OCTOBER 2004 VOLUME 29 NUMBER 5

# ;login:

THE USENIX MAGAZINE

## USENIX '04 Conference Summaries

SEE PAGE 41



USENIX
The Advanced Computing Systems Association

# USENIX Upcoming Events

## 18th Large Installation System Administration Conference (LISA '04)

SPONSORED BY USENIX AND SAGE

November 14–19, 2004, Atlanta, GA, USA
http://www.usenix.org/lisa04

## 6th IEEE Workshop on Mobile Computing Systems and Applications (WMCSA '04)

SPONSORED BY IEEE TCOS/TCI IN COOPERATION WITH USENIX AND ACM SIGMOBILE

December 2–3, 2004, English Lake District, UK
http://wmcsa2004.lancs.ac.uk/

## First Workshop on Real, Large Distributed Systems (WORLDS '04)

December 5, 2004, San Francisco, CA, USA
http://www.usenix.org/worlds04

## Sixth Symposium on Operating Systems Design and Implementation (OSDI '04)

SPONSORED BY USENIX IN COOPERATION WITH ACM SIGOPS

December 6–8, 2004, San Francisco, CA, USA
http://www.usenix.org/osdi04

## USENIX '05 Annual Technical Conference

April 10–15, 2005, Anaheim, CA, USA
http://www.usenix.org/usenix05
Paper submissions due: October 22, 2004

## 2nd Symposium on Networked Systems Design and Implementation (NSDI '05)

SPONSORED BY USENIX, IN COOPERATION WITH ACM SIGCOMM AND ACM SIGOPS

May 2–4, 2005, Boston, MA, USA
http://www.usenix.org/nsdi05
Paper submissions due: October 8, 2004

## 3rd International Conference on Mobile Systems, Applications, and Services (MobiSys '05)

JOINTLY SPONSORED BY ACM SIGMOBILE AND USENIX, IN COOPERATION WITH ACM SIGOPS

June 6–8, 2005, Seattle, WA, USA
http://www.usenix.org/mobisys05
Paper submissions due: May 26, 2004

## 10th Workshop on Hot Topics in Operating Systems (HotOS X)

June 12–15, 2005, Santa Fe, NM, USA
http://www.usenix.org/hotos05
Paper titles and abstracts due: February 1, 2005
Full paper submissions due: May 2, 2005

## 14th USENIX Security Symposium

August 1–5, Baltimore, MD, USA
http://www.usenix.org/sec05
Paper submissions due: January 18, 2005

# contents




;login:

**VOL. 29, NO. 5, OCTOBER 2004**

EDITOR
Rob Kolstad
kolstad@usenix.org

CONTRIBUTING EDITOR
Tina Darmohray
tmd@usenix.org

MANAGING EDITOR
Jane-Ellen Long
jel@usenix.org

COPY EDITOR
Steve Gilmartin
proofshop@usenix.org

PROOFREADER
proofshop
proofshop@usenix.org

TYPESETTER
Star Type

USENIX ASSOCIATION
2560 Ninth Street,
Suite 215, Berkeley,
California 94710
Phone: (510) 528-8649
FAX: (510) 548-5738

office@usenix.org
login@usenix.org
conference@usenix.org

http://www.usenix.org
http://www.sage.org

*;login:* is the official magazine of the USENIX Association.

*;login:* (ISSN 1044-6397) is published bi-monthly by the USENIX Association, 2560 Ninth Street, Suite 215, Berkeley, CA 94710.

$80 of each member's annual dues is for an annual subscription to *;login:*. Subscriptions for nonmembers are $110 per year.

Periodicals postage paid at Berkeley, CA, and additional offices.

POSTMASTER: Send address changes to *;login:*, USENIX Association, 2560 Ninth Street, Suite 215, Berkeley, CA 94710.

ROB KOLSTAD

# motd

Dr. Rob Kolstad has long served as editor of ;login:. He is SAGE's Executive Director, and also head coach of the USENIX-sponsored USA Computing Olympiad.

kolstad@sage.org

EVERY YEAR, SAGE ADMINISTERS A salary survey. This year's survey garnered over 4,000 responses and included an Unemployment Survey for those who were out of work for more than half the year. Some of the statistics turned out to be interesting or counterintuitive, and the respondents' comments make thought-provoking reading.

Of those respondents employed more than six months during 2003, over one in ten (10.9%) was unemployed for at least a week (by simple arithmetic, almost 90% were employed steadily throughout the year).

