;login:
THE MAGAZINE OF USENIX & SAGE
August 2002 volume 27 • number 4

inside:

CONFERENCE REPORTS
SANE 2002

USENIX & SAGE
The Advanced Computing Systems Association &
The System Administrators Guild

This issue's reports focus on the Third International System Administration and Networking Connference (SANE 2002),

Our thanks to the Summarizer:

Diomedis Spinellis

conference reports

Third International System Administration and Networking Conference (SANE 2002)

Maastricht, The Netherlands

May 27–31, 2002

Summarized by Diomidis Spinellis

dds@aueb.gr

SANE, co-sponsored by USENIX and the NLnet Foundation, has evolved to be the European equivalent of the US-based system administration conference (LISA).

A lively and colorful crowd of systems-related attendees (including the obligatory UNIX elders), copious amounts of food, wireless Internet connectivity, interesting poster presentations, and a technical exhibition made the conference a fun place to be. Two parallel tracks of very interesting papers made the selection of presentations a real challenge. The following summaries are therefore only a subset of the conference's presentations. See the conference's Web page at *http://www.nluug.nl/events/sane2002/*.

The conference's keynote address was made by Bill Cheswick who described his Internet mapping work at Bell Labs that resulted in founding the Lumeta startup company. Mapping the Internet is becoming more and more difficult. Drawing routes from one point to another on a geographical map does not reveal any useful information in densely wired areas like North America and Europe: all that appears is a solid blob. More interesting are diagrams that depict the routes between different networks, arranged in a spring-like fashion, with well-connected networks appearing in the diagram's center and leaves at its periphery.

Color is used to distinguish different network providers, network addresses, or administrative domains (e.g., countries). Interestingly, directly representing the network IP address using a (red, green, blue) triple results in a map drawing where a mouse can directly determine the address of a given network by hovering over a particular color.

The applications of this research are numerous. Bill described how, after the events of 9/11, he ever more frequently finds himself at meetings in Virginia or Baltimore with individuals who refuse to identify themselves or the government branch they are working for! An animated map of Yugoslavia network connectivity during NATO's bombardments was especially interesting. We could see network links appearing and disappearing and total connectivity dropping as NATO started targeting Yugoslavia's infrastructure. "Son, you are making remote damage assessment from your basement," remarked one general about Bill's work.

Of course, viewing only Internet addresses and domain names has its limits. The excitement of a discovery of a group of Yugoslavian hosts (.yu) that proved to be extremely well connected, and resided somewhere in Virginia, was slightly tempered when he found out that he had only discovered the Yugoslavian embassy in the US.

Christine Hogan, co-author (with Tom Limoncelli) of the book *The Practice of System and Network Administration*, gave a talk on "Scheduled Maintenance Windows." This concept allows you to proactively plan your system maintenance (and thereby manage your users' expectations). The system administrator's role in such an exercise is similar to that of a flight director, the person you see in historical space-flight films managing the entire operation in the flight control center. The flight director knows

the mission details but does not participate in the actual operations, thus distancing herself from the task and keeping the clear head needed during the maintenance window's stressful hours. She is the one who will notice that the maintenance operations are running behind schedule and will command that the system should revert to its previous state (you did keep a backup, didn't you?), thus averting service disruption.

Mark Burgess, from Oslo University College, gave a thought-provoking talk titled "System Administration as Communication over a Noisy Channel." Mark believes that Shannon's communication theory can be used as the underlying foundation for explaining and predicting a number of phenomena related to system administration. And this is what science is all about. Specifically, Mark considers that a system's policy is communicated over a noisy channel, in which the users of the system represent noise. A significant result of this view is that error correction techniques are needed to create stable system administration tools.

Computer forensics, the study of the legal aspects of digital evidence, are increasingly important to system administrators who will be called to testify as experts in a court of law. Vlasti Broucek, an experienced system administrator who is currently researching this issue at the University of Tasmania, outlined in "Bridging the Divide: Rising Awareness of Forensic Issues amongst Systems Administrators" (co-authored with Paul Turner) the main challenges and techniques for preserving and effectively presenting forensic evidence.

Computer forensics differs from IT security in that it is typically conducted after an attack and its results will be presented to a non-IT-literate audience. Important aspects of digital evidence include its legal admissibility, its validity, and the conduct of the forensic analysis. System administrators faced with the task of collecting evidence should therefore minimize the handling of the original data, account for any changes, comply with the rules of evidence, and avoid embarrassment by not exceeding their knowledge and skills.

