-server-1.7.0-34.noarch.rpm http://ks.tier2.hep.manchester.ac.uk/linux-install/yum/dcache/dcache-client-1.7.0-31.i586.rpm"

    "/usr/bin/yum -y install dcache-server-1.7.0-34 dcache-client-1.7.0-31"
  dcache_upgrade.dcache::
    "/opt/glite/yaim/scripts/configure_node /opt/glite/yaim/etc/site-info.def glite-SE_dcache_admin_postgres" umask=000
  dcache_upgrade.!dcache::
    "/opt/glite/yaim/scripts/configure_node /opt/glite/yaim/etc/site-info.def glite-SE_dcache_pool WN_torque" umask=000
```

# Classes in Actionsequences

control:

do\_dcache\_upgrade::

```
actionsequence = ( editfiles.dcache_upgrade packages.dcache_upgrade  
shellcommands.stop_dcache copy.dcache_upgrade  
shellcommands.upgrade_dcache_database shellcommands.update_pkg  
copy.dcache_upgrade shellcommands.dcache_upgrade )
```

# Resources

- Mailing Lists
  - <http://www.cfengine.org/discuss.html>
  - [hep-project-grid-cfengine@cern.ch](mailto:hep-project-grid-cfengine@cern.ch)
- Web Sites
  - <http://www.cfengine.org/>
  - <https://www.gridpp.ac.uk/wiki/Cfengine/>