This year's survey asked respondents what they thought about certifications. We received a number of extremely vehement responses from an obviously vocal group. The question was, "Do you think certifications are a good thing for the profession in general?" 48% answered, "Sometimes, it depends"; 19.1% checked "Rarely, a few are good"; 9.3% checked "No, generally they are worthless." 11.7% thought they were a good thing, and 10.0% checked "Usually, most are pretty good." Tallied up, almost 70% think that certifications have at least a somewhat positive impact.

Subtracting years of experience from a person's age gives a suggestion as to what age they entered the field. Figure 1 shows that over a third of the respondents entered the field when they were 25 or older. Figure 2 shows the bell-curve-like distribution of experience, including the big bubble of people entering the field during the dot-com boom. Mean experience was 8.01 years; median was 5 years.


FIGURE 1

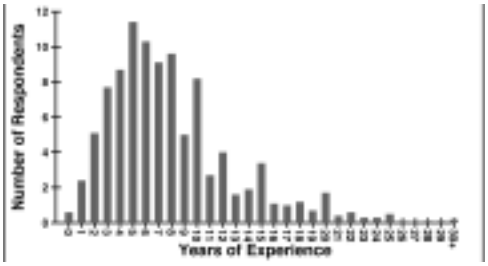


FIGURE 2

Many sysadmins were not educated in a field they thought was relevant (e.g., computer science). Only 7.4% had a bachelor's degree or higher in a relevant field. See Figure 3 for the surprising distribution.

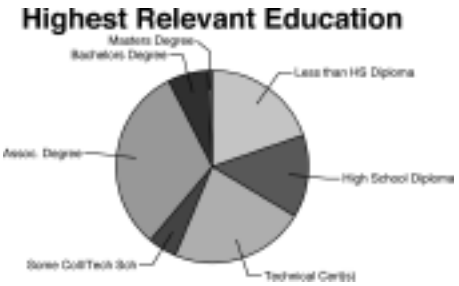


FIGURE 3

74.6% of those responding increased their salary during 2003. The average increase of those who improved their lot was 8.18%; the median was 6.06%. 15.4% saw a 0–0.99% increase, while 10.8% earned less in 2003 than in 2002. See Figure 4 for the distribution. Respondents cited a variety of reasons for their improved salaries. Those items checked by 5% or more of the respondents (each of whom could check as many as three reasons) included:

- Performance
- Goal achievement
- Standard raises/cost-of-living adjustments
- Increased responsibilities
- Hard work
- System stability
- High-profile projects


FIGURE 4

Several respondents were frustrated with the profession, as indicated by their responses in the comments section:

> Future not so bright, shades no longer required.
>
> I'm depressed. I'm now making less than when I first started years ago, as if my career has been wiped away.
>
> Too many people with MCSEs who think they know everything. The management in my department has threatened to replace me with a lower-paying position. The secretaries and administrative support folks make as much money as I do, even though they are not as skilled or required to keep up with technologies. I feel underappreciated as a result.

They also had comments about system administration as a profession.

> UNIX is cool, sysadmin is cool, large-scale corp IT is not. To paraphrase a book I read recently: "IT is not like football. In football the coaches are proud of star players; in IT the managers/execs want to be the star players as well."
>
> Employers need to have a better grasp of exactly what it is that system administrators do and why our job function is so important.
>
> The biggest problem I'm experiencing currently is now that technology is so widespread and that non-IT managers read the latest "trends" in magazines such as *Business 2.0*, many managers are forcing themselves into technology decision-making roles when they have absolutely no competence to do so. This is disrupting workflow, time, and wastes tons of money.

This one was particularly interesting:

> Open Source Software is eroding the desirability and marketability of traditional UNIX sysadmins. Employers are increasingly more reluctant to pay high salaries to manage what they see as a "free" software. As more high-end UNIX installations (enterprise-class servers) are replaced by racks of whitebox Linux servers, companies are looking for cheaper manpower to manage them.

A subsequent query to the sage-members email list revealed that this attitude, while not widespread, does appear in some companies.

Many had advice. This paragraph was illuminating:

> As an industry, we need to clear up the "old school thinking" and make a concerted effort to understand the businesses we work for—try to align strategies, to use technology to best make the business work, etc., but most of all, we need to educate the business about what we do, how we do it, and why, and to earn mutual trust and respect so that they direct us in what they want achieved and let us decide how to do it.

All in all, system administration doesn't appear to be doing tremendously badly. Nevertheless, the same complaints about misunderstood job functions continue to appear year after year.

