



System Configuration and LCFG

Paul Anderson <dcspaul@inf.ed.ac.uk>

<http://homepages.inf.ed.ac.uk/dcspaul>

LCFG Tutorial, June 2007

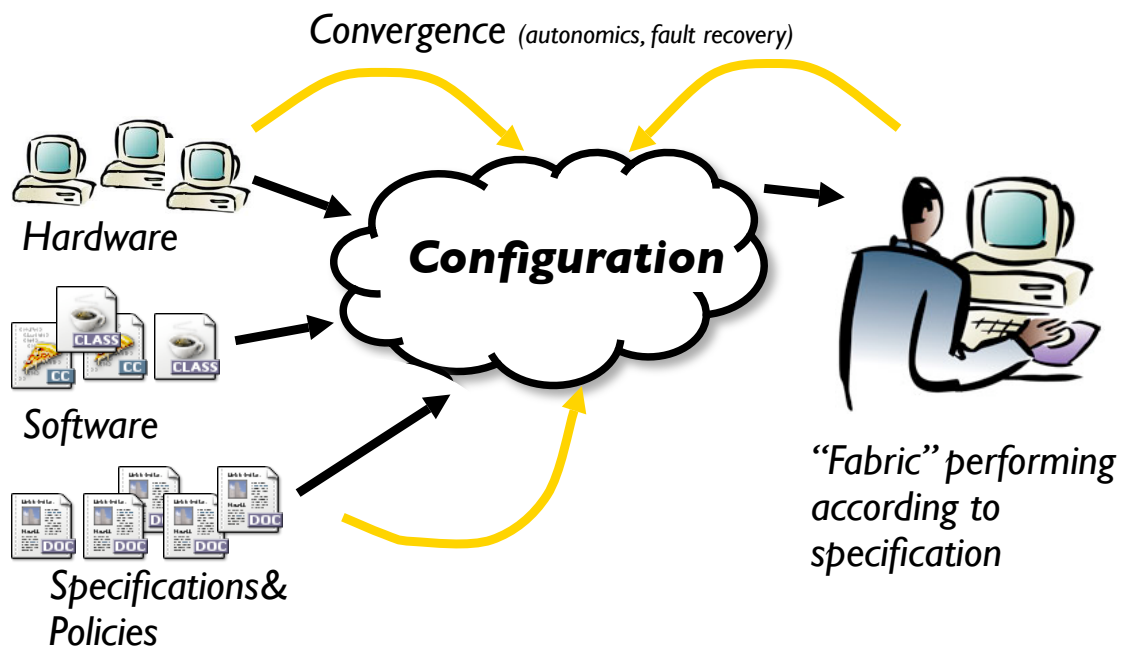
Overview

- System configuration
 - *what exactly is “System Configuration”?*
 - *the manual process*
 - *complicating factors*
 - *the “state of the art” in configuration automation*
 - *the way forward*
- LCFG

System Configuration

- Starting with:
 - *several hundred new machines with empty disks*
 - *a repository of all the necessary software packages*
 - *a specification of the required service*
- Load the software and configure the machines to provide the service
- This involves many internal services:
 - *DNS, LDAP, DHCP, NFS, NIS, SMTP, Web ...*
 - *the relationships are most important*
- Maintain the specification when things change
 - *either the requirements, or the system (failures)*

System Configuration



Complicating factors

- Managing relationships - *clients & servers*
- Managing change - *failures & changing requirements*
- Managing diversity - *servers to laptops*
- Devolved management - *aspects*
- Distributed systems - *communication failures & latency*
- Usability - *multiple levels of experience*
- Autonomics - *unattended reconfiguration*

An example...

Create a new web server

- Infrastructure support
 - *create a DNS entry*
 - *create a DHCP entry*
 - *create holes in the firewall*
 - *create and sign SSL certificate*
 - *add to backup system*
- Configure the machine
 - *Configure disks, load the software etc ..*
 - *Configure dns, networking, timeservice, apache etc ..*

Some difficulties

- The relationships have to be maintained manually
 - *the machine running the web server & the firewall*
- How do we know the configuration is correct?
 - *there is no explicit representation*
- How do we configure a replacement?
 - *some “aspects” of a replacement machine may be different (eg. disk layout, network devices)*
 - *this makes simple restore from backups insufficient*
 - *we may have to update all the associated services*
 - *different people may be responsible for different services*

The “state of the art”

- Most automation of system configuration grows from the “bottom up”
 - *low-level manual procedures are automated*
 - *higher-level decision making is still manual*
- Most current tools are low-level tools which are designed to interface with humans
 - *this makes it harder to incorporate them in a complete “stack”*
 - *no representation of “relationships”, for example*
- Administrators must deal with low-level issues
 - *like assembly-code vs high-level languages*

Higher-level automation

Why?

- Efficiency - *obviously*
- Manage the complexity
- Have confidence in the correctness
 - *What happens if one machine in 1000 node cluster has the wrong version of a maths library?*
- Security
 - *Can we be confident there are no insecure dependencies?*
- Reliability - *autonomic reconfiguration*

The way forward

- Think about the whole “stack”
 - *even if you can only automate the bottom layers*
- Use tools to automate the bottom layers
 - *but make sure that these tools present a clean interface, so that there is potential for higher-level automation*
- Think in terms of “declarative specifications”
 - *ie. “what we want to be true”*
 - *let the tool compute the “how to”*
 - *avoid undisciplined “scripting”*
- Beware of marketing - *think about these issues!*

LCFG

- Some properties
 - *a clean interface to the lower layers*
 - *modular components for different subsystems*
 - *support for fully “prescriptive” configuration (no manually configuration or installation)*
 - *support for automatic maintenance of relationships*
 - *support for “aspects” managed by different people*

References



- Slides for this talk:
<http://homepages.inf.ed.ac.uk/dcspaul/publications/lcfg2007a.pdf>
- System Configuration, Paul Anderson
Short Topics in System Administration, ISBN 1-931971-42-0
http://www.sage.org/pubs/14_sysconfig/
- Paul Anderson
<http://homepages.inf.ed.ac.uk/dcspaul>
<dcspaul@inf.ed.ac.uk>