A paper co-authored by Giorgos Gousios of the University of the Aegean and your correspondent, "A Comparison of Portable Dynamic Web Content Technologies for the Apache Web Server," presented the main technologies for providing dynamic content on the Web (CGI scripts, PHP, mod_perl, mod_python and Java Servlets) and outlined the results of a series of benchmarks that measured their performance. FastCGI followed by mod_perl appeared to score best in moving data out of the server, but Java servlets proved to be the most resilient. The paper received the conference's best refereed paper award.

Diane Lark from Hewlett Packard presented work on a similar problem in a talk titled "Simulating Web Workloads." The major insight behind her and her colleagues' work was the similarity behind many Internet traffic patterns. They therefore used the SURGE network traffic generator to overcome the deficiencies of SPECWeb96 and Webstone. Through those tests they observed that Web serving is a memory-intensive operation that puts relatively less stress on the processor. They found out that a 1GHz server processor can serve the equivalent of 3000 users.

Your correspondent presented his work relating to the integration of home appliances, in a paper titled "The Information Furnace: User-Friendly Home Control." The Information Furnace is a basement-installed PC-type device that integrates existing consumer home-control, infotainment, security, and communication technologies to provide transparent user-friendly access and value-added services.

A modern home contains a large number of sophisticated devices and technologies. Access to these devices is currently provided through a wide variety of disparate interfaces. As a result, end-users face a bewildering array of confusing user-interfaces, access modes, and affordances. In addition, as most devices function in isolation, important opportunities to exploit synergies between their functionalities are lost. The Information Furnace distributes data, provides services, and controls an apartment's digital devices. Emphasis is placed on user-friendliness and on exploiting the synergies that inevitably come up when these technologies and services are housed under a single roof. The prototype implementation outlined integrates on a FreeBSD server the distribution of MP3-encoded music to DNARD/NetBSD thin clients, an answering machine, a burglar alarm, an Internet router, a fax server, a backup server, and intelligent control of a PBX.

A highlight of the conference, as indicated by the number of attendees, was the talk of Kirk McKusick (chief architect of the Berkeley UNIX and co-author of *The Design and Implementation of the 4.[3/4] BSD UNIX Operating System* books) titled "Running fsck in the Background." Kirk has a talent for simplifying the presentation of highly technical information. The fsck program verifies and fixes the integrity of UNIX file systems. Running such a program on a large disk (e.g., 100GB) can take hours, an unacceptable proposition for production servers. The problem was solved by taking a virtual snapshot of a disk (by temporarily suspending running system calls), maintaining the snapshot current by monitoring disk updates, and running fsck on that frozen snapshot. Snapshots could also be useful for

backing-up system state, and running a dump (backup) operation on a live system. Memorable quote: "I could write a special version of fsck, but I've already written fsck once, and I don't want to do that again."

Cor Bosman from XS4ALL gave a talk on installing and maintaining clusters of servers using PXE and Rsync. Installing software on a large number of servers can be a tricky proposition. Cor explained how he used the PXE remote booting standard supported by most modern Ethernet cards to transparently load and install FreeBSD on server clusters. The procedure is so smooth that having PXE booting enabled on the BIOS of a Windows machine will make it install FreeBSD on the fly — "a software upgrade," as Cor described it.

Mark Overmeer gave a talk on email processing with Perl. It turns out that many of the Perl modules that deal with mail are unsupported, buggy, and lack important functionality. In addition, correctly processing mail elements is a lot more difficult that what it appears to be. MIME encapsulation, multi-part messages, different presentation mechanisms, varying mail user agents, latitude in the mail header specification, and nonconforming implementations conspire to make the implementation of robust mail processing software a Herculean task. Mark worked on overcoming this situation by implementing a complete, robust, and reusable mail processing module (available on *http://www.cpan.org*) that other developers can import when building mail-handling applications. Thanks Mark!

The conference ended with an entertaining talk by Jos Visser titled "Welcome to the Tribe: Socio- and Anthropological Phenomena at UNIX Hacker Conferences," discussing the audience, conference, elders, mythology, economics, ethics, values, moral code, humor and entertainment, enemies, nutcases and outcasts, gadgets, women (section intentionally left blank), and status aspects of hacker culture. At the same time, Peter Salus, on a more somber note, gave a talk on "The Types of Internet Trauma: 1994–2002," where he showed how the Northridge earthquake, hurricane Floyd, fiber cuts, denial of service attacks, and the 9/11 WTC incident affected the Internet's connectivity. Overall the network fared well re-routing packets and compensating in real time.