LISA 2005 Conference Summary

John Borwick, <borwicjh@wfu.edu>

Summary

1,300 system administrators attended LISA 2005, from education, government, and industry.
I want to review ITIL, a British IT management standard, to see what it recommends that would be helpful in our environment. We should rearrange our administration calendar cycle so that we schedule more downtime between semesters and less during semesters.
I believe we should replace our real-time server monitoring systems with Nagios, a popular open-source server monitoring tool. We should buy a syslog server with lots of cheap hard disk space and make all our machines log to that server in addition to their current logging.
We should consider backing up to cheap disk, and archiving all incoming email so that it can be retrieved on a message-by-message basis.
The PMO office should create a "unified project view" that everyone in the department can see. This view would include deadlines for each person in the department, so it is easier to tell when people are under- or over-committed.

Itinerary

I attended tutorials, which were half-day and full-day classes; technical sessions, which were papers, "Guru sessions," and invited talks; and "Birds of a Feather" meetings at night.

Tutorials

Building a Logging Infrastructure and Log Analysis for Security

Go to `http://www.loganalysis.org`. Abe Singer has 2,000 hosts logging to one syslog server. That server has a lot of bandwidth and disk space.
You can add tcpwrappers to programs to make them log. You can log via iptables. UDP can be dropped easily; TCP has a high overhead. You can combine mark messages with Nagios to make sure all hosts are sending logs.

Security without Firewalls

To analyze security, you need to analyze trust relationships. Trust relationships are any assumptions a machine makes that are dependent on external sources. You can put your root password in an envelope in a safe, and then when someone needs it you can tell that they broke the envelope's seal and therefore that you need to change the password.
When auditing hosts, look for SUID/SGID scripts, world-writable files, user accounts, that everything to ~root is owned by root, and look at daemons.

SDSC treats wireless guest logins as off-campus connections.

Welcome to My ~/bin

Chip Salzenburg, Perl pumpking for Perl 5.004, taught this class. Most everything in the class was esoteric, but I liked it.

Over the Edge System Administration, Vol. 1

David Blank-Edelman printed out an mp3 with lpr to show that traditional tools can be used in new ways. WWW::Mechanize::Shell can give you a command-line interface to web-only tools. Digital cameras can let you see things you otherwise couldn't.

Databases: What You Need to Know

This class was not as helpful for me.

Papers and "Guru Sessions"

Scaling Search Beyond the Public Web

A Yahoo representative talked about how the web is being divided into three search spaces: public spaces, communal spaces, and private spaces.

Backups

Backups are different than archives. Do not treat your archives the same way you treat your backups. For example, if you treat them as the same thing, you might find yourself having to restore 52 backups to respond to a lawsuit asking for all email over the last year.

Practically speaking you can only get around 50 megabytes per second of data on a gigabit Ethernet connection. LTO drives can write 160 megabytes per second. You should always back up to disk, then to tape. Your OS may not even be able to write at full speed to an LTO drive.

Bacula, amanda, rsync, rdiffbackup, and BackupPC are all good open-source tools for performing backups.

Backing up filers can be difficult, because the NDMP protocol does not specify a backup format. You may find yourself begging an ex-vendor for a piece of equipment to let you restore your NDMP files.

You can create a "hard link farm" for your backups. /backups/backup.0 can be hard links to all your files for that day. Then, if you delete a file, you can find it in your link farm.

"Data reduction" backups can reduce the amount of data stored up to 20:1.

"RTO" stands for recovery time objective. "RPO" stands for recovery point objective.
Never ever sell used backup tapes. They can't be degaussed *and* used again.
You can turn on encryption in Oracle to protect your data there.
Use disk for on-site backups and tape for off-site backups.

Configuration Management Theory

At its highest level, configuration management should let you say "I need a mail service with less than one second of latency" and the tool should go find the appropriate equipment and set it up.

Under 200: Applying Best Practices to Small Companies

A ticketing system, users feeling that the system works, and quarterly meetings about important issues can all help reduce response time and justify hiring.

Spam

The best spam blacklists are SBL/XBL. They're free.
You can turn SPF-checking on for just MSN, AOL, and Yahoo.

Theory

There are some sites for which there are too few administrators. In these sites, you can watch the number of tickets over time: they grow.

Project Management

We should create a "unified project view" of all people on all projects, along with deadlines, to show who is needed when for each project.
In any project, list your scope, schedule, and resources.
Denote the skill sets needed by your team. Showing that you're lacking in one can help with hiring and training.
With scope creep, either close the current project and create a new one for the new scope, or impose a change control process where Important People have to sign off on the change.
Project tasks should be no longer than 2-3 days, or should be deliverables such as someone signing off on a project.
Project reports should include the scheduled completion date and the actual completion date. They should then summarize why these dates are not the same.
For really long projects, have "gateway meetings" each three months to reassess the project scope.

Time Management

43Folders.com is a neat site.
If something is not a high priority, then let it go.
Work on hard high-impact projects rather than hard low-impact projects.
Read trade magazines and listserves no more than 1 hour a week, then throw away whatever you didn't get to read.
Don't make it a habit of taking work home.
If something looks incorrect, and you've brought up the issue twice to management with informed consent (so that you've told them what could break), and they accept the risk, then drop your concern. Management has made their decision and you will waste time trying to change their minds.
Read his book. With the free time you get, don't work more: improve yourself outside of work.

Picking Locks with Cryptology

The speaker talked about ways to defeat people who are tapping your phone.

Change Management

ITIL is more relevant for us. ITIL is really complex. The speaker likes the "BS15000" view of ITIL.
The three big components of ITIL, for us, are

- Release (pre-production): the best organizations concentrate resources here
- Control (production): NO DEVELOPERS ALLOWED
- Resolution (problems)

It's almost impossible to get above 85% uptime with developers in production.
One site's highest availability was when all the DBAs went to Oracle World.
80% of outages are due to changes.
With a problem dependency tree, you can fix 50% problems. You can fix 80% of the problems without actually logging in or rebooting.
A bureaucratic change management process decreases success rates by 15%.
It should be cheaper to rebuild a system than to repair it.

Security/Cryptography

CAcert.org may be a good solution to letting us encrypt intra-campus email.

Apache/OpenSSL/DNSSEC

DNSSEC isn't worth implementing.
Name-based SSL hosting may be worthwhile. You can get vendors to sign an SSL certificate with several cn's in it, allowing you to do name-based hosting.

The director of the Apache Software Foundation does not use Apache 2.
A binary diff analysis can tell you what holes Microsoft is patching.

Work-in-Progress Reports

One presenter wrote "VNC Manager," a simple GUI that lets you keep track of your VNC sessions.
A Cisco administrator runs one Nagios server watching over 15,000 hosts.

"Birds of a Feather" Sessions

Open Source in the Workplace

Scripts and Tips

Server Automation

Puppet

LOPSA

cfengine

Oslo's Masters Degree Program

CMS Options

Cool UNIX Tools

- dtrace
- fping
- HTTP::Recorder
- MIME::Lite
- perl -i~
- syslog-ng
- tail -F
- WWW::Mechanize::Shell