

SANE 2004 Conference Summary

borwicjh@wfu.edu

October 12, 2004

“We tend to react ‘organically’ to organic problems, worsening the situation.
We need to engineer solutions.”
—Geoff Halprin, president of the System Administrators’ Guild

1 Summary

SANE 2004 was held in Amsterdam. The United States attendees were Wake Forest University, Ohio State University, VERITAS Software, Google, Roble Systems, USENIX, and UC Berkeley. I attended three days of tutorials and two days of presentations.

Wake Forest University could learn a lot from what was presented at SANE. I have included occupation-specific ideas below. “Soft topics” like administration policy and “creating happy users” were covered along with technical topics like spam reduction methods, IP telephony, and change management software.

System administrators’ jobs are initially considered reactive. “Anyone can install a machine,” says Geoff Halprin; it is easy to misunderstand our role in an organization. System administrators integrate products and make disparate services work together. We “manage change.” Our challenge is to become proactive. Tools like the SA-BOK, at <http://www.sysadmin.com.au/sa-bok.html>, explain what we do and can help describe why we need time to perform preventative maintenance and planning.

After becoming a proactive entity, Tom Limoncelli describes what is necessary to “perform triage” and prepare for progress: stable email, repeatable processes, physical issues like security and labeling, real-time and historical monitoring systems, collaborative documentation, remote server control, and stable backups. Tom also cited studies showing that people become more productive as their workplace becomes quieter.

eBay’s system administrator, Paul Kilmartin, gave the keynote. eBay processes \$1,015 per second. They have a few policies: avoid the cutting edge, acknowledge they have to re-invent their infrastructure every three years, and deploy processes that scale in a maintainable way.

Email improvements immediately make presidents, CEOs, and customers happy. I learned several technical ways to improve email transmission and delivery. Postfix, an open-source alternative to Sendmail, can reduce demands on our server and check for common spam methods. “Maildir” is a way to deliver and retrieve mail to a directory rather than to an “mbox”-style file, which reduces mail server load and lets mail delivery scale.

2 Itinerary

I attended the following tutorials:

- *Running IP telephony on your network*
A Cisco engineer taught this class. Topics included power over Ethernet, Quality of Service, and Tail-End Hop Off for making long-distance calls local. Notably, he mentioned that SIP is a new protocol not widely deployed or tested.
- *But is it UNIX? A Mac OS X administrator's survival guide*
Aleen Frisch of “Essential System Administration” fame taught this class.
- *Practical Postfix*
The basics and some intermediate concepts were introduced in this class, taught by a Postfix developer.
- *Linux 2.6 processor and process management*
We covered some very detailed examples of how Linux 2.6 deals with CPUs, processes, threads, and scheduling.
- *Creating Happy Users: A “How-To Guide” for Sysadmins in a Hurry*
Tom Limoncelli taught this full-day class. He described his help desk process, the importance of first impressions, thinking as “advocates” instead of “clerks,” how to keep in contact with users, how to effectively contact users, how to create SLAs, and how to deal with political issues. Several people recommended using RIS—remote installation services—for Windows installations.

I attended the following presentations:

- *Inside eBay.com: the System Administrator's Perspective*
- *MySQL Roadmap*
- *Using Virtual Machines for System and Network Administration Education*
- *Managing Large-Scale Samba Deployments*
- *NFS, Linux, and Clusters*
- *TCG 1.2—fair play with the ‘Fritz’ chip?*
- *Distributed Software Deployment with Subversion and Submaster*
- *The Changing Face of System Administration*

- *High-Availability Load-sharing with OpenBSD*
- *Development of World-Wide IDS Networks*
- *ISGTC: an alternative to bofh/bin*
- *The Effects of ‘Tit for Tat’ Policy for rejecting spam or denial-of-service attacks*
- *UNIX and the ARPAnet/Internet at 35: Linux a teenager; still in court*

Also, I heard Richard Stallman speak about the origins of the GNU Public License and GNU/Linux.

3 Notes by Job Type

3.1 Communications

- You should talk to new users twice, once on their first day and again a week later after they’re no longer overwhelmed
- Introductory documentation should be 1-2 pages long
- If someone’s first week is a bad IT experience, they will always dislike IT
- Consider yearly “Town Meetings;” IBM’s first year they had huge numbers of complaints, but things got better and users were happier
- Tom Limoncelli conducted one-minute meetings with relevant managers (cf. department heads) to listen and announce future projects
- One site took every x-thousandth help desk user out to lunch

3.2 DBAs

- eBay used solid-state drives (giant RAM disks) for their REDO logs
- “Don’t run batch jobs on the transaction engine!” says the eBay system administrator
- Vendors like Cisco incorporate MySQL into their products
- MySQL can be ACID-compliant, with transactions
- MySQL is working on creating a version of PL/SQL for version 5.0
- MySQL supports point-in-time recovery
- MySQL supports dynamic views

