

# Why the Sysadmin Hates Your Software

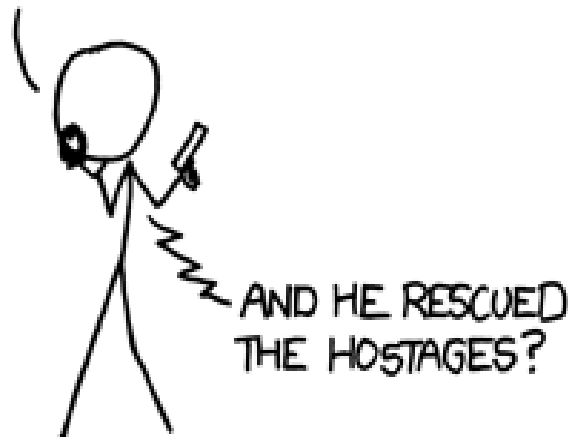
Steve VanDevender  
University of Oregon

# Because we're cranky?

WE TOOK THE HOSTAGES,  
SECURED THE BUILDING, AND  
CUT THE COMMUNICATION  
LINES LIKE YOU SAID.



BUT THEN THIS GUY CLIMBED UP  
THE VENTILATION DUCTS AND WALKED  
ACROSS BROKEN GLASS, KILLING  
ANYONE WE SENT TO STOP HIM.



NO, HE IGNORED THEM.  
HE JUST RECONNECTED  
THE CABLES WE CUT,  
MUTTERING SOMETHING  
ABOUT "UPTIME".



# Sympathy for the Devel<sub>o</sub>pers

- We want similar things
  - Simplicity
  - Modularity
  - Robustness
  - Functionality
- Complexity is hard, but we have to cope anyway
- We want people to use our software and systems

# Sysadmins work at a different level

- Sysadmins
  - Software and configurations
  - Hosts
  - Networks
  - Services
  - User experience of whole systems
- Developers
  - Programming languages
  - Libraries
  - APIs
  - Build systems
  - User experience of particular product

# Sysadmins interact with software differently

- Installation
- Maintenance
- Troubleshooting
- Maintaining many, many items of software
- Sysadmins are often **not** users of the software they're maintaining
- Sysadmins help make your software available to users

# System administration is changing

- Larger, more complicated systems
- Many more hosts and networked systems
- Need for replication
- Increasing use of automation
- Less time for elaborate hand-crafting of software

# Tools of the modern sysadmin

- Version control for system configuration
- Configuration management systems
  - Examples: Puppet, cfengine
- Automated installation systems
  - Examples: Kickstart, Jumpstart
- Virtual machines
- Sometimes tools like CMS or VMs are used to prop up poor software