ROB KOLSTAD

# USENIX Haiku Contest

Dr. Rob Kolstad has long served as editor of ;login:. He is SAGE's Executive Director, and also head coach of the USENIX-sponsored USA Computing Olympiad.

kolstad@sage.org

I WAS PRIVILEGED TO JUDGE THE HAIKU contest at the recent USENIX Annual Technical Conference (see p. 41 for the conference summaries). Attendees were asked to build haiku from any form of the words participating attendees wrote on a ribbon attached to their badge, plus articles, pronouns, and conjunctions.

Here are the best (according to me) of the dozens submitted. Thanks to all the entrants, and congratulations to these published poets.

David R. Linn wrote of the Joy of Programming:

> Even an old tool,
> in the hands of a hacker,
> can do cool new things

Presumably, he meant the good kind of hacker.

Larry Snider played on some advertising:

> Life in the cold room
> Maintaining and defending
> "We" are the network

Kenneth Geisshirt was among several who advocated a favorite technology:

> All those penguins
> Drinking free beers with Beastie
> Making a better world

Likewise Carla Austin-Silvani:

> A cynic is more
> likely to panic than the
> Andrew File System

Peter Mager wrote of grimmer technology happenings, though with a happy ending:

> SCO sells out
> Lawyers threaten all progress
> But truth wins at last

Two entrants wrote of the bleak world of spam. Miranda Mowbray's entry echoed Monty Python's sentiments with some extremely colorful imagery:

> All your email drowns
> in a bright pink foaming sea
> spam spam spam spam spam

while Justin Ferguson's entry was a bit more frustrated:

> Junk email rise brings
> The eternal September—
> The Internet's dead.

Charles Peterman summed up the cultural Zeitgeist of some aspects of system administration in just seventeen syllables:

> Raid driver crashes
> Filenames transform yet again
> Device Driver Dead

Finally, Ming Chow's joyous ode to the USENIX conference itself:

> USENIX in Boston
> Hackers and experts unite
> Technology rules!

# letters to the editor

**TO ROBERT HASKINS:**

I read your article in the August '04 *;login:*. I have some comments.

Under your discussion of SPF, in the last paragraph, you state, "The real way to fix email is to replace RFC 822 with a more secure protocol." Of course, RFC 822 doesn't specify a protocol, it specifies a message format. Also, as I'm sure you're aware, this protocol has been updated. I think citing RFC 2821 would have been more appropriate.

Parenthetically, I think replacing SMTP is unrealistic. We can't just have a "Flag Day" like this, which means there would have to be "translators" between the old and new protocols. These would continue to pass spam until they all went away. When would they go away? My guess is at about the same rate that SPF would be adopted.

Also, you state that SPF doesn't solve the "spam zombie" problem. It does prevent spam zombies (and anyone else) from forging another's domain without their explicit or implicit consent. I believe this would be a big help.

I also don't see why SPF is more likely to be adopted by the big email hosting providers than the small ones. Small ones probably have more to gain ('cuz they can less afford to be joe jobbed), and it's not that much work. It's less likely to be adopted by very decentralized domains, but that's an issue independent of size. I've published SPF records for a tiny number of email subscribers, and I know I'm not alone.

While SPF is imperfect (you're absolutely correct that email forwarding is problematic), I still believe it's a good step forward. It's designed to stop "joe jobbing," and it can succeed in doing that if listening to SPF records is widely adopted. The uptake on publishing SPF records can be slow and SPF could still be effective.

Best,

NICK CHRISTENSON
npc@gangofone.com

**ROBERT HASKINS REPLIES:**

You are indeed correct about RFC 822 and RFC 2821. As for my comments on replacing the underlying protocol, of course replacing it is unrealistic. But solving the problem of spam once and for all will require a re-engineering of the SMTP protocol in my (and other people's) opinion. If the rate of spamming keeps increasing at the current pace, we will have no choice but to either throw out the email functionality altogether or do something really drastic like re-engineer the SMTP protocol. Unfortunately, every other anti-spam solution is just a band-aid, and SPF is no exception.

SPF doesn't help the zombie problem directly, but I suppose it does help it indirectly.

As for adoption rates, the providers who host the largest number of mailboxes (and whom the spammers most often use for their joe-job attacks) have the most to gain from SPF. I'm not saying that smaller providers won't have *anything* to gain, just that big providers have *much more* to gain.

As far as SPF record enforcement/listening goes, I think that the MTA/SPF integration needs to get much better before SPF enforcement will be widely adopted. I will admit, I haven't set up an MTA for SPF enforcement. But from a quick perusal of instructions for doing so, it does not look like an easy, straightforward task.