3.3 Help Desk

- Many customers think the Help Desk is the entire organization
- Staff need to know the metrics being used
- One system only has “good,” “ok,” and “bad” links in the completion email for users to click on, to make sure response is high
- Cisco’s TAC escalation policy is a good example; problems theoretically go all the way to the CEO
- RT at <http://www.bestpractical.com/rt/> has some ideas for help desk ticket tracking
- Tom Limoncelli has a good model for help desk problems:
 1. The Greeting
 2. Problem Classification
 3. Problem Statement
 4. Problem Verification
 5. Solution Proposals
 6. Solution Selection
 7. Execution
 8. Craft Verification
 9. Customer Verification

Different people, or groups, can be responsible for each step. You can always move backwards in the process, e.g. from “solution proposals” to “problem statement.” He has a useful flow chart.

- Walking users through a problem tends to take nine times as long as doing it yourself
- User satisfaction is different than whether the problem was solved
- Caller ID can be used to send users to the same people

3.4 Networking

- IP phones can use power over Ethernet and can trunk VLANs
- Voice’s maximum one-way latency is 150ms
- Only one packet can be lost at a time before clipping occurs
- Spanning tree algorithms like 802.1w or 802.1s should be used to reduce delays for voice

- Some codecs like G.729 are bad for non-voice, such as hold music
- SIP is very new
- You can trade between phone bandwidth cost and hardware cost: cheaper phones need more bandwidth
- All phone numbers should be normalized to E.164 for routing purposes
- Attacks like MAC flooding, DHCP gobbling, and ARP spoofing have new meanings when it comes to voice
- Monitoring should use less than one percent of resources
- SNMP discovery can reveal the community string to snoopers
- The community string should be changed regularly and you should try old and common passwords to see if any devices are using them

3.5 Systems: General

- Use RIS for Windows to install machines
- You need “repeatability” for everything, so you don’t end up with “accidents of history”
- Think of the Help Desk as an “interrupt shield;” ask “what did the help desk say?”
- Allow interruptions to *record* requests, but don’t fix them immediately
- Prioritize problems
- You can be a “clerk” or an “advocate”
- Many people recommended *The Practice of Programming*, 1999
- Do a task manually once, then record the steps and automate the task
- Salary and management do not improve productivity as much as quiet space
- “Tier 1” employees should be easy to find, while “Tier 2” employees should be secluded
- SLAs should drive funding levels
- In general, the stages for fixing system administration is triage, then fixing things, then growth
- When you find a problem, write a script to test for it, and incorporate that into real-time monitoring

- You need a “snooze” button for monitoring systems
- You should monitor as much as possible
- Monitor as soon as a server is installed
- We may need “fire drills” for critical servers, to find and verify the time to recovery
- eBay doesn’t use SAN zoning *AT ALL* and they have over 1000 machines on their SAN
- eBay forces their backup site to have traffic so they know it will work in a disaster
- Make vendors fix your problems
- Make vendors give references for “best practices”
- One eBay slide was a big hit: “Certification != Competence != Expertise”
- “Trusted Computing” is really scary
- There are over 750,000 system administrators in the USA
- We need to change from reactive to proactive workers, and need recognition for being proactive
- We have a large sphere of influence because we affect everyone in the organization
- Describe things in terms of how we’ll save money, make money, or avoid legal problems

3.6 Systems: UNIX

- Sometimes you need to use `star` to archive file ACLs
- CUPS is awesome; use `http://localhost:631` for management
- The author of Postfix also wrote SATAN and TCP wrappers
- Postfix has lots of smaller programs to do its bidding, only a few of which are SUID
- We should consider using “maildir” instead of mbox for storing user email
- If and only if you say “250 OK” in an SMTP session, you’re responsible for the email’s delivery
- Postfix has several queues for incoming and deferred mail to help performance

- You can totally mess up a system with `chrt`—“change real-time attributes”
- Processes in non-interruptible sleep cannot be killed
- Linux 2.6 will be able to understand “node” and “SMP” affinity through “scheduling domains”
- Real-time and historical monitoring should be separate. We should investigate Nagios for real-time monitoring and “cricket” for historical monitoring
- 10 gig E is coming
- NFSv4 is cluster-friendly
- NFSv4 may be able to replace AFS
- NFSv4 has Windows-native features
- NFSv4 may support iSCSI
- We should replace CVS with Subversion
- You can do some *RIDICULOUS* load-sharing things with OpenBSD, duplicate MAC addresses, and a balancing program
- We need a policy for how shell scripts are written and how they are used; see <http://isg.ee.ethz.ch/tools/isgtc> for a tool package