

# A (very) brief introduction to cfengine.

Matt Doidge, Lancaster University.  
GridPP Storage Workshop, July 2<sup>nd</sup> 2009.

# What is cfengine?

- Cfengine is an open source fabric management and “automatic maintenance” tool.
- It has been around for a while, first released in 1993.
- Class based system.
- Handles configs, packages, environments and events.
- Works using a “final state” paradigm.

# Cfengine Workflow.

- Define the desired machine final state within the cfengine configs- easier said than done!
- Store these on your cfengine nameserver.
- Install cfengine on your node. (we do it from dag).
- Copy over the update.conf.
- Run the cfagent.
- Profit. Hopefully.

# Gubbins.

- We use an svn “backend” to manage config versions and changes.
- Connections between the nameserver and client machines are authenticated using public/private keys.
- Main configs to worry about are update.conf, cfagent.conf and class-specific \*.cf files.
- Invokes 3 daemons: cfservd (server daemon), cfexecd (oversees periodic running of cfagent) and cfenvd (which performs anomaly detection).

# The Good.

- Combined with PXE/Kickstart really takes the sting out of re-installing nodes.
  - A great help during Step'09 when a pool node blew a system disk.
- Takes work away from yaim.
- Makes rolling out small changes a breeze.
- A major asset in keeping your workers homogeneous.

# The Bad.

- It ain't intuitive. At all.
  - Setting up a new server type can be a pain.
- The chances of rolling out a bodge-up across your whole cluster is increased.
  - Luckily rolling back is also easier.
- Has some odd nuances that can lead to peril (- is translated to \_, IPADDR range group definitions are dodgy, possible to break your nameserver).