# What do configuration management systems do?

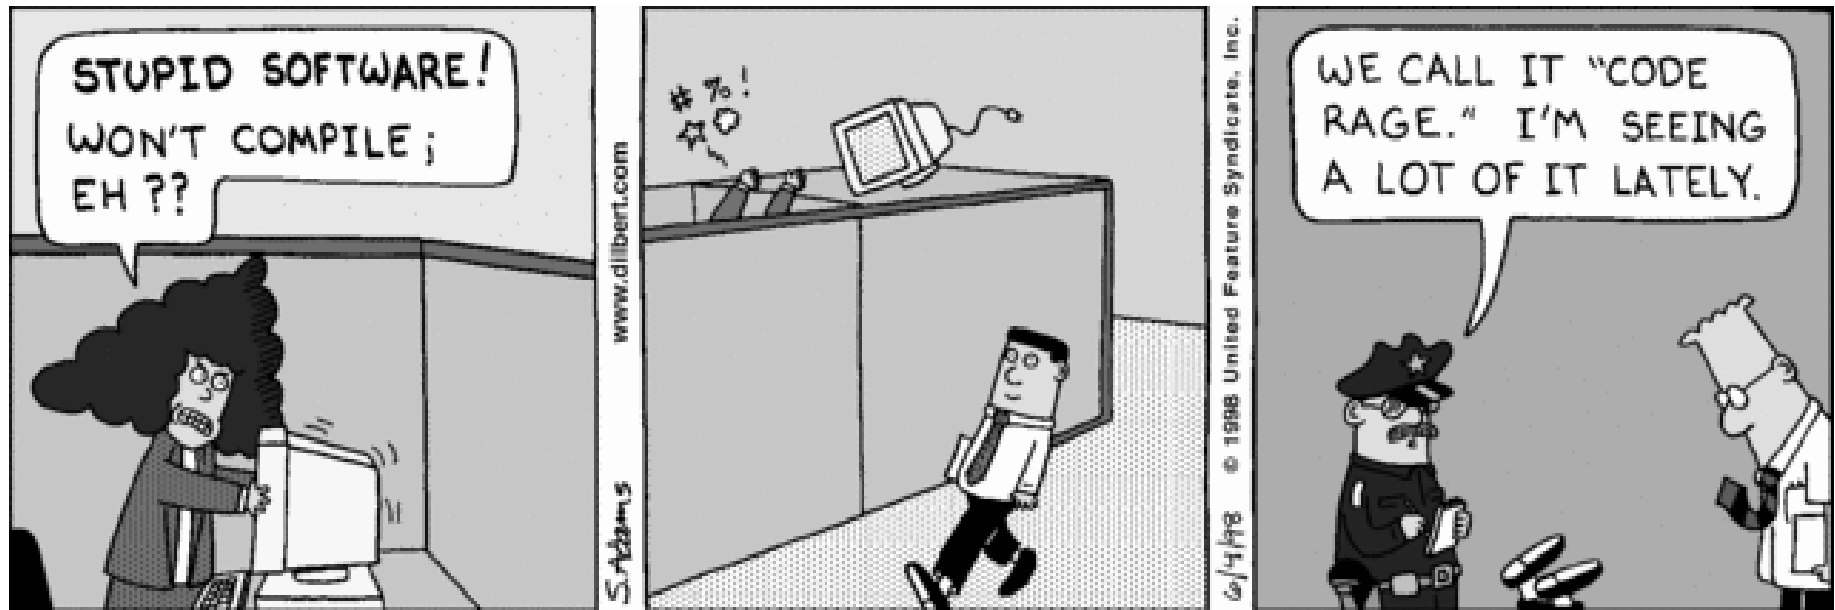
- Specify intended state of system resources
- Continually inspect system state, modify to match specification
- Configuration management is almost always closely integrated with version control
- Essentially using software to manage other software through the entire process of installation, configuration, and operation



# Puppet configuration example

```
package { "sendmail":  
  ensure => installed,  
}  
  
file { "/etc/mail/sendmail.cf":  
  owner   => root,  
  group   => root,  
  mode    => 444,  
  source  => "puppet:///sendmail/sendmail.cf",  
}  
  
service { "sendmail":  
  enable      => true,  
  ensure      => running,  
  hasrestart  => true,  
  hasstatus   => true,  
  require     => [ File["/etc/mail/sendmail.cf"], Package["sendmail"] ],  
  subscribe   => File["/etc/mail/sendmail.cf"],  
}
```

# So, why does your software suck?



# Labor-intensive manual installation

- “Just specify these 37 simple configuration items in our GUI/browser interface/install script”
  - Aaarrggg. Kill me now.
- “Copy this file there, copy that file over there, run this, edit that”
  - Yuck.
  - But it may be easier to automate.
- What we want is automated installation **and** uninstallation, as well as ways to automate configuration and customization

# Better approaches to installation

- `./configure; make; make install`
  - Less manual effort
  - Not always easy to uninstall or upgrade
  - At least maybe I can build my own package
- `dpkg -i foo_1.2.3_i386.deb`
  - Least manual effort for sysadmin (or end user) to install
  - Easy to uninstall (`dpkg -r`)
  - **If** package dependencies have been specified properly and package was built to standards



# Dependency hell

- “. . . lives high on the software food chain”
- Reusing modular components is (usually) good
- Increased installation footprint and management complexity is bad
- Try not to bundle specific library versions, use system libraries as much as possible
- Put careful thought into library APIs for portability, backward compatibility

# Horrible error messages

- “Invalid Database Collation”
- Tell me what you want, not that something went wrong
- Have messages that are of some help to end users
- Java tracebacks are not useful error messages

# Why sysadmins hate browser-based configuration interfaces

- “You can't grep a web page”
- Replication of configuration is tedious
- Version control is nearly impossible
- Automated configuration management is nearly impossible