My point wasn't to discount SPF totally but, rather, to provide an attempt at a balanced introduction to SPF. I don't think I achieved my goal!

ROBERT HASKINS
rhaskins@cnetwork.com

STEPHEN B. JENKINS

# offline programmatic generation of Web pages

Stephen is the senior programmer/analyst at the Aerodynamics Laboratory of the Institute for Aerospace Research, National Research Council of Canada. For more information, see http://www.erudil.com.

Stephen.Jenkins@nrc-cnrc.gc.ca

AS PROGRAMMERS, WHEN WE NEED to provide Web-accessible information, two methods usually come to mind: a static one—creating Web pages in an editor or Web development tool, and a dynamic one—creating CGI programs to generate HTML. There is, however, a third, often overlooked, option: offline programmatic generation of Web pages (OPG). By OPG, I mean writing programs to generate HTML documents at the time and location of your choosing, as opposed to CGI programs, where the pages are generated at access time on the computer hosting the Web server.

## When to Use OPG

While it may appear, at first glance, that OPG has little to offer over the other two methods, this is not the case. Its primary advantage is that complex HTML documents can be quickly and easily modified, without the need for CGI programs. This is an absolute necessity for people using the services of many of the largest ISPs, since those companies typically only provide a small number of "canned" CGI scripts (e.g., formmail) and do not allow user-written programs.

Even if you do have complete access to your Web server, OPG offers a significant benefit in performance: Web pages can be generated at times when CPU and IO loads are low. This can be especially significant for large Web pages that take a considerable time to generate, such as log file summaries. Rather than create the documents on demand, as CGI programs do, the pages can be generated once (e.g., each night) via a crontab entry. This is also useful for information that is rarely modified (staff email address, phone lists, etc.). The Web pages only need to be generated as often as the data changes.

The third place that OPG makes sense is for pages containing large/multiple tables of data. Even if the information is allegedly unchanging (we've all heard that before!), creating and modifying large tables by hand can be tedious and error prone. Also, as programmers, many of us would rather spend the time writing code to perform a task rather than do it manually.

One final issue is security. While CGI programs can be made as secure as any other software on the Net, inexperienced coders can inadvertently leave themselves open to malicious attacks. For the wary (and the

downright paranoid), OPG offers many of the benefits of CGI, while avoiding all of the potential risks.

## A Simple Example

By way of example, I thought I'd show a simplified version of a program I wrote for my wife, Christine, who gives private music lessons. She needed a way to display her schedule and to show which lesson times were available to potential students visiting her Web site. Since her HTML skills are rudimentary and her ISP has difficulties with custom CGI programs, I decided to write some Perl code to generate the Web pages on our home computer. When her timetable changes, Christine modifies the data, double-clicks the program's icon, and then uses a graphical FTP program to upload the newly created Web pages to her service provider.

As so often happens with these kinds of small projects, I decided to add features after I started writing the program. Rather than just generate a public Web page showing the available time slots, I decided to have the program also generate a private page to show such information as the student's initials, other musical commitments, and time off. To keep things as simple and compatible as possible, I decided to put the schedule information in a "DATA" segment at the end of the program, and chose not to use external Perl modules or Cascading Style Sheets (CSS). Figures 1 and 2 show the public and private Web pages generated by the example program shown in Listing 1.


