

Automated System Administration

We are all individuals

- Unique systems
- Hard to replicate
- Un-necessarily complex
- Mostly undocumented
- Painful to manage

I am not

- We actually have lots of identical machines
- We have problems keeping them identical
- Sysadmins need to stop managing individual machines and manage groups instead

Solutions

- System Administrators Anonymous
 - The obvious three step solution
 - Admit there is a problem
 - Hire more people
 - Add more complexity
 - Go to step one.

Processify!

- Make a process
 - Follow it religiously
 - If it fails, it's the fault of the implementors
 - If it succeeds, you did things right

ITIL

- A framework of best practice approaches
 - A set of processes
- Requires something termed as a Configuration Management Database (CMDB)
- Everything gets defined in the CMDB
 - If we don't understand it, it goes into the CMDB
 - That's an implementation issue

Go away, or I will

- ... replace you with a small shell script
- If it's a process, it can be scripted
 - Lots of people don't recognise this
 - Also see the Agile programming philosophy
- Very few people do this in sysadmin land
 - But lots of programmers do

Benefits

- Auditability
 - We know who changed what and when
- Documentation
 - Everyone needs it, everyone hates writing it.
 - Why document twice? The code is the documentation
- Make work consistent
 - Of course, we can rebuild systems any time we like
 - We have backups of everything

More benefits

- Make work testable
 - Of course, we all have test infrastructures for everything
 - All our admins test out configuration changes
- Make it work faster and better
 - Get paid to have a vacation

We need tools

- Managing systems by hand is like writing programs in assembly
 - We have better tools available now
- No one really does that any more
 - John Carmack is odd, and even he doesn't write everything in assembly.

Here's a list

- A few examples
 - <http://www.cfengine.org>
 - <http://reductivelabs.com/projects/puppet>
 - <http://www.lcfg.org/>
 - <http://www.isconf.org>
 - <http://rsug.itd.umich.edu/software/radmind/>
 - <http://www.mcs.anl.gov/cobalt/bcfg2>

We use Puppet

- The next few slides will refer to Puppet configs.
- You need to make your own choices.
- If you need commercial tools, there's a list available in the config-mgmt FAQ at <http://lopsa.org/ConfigMgmtFAQ>
- Open source tool comparison: http://en.wikipedia.org/wiki/Comparison_of_open

Sample config

- Host definitions

```
node 'Rclub_LB.mailbox.inbound.us.1' {  
    include    mta  
}  
node 'Rclub_LB.mailbox.inbound.us.2' {  
    include    mta  
}
```

Example config

```
class mta {  
    package "postfix" : {  
        ensure => latest  
    }  
    package "ntp" : {  
        ensure => latest  
    }  
    file {"/etc/ntp.conf":  
        source => puppet://etc/ntp.conf,  
        owner => "root",  
        group => "root",  
        mode => 644  
    }  
}
```

Changing configuration

```
class mta {  
  package "caching-nameserver": {  
    ensure => 9.3.4  
  }  
  file { ["/etc/postfix/blacklist"]:  
    source => puppet://etc/postfix/blacklist  
    ensure => present  
  }  
  exec { ["postmap /etc/postfix/blacklisted_clients":  
    path => ["/usr/bin", "/usr/sbin"],  
    subscribe => file["/etc/postfix/blacklisted_clients"],  
    refreshonly => true,  
  ]  
  ...  
  ...  
}
```

Alternatively

```
node 'rclub_lb.mailbox.inbound.us.1' {  
    include mta  
    include ntp  
}  
  
node 'rclub_lb.mailbox.inbound.us.2' {  
    include mta  
    include ntp  
}
```


Classes

```
class mta {  
    package "postfix": {  
        ensure => "latest"  
    }  
}  
  
class ntp {  
    package "ntp": {  
        ensure => "latest"  
    }  
    file { "/etc/ntp.conf" :  
        source => puppet://etc/ntp.conf  
    }  
}
```

OS changes?

- The previous config will work with CentOS, RedHat, Debian, FreeBSD and Solaris
- Migration between operating systems is easy
- So is adding new hosts of the same type
- Administrators just need to express policy, not implementation details
- Normalisation helps a lot, just ask your DBAs

Basic thought processes

- Lower admin headcount
 - Communication between humans is expensive
- Describe the solution, don't prescribe
 - Prescriptive solutions are easier to imagine, but less flexible
- Pull data, don't push it out.
 - Simplifies the code

Recommendations

- At least one company manages 2500 servers with 4 administrators total and Puppet
- Lindenlabs (Second Life)
- Google has their own homegrown toolkit, but they follow the same philosophy.

Problems

- Technology
 - Not all things can (or should) be automated
 - Not all tools are equal, research is needed to choose the one which fits your situation best
- Politics
 - People don't like changing the way they work
 - This is a big change

More problems

Skills

You need fewer administrators, but with a far higher skillset

Time

Implementing automation takes time and effort

Management buy-in

Hard to get resources