# Why sysadmins love config files

- Ideally, encapsulate configuration state in one object
- Easy to search and compare
- Easy to distribute and replicate
- Lots of tools for automated editing
- Easy to put in version control
- Easy to directly manage with configuration management systems

# Fighting over config files

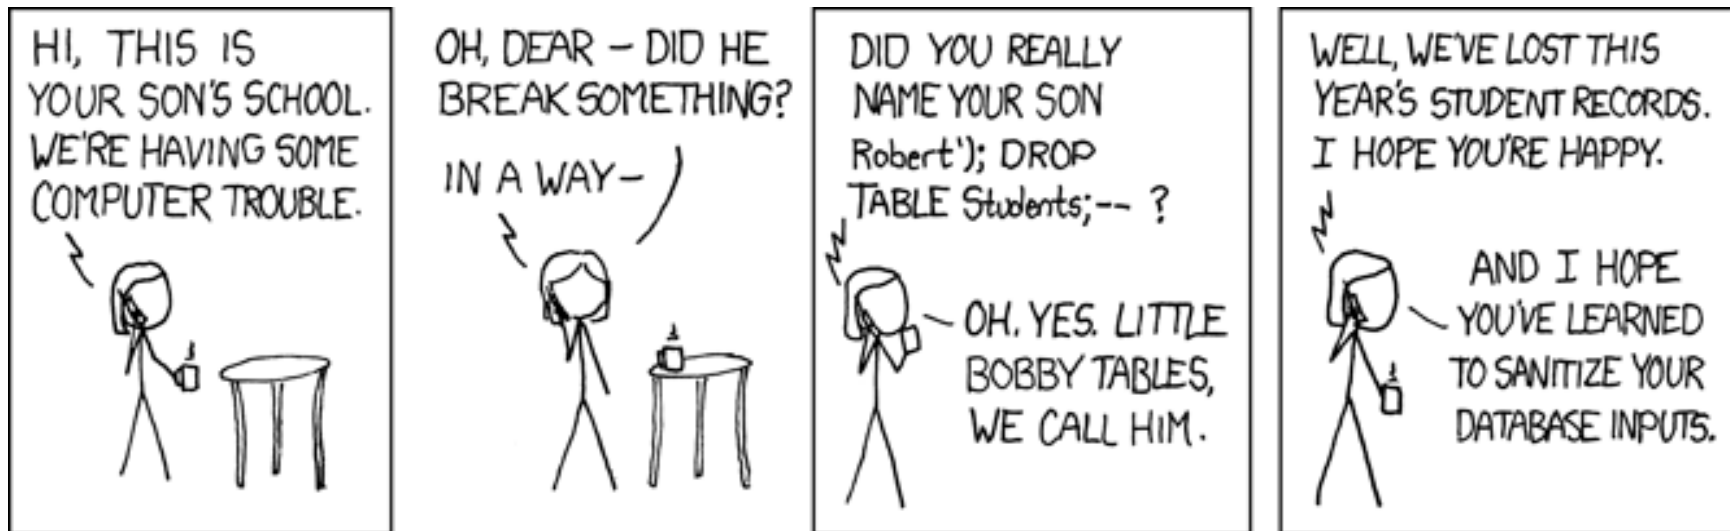
- Software writes its own config files (usually from GUI/browser-based interface)
- Sysadmin cheats and grabs underlying config files to put in configuration management
- Software rewrites its config file
- Configuration management system sees different config file contents, reinstalls its own file
- Software freaks out over unexpected change

# Making everyone happy

- Casual users tend to prefer GUI or browser-based interfaces for configuration
- Sysadmins want easy ways to automate and replicate configuration
- Can you provide a low-level interface that is more satisfactory to sysadmins?
- But please, not yet another database or binary format that requires special editing tools

# Bad security choices

- Naively ship with (or require) world-writable files, administrator or superuser privileges
- Default passwords or unrestricted access
- Please not another separate authentication database
- Poor input validation leading to security exploits



# Documentation

- At least try to explain what your software does
- Highlight unusual installation requirements
- Have error messages that line up with documentation

# Summary

- Be good at doing things you're supposed to be doing anyway
  - Modular, flexible design
  - Software that fits well into its intended environment
  - Helpful documentation and diagnostics
- What sysadmins want is different from what users want
  - Automation of all aspects of software management