FIGURE 1 Public Schedule Showing Only Available Time Slots

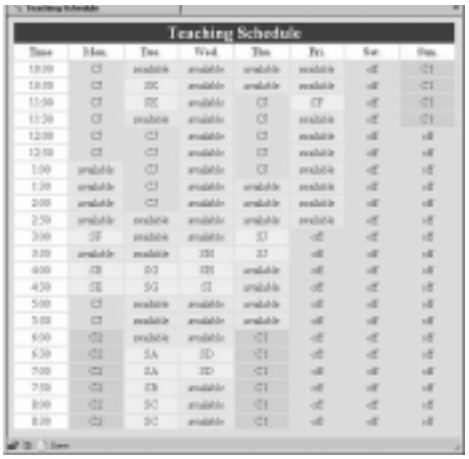


FIGURE 2 Private Schedule Showing All Information

```
#!/usr/bin/perl
use strict;
use warnings;

my $title  = 'Teaching Schedule';
my $colwidth = 'width=75';
my %colorfor = ( 'bg'    => '"#D8E8D8"',
          'hddark'  => '"#336666"',
          'hdlight' => '"#FFFFFF"',
          'choir'   => '"#FFCCCC"',
          'student' => '"#CCFFCC"',
          'avail'   => '"#FFE7CC"',
         );

my $html1 =<<EOF;
<html><head><title>$title</title></head>
<body bgcolor=$colorfor{'bg'} text=$colorfor{'hddark'}>
<table border=0>
<tr><th colspan=8 bgcolor=$colorfor{'hddark'}>
<font color=$colorfor{'hdlight'} size="+2">$title</font></th></tr>
<tr>
EOF

foreach ( qw( Time Mon. Tue. Wed. Thu. Fri. Sat. Sun. ) ) {
  $html1 .= "<th $colwidth bgcolor=$colorfor{'hdlight'}>$_</th>";
}
$html1 .= "</tr>\n";

my $pri = "";
my $pub = "";

while( <DATA> ) {
  next unless /^\d/;

  my $time = substr($_, 0, 6, "");
  $pri .= "<tr><td bgcolor=$colorfor{'hdlight'} align=\"center\">$time</td>";
  $pub .= "<tr><td bgcolor=$colorfor{'hdlight'} align=\"center\">$time</td>";

  my @days = /.{1,4}/g;
  @days = splice @days, 0, 7;
  foreach ( @days ) {
    my $bgc = $colorfor{'avail'};
    if( /\S/ ) {
      if( /CJ|off/ ) { $bgc = $colorfor{'bg'};  }
      elsif( /C\d/ ) { $bgc = $colorfor{'choir'};   }
      else      { $bgc = $colorfor{'student'}; }
      $pri .= "<td bgcolor=$bgc align=\"center\">$_</td>";
      $pub .= "<td></td>";
    } else {
      $pri .= "<td bgcolor=$bgc align=\"center\">available</td>";
      $pub .= "<td bgcolor=$bgc align=\"center\">available</td>";;
    }
  }
  $pri .= "</tr>\n";
  $pub .= "</tr>\n";
}

my $html2 = "</table></body></html>\n";

open  PRI, ">private.html" or die "Oops: $!";
print PRI  "$html1$pri$html2\n";
close PRI;
open  PUB, ">public.html" or die "Oops: $!";
```

```
print PUB  "$html1$pub$html2\n";
close PUB;

__DATA__
Time   Mon  Tue  Wed  Thu  Fri  Sat  Sun
10:00  CJ                      off  C1
10:30  CJ   SK                 off  C1
11:00  CJ   SK        CJ   SF   off  C1
11:30  CJ             CJ        off  C1
12:00  CJ   CJ        CJ        off  off
12:30  CJ   CJ        CJ        off  off
1:00        CJ        CJ        off  off
1:30        CJ                  off  off
2:00        CJ                  off  off
2:30                            off  off
3:00   SF             SJ   off  off  off
3:30             SH   SJ   off  off  off
4:00   SE   SG   SH        off  off  off
4:30        SE   SG   SI   off  off  off
5:00        CJ             off  off  off
5:30        CJ             off  off  off
6:00   C2             C1   off  off  off
6:30   C2   SA   SD   C1   off  off  off
7:00   C2   SA   SD   C1   off  off  off
7:30   C2   SB        C1   off  off  off
8:00   C2   SC        C1   off  off  off
8:30   C2   SC        C1   off  off  off
```

LISTING 1 Schedule Web Page Generator

The first three lines start the program by invoking Perl with the `strict` and `warnings` pragmas. The next nine lines set up some HTML parameters for later use. After that, a "here document" is used to define the HTML head and title elements, as well as setting up the beginning of the main table that will hold the schedule. Next, a `foreach` statement is used to create the table header entries. Until this point, the HTML code has been common to both the private and public pages, but now we need to define two scalars to hold the HTML elements that are unique to each.

At this point, I'll jump down to the end of the program to describe the DATA segment. It consists of a header line followed by multiple lines of data in a simple table, one time slot per row and one day per column. Available slots are expected to be blank, but occupied slots are expected to contain either the string "off", a "C" followed by a number for a choir, or a set of initials, set here to "SA" to "SJ" to denote 10 students or to "CJ" to denote Christine.

A `while` loop reads the rows of the DATA block. Lines that do not begin with a numerical digit are ignored (to allow for blank, formatting, or comment lines). The first six characters are removed and are used as the time string for this row of the HTML table. Next, a regex breaks the rest of the line up into an array of four-character strings. The `@days` array is truncated to seven elements, just in case some extra characters were placed in the data by accident. A `foreach` loop examines each day's entry, and the table cell background color is set to the "available" color. If the entry contains non-whitespace characters, it is compared against two regexes to determine the appropriate background color. For the private page, the schedule's data is written into the table cell, but the public page's cell is left empty to show an unavailable time slot. If the day's entry was only whitespace, the table cells of both private and public pages are set to "available". The table row tags are then closed, and the while loop repeats until all of the

time periods have been processed. In the final few lines of the program, the HTML closing tags are added, and the files are written to disk.

## Other Examples

I've used OPG several other times in the past year or so; once was for a local minor hockey league. They wanted to put player statistics (goals, assists, goals per game, points per game, etc.) on their Web site. I wrote two small Perl programs: one for the goalie stats and one for the other players. The data comes from tab-separated text files maintained by one of the league officials. He updates the data files on his home PC and then runs the programs to calculate the stats, sort the player rankings, and generate the Web pages. He then uploads the HTML documents to the league Web server.

As part of my job at the Aerodynamics Laboratory, I've created an event-logging system that receives and records software events from a number of independent computers and stores them in a log file similar to the type used by the Apache Web server [Jenkins]. In order to provide an executive summary of the events that occurred during the previous day, week, month, etc., I wrote a Perl program that reads and analyzes the log files, then generates two Web pages each day: a daily summary and an index of all of the available daily summaries. Since this program can take several minutes to run, and heavily exercises the hard drives, I didn't want it to run on-demand as a CGI would. Instead, I set up a crontab entry to run the process daily (shortly after midnight) and place the output pages in a Web-accessible location.

For a final example, I'd like to talk about three somewhat more complex programs that I wrote for the Fifth International Colloquium on Bluff Body Aerodynamics & Applications (BBAA V). As anyone who's ever been on the organizing committee for a conference will tell you, many of the meeting details (the list of papers, the paper titles, the authors, and the program schedule, just to name a few) are far from static. I wrote Perl programs to take information from the master "database" (an Excel spreadsheet that I read using `Spreadsheet::ParseExcel::Simple`) and generate a list of presentations by topic, an author's index, and the daily program schedules (see Figure 3 for a facsimile).


FIGURE 3 A Portion of the Conference Presentation Program

As new information arrives, the conference administrative assistant updates the spreadsheet and forwards me a copy via email. Within minutes, I can run my programs, update the Web site, and reply that the changes have been made public.

## Concluding Remarks

One of the beauties of such a simple concept is that it is completely OS-, Web server-, and Web browser-independent. While OPG is also independent of implementation language, it definitely works best with a language such as Perl, which was designed to manipulate text. Perl also has many other benefits such as low/no cost (open source), portability (available for most modern OSes), easy access to databases through DBI, and a large, freely available archive of modules in CPAN.

For me, the best things about offline programmatic generation are summed up in the first two of the three great virtues of a programmer: laziness and impatience. It enables me to be lazy because with only a few hours' work, I can write a program that enables nonprogrammers to maintain their complex Web pages, putting the responsibility on them. It enables those users to be impatient because rather than wait for someone else to update a Web site, they can do it themselves, immediately. If necessary for the most naïve users, I could even automate the FTP process using `Net::FTP`. This means that to modify their Web site, they would only have to update their text files, Excel spreadsheets, or database entries, and double-click a desktop icon on their PCs.

REFERENCE

Jenkins, S.B., "A Web-Based Environment to Support Aerodynamic Testing," *IEEE Aerospace and Electronic Systems Magazine* 19:1 (January 2004), p. 3.

CLIF FLYNT

# the tclsh spot

Clif Flynt is president of Noumena Corp., which offers training and consulting services for Tcl/Tk and Internet applications. He is the author of Tcl/Tk: A Developer's Guide and the TclTutor instruction package. He has been programming computers since 1970 and a Tcl advocate since 1994.

clif@cflynt.com

THE PAST FEW TCLSH SPOT ARTICLES described a software architecture in which the main-line code is compiled, and then Tcl/Tk scripts are invoked to read input and display results. A major strength of this architecture is that the look and feel of an application can be drastically modified without touching a line of compiled code.

An old FORTRAN TTY-based program, the Lunar Lander, was modified to run in real time, with a display that looks like this:

This article will show how to replace that GUI with one that shows the state of the lander graphically at each second of the descent.

The first few seconds of a landing on Earth might look like this:

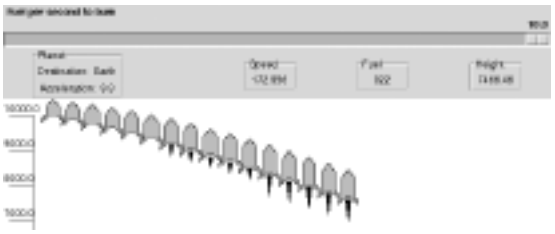

The numeric values are displayed using the Tk label widget, as in the previous GUI. The area below the labels is a graph displaying the height of the lander on the Y axis, and time into the landing on the X axis. The points on the graph are drawn with rockets.

The rockets show:

- The current height of the rocket (the Y axis of the graph)
- The number of seconds into the landing (the X axis of the graph)
- The current speed (the height of the rocket)
- The remaining fuel (the width of the rocket)
- The amount of fuel burned in this timestep (the height of the flame below the rocket)

You can see in the display that the rocket falls a little faster each second when there is no thrust applied. With half-thrust it still accelerates, while reducing the amount of fuel. At full thrust, the speed is close to constant (in fact, there is a very slight negative acceleration), and the fuel is consumed more rapidly. This is

not as extensive as Charles Joseph Minard's graph of Napoleon's Russian campaign, but it does convey five dimensions of data in two dimensions.

The Tk canvas widget makes this sort of data representation easy. The Tk canvas widget is an object-based drawing surface inspired by Joel Bartlett's ezd program. It enables the programmer to define a window into an arbitrarily large drawing surface and place graphic objects on that surface. The graphic objects are each defined by a location and a set of configuration options specific to the type of object.

For example, a text object can be configured to display a certain set of text in a certain font, while a rectangle object can be assigned height, width, and colors.

Once an object is created, it can't change type, but its location and configuration options can be modified as necessary. Thus, the words displayed by a text object can be changed, as can the font or color.

A new canvas widget can be created and displayed with the same syntax as other Tk widgets. The command is `canvas`, followed by the name of this canvas, followed by a list of key/value option pairs. The newly created canvas is displayed using either the `pack`, `place`, or `grid` geometry manager.

**Syntax:** canvas canvasName ?options?

canvasName — The name for this canvas

?options? — Some of the options supported by the canvas widget are:

`-background` *`color`*
: The color to use for the background of this canvas. The default color is light gray.

`-height` *`size`*
: The height of the displayed portion of the canvas. If `-scrollregion` is declared larger than this, and scrollbars are attached to this canvas, this defines the height of the window into a larger canvas.
The `size` parameter may be in pixels, inches, millimeters, etc.

`-width` *`size`*
: The width of this canvas widget. This may define the size of a window into a larger canvas.

`-scrollregion` *`{left top right bottom}`*
: The region of a larger canvas for the window to scroll over. These coordinates define the area of a canvas that can scroll into view when the canvas is attached to a scrollbar widget.

Like other Tk widget creation commands, the `canvas` command returns the name of the canvas it created, and also creates a new procedure by that name to use to interact with the canvas.

For example, this command:

```
set cvs [canvas .c -height 500 -width 700 -background white]
```

creates a new canvas named `.c` and a new procedure named `.c` and assigns the value `.c` to the variable `cvs`.

The new `.c` procedure supports several subcommands, including:

`create`
: Create a new object on the canvas. Returns a unique ID for the new object.

`configure`
: Query or set canvas configuration options.

`itemconfigure`

Query or set configuration options for an item on the canvas.

`xview`

Define the window into a canvas to be displayed.

`bbox`

Returns the bounding rectangle that encloses a set of canvas items. The command *`canvasName`* `bbox all` returns the bounding rectangle that includes all items displayed in a canvas.

`bind`

Assign a binding to an item on the canvas.

The most used canvas command is the `create` command. The syntax for this command is:

Syntax: *`canvasName`* `create` *`itemType coords ?options?`*

`itemType`

The type of item to create may be `arc`, `bitmap`, `image line`, `oval`, `polygon`, `rectangle`, `text`, or `window.`

`coords`

The coordinates for this item. The coordinates are X/Y pairs. All `itemTypes` require at least one X/Y pair. Some `itemTypes` (ovals, rectangles) require two pairs to define the opposing corners of a bounding rectangle for the object. Lines and polygons can have multiple X/Y pairs to define the corners of the figure.

The `canvas` coordinate system places 0,0 on the upper left corner. X values increase toward the right, and Y values increase toward the bottom.

`options`

Keyword/value pairs to define configuration options for a graphic item. The supported keywords are different for different types of graphic items.

The axes for a graph can be created with line and text objects. The commands to create a canvas, short horizontal line, and text would look like this:

```
set cvs [canvas .c -height 500 -width 700 -background white]
$cvs create line 10 10 20 10
$cvs create text 25 10 -text "10000"
```

By adding a couple of loops, the code to draw the X and Y axes and ticks looks like this:

```
  # Create the axis lines
  .c create line 40 20 40 480
  .c create line 40 460 4000 460
  # Create and label the ticks on the Y axis
  for {set i 0} {$i <= $height}\
       {set i [expr {$i + $height/10}]} {
    set y [expr {460 - ($i * .044)}]
    .c create text 3 $y -text $i -anchor sw
    .c create line 10 $y 40 $y
  }
  # Create and label the ticks on the X axis
  for {set i 40; set j 0} {$i < 4000 }\
       {incr j; incr i 25} {
```

```
    .c create text $i 482 -text $j -anchor nw
    .c create line $i 460 $i 475
  }
```

The horizontal X axis line goes from pixel 40 to pixel 4000. That's a bit longer than most monitors. The canvas widget has commands that make it easy to attach the canvas to horizontal and vertical scrollbars.

Create a scrollbar with the `scrollbar` command:

`Syntax: scrollbar` *`scrollbarName ?options?`*

`scrollbar` Create a `scrollbar` widget.

*`scrollbarName`* The name for this `scrollbar`.

*`options`* This widget supports several options. The `-command` option is required.

`-command` *`"procName ?args?"`*
This defines the command to invoke when the state of the `scrollbar` changes. Arguments that define the changed state will be appended to the arguments defined in this option.

`-orient` *`direction`*
Defines the orientation for the scrollbar. The *`direction`* may be `horizontal` or `vertical`. Defaults to vertical.

`-troughcolor` *`color`*
Defines the color for the trough below the slider. Defaults to the default background color of the frames.

The `wish` interpreter handles the interaction between the `canvas` and `scrollbar` by registering a callback procedure with the `scrollbar` and `canvas` widgets. Whenever one of these widgets changes state, it will evaluate the registered script to update the other widget.

The code to create and display a canvas and scrollbar resembles this:

```
canvas .c -height 500 -width 700 -background white \
  -xscrollcommand {.sb set}
scrollbar .sb -orient horizontal -command {.c xview}
grid .c -row 4 -column 1
grid .sb -row 5 -column 1 -sticky ew
```

The canvas `xview` and `yview` subcommands are invoked by a scrollbar when it changes state (e.g., a user drags the slider). The scrollbar `set` command is invoked by the canvas when it changes state.

After creating the canvas and drawing the X and Y axes with the code above, the canvas widget would show the X axis from pixel 0 to 700, and the scrollbar would have a slider that extended from edge to edge.

By default, the canvas widget is set to scroll for the width of the canvas, i.e., not to scroll at all.

The `-scrollregion` option defines the portion of the canvas that can be scrolled into. The argument to the `-scrollregion` option is a list of the left, top, right, and bottom coordinates of the rectangle that can be scrolled into. Adding the command

```
.c configure -scrollregion {0 0 4000 500}
```

after drawing the axes causes the canvas to change its state, which invokes the scrollbar's `set` command to make the slider show that the leftmost portion of a larger window is being displayed.

In this case, we know the area of the canvas we need to scroll around in. If your application is creating objects without knowing the boundaries, the `bbox` command can be used to find the bounding rectangle.

Syntax: *`canvasName`* `bbox` *`tagOrId`*
`bbox`

Return the coordinates of a box that would enclose the item, or items with the same tag.

*`tagOrId`*

A tag or unique ID that identifies the item. If a tag is used, and multiple items share that tag, then the return is the bounding box that would cover all the items with that tag. The tag `all` can be used to return the bounding box for all items on a canvas.

The `bbox` command returns the bounding rectangle in the same format that the `-scrollregion` configuration option requires, so a command like:

```
$cvs configure -scrollregion [$cvs bbox all]
```

can be used to set the scrollregion for a canvas without tracking where graphic items have been placed.

The `xview` and `yview` commands can be used to bring newly placed objects into view.

Syntax: *`canvasName`* `xview moveto` *`fraction`*
*`fraction`*

The fraction of the scrollwindow to place to the left of the leftmost edge of the viewed window. A value of 0 will display the leftmost area of the scrollregion, while a value of 0.5 would not display the left half of a scrollregion.

As the Lander program runs, it will start drawing rockets to the right of the 700 pixels that are displayed by default. The user can scroll to the latest rocket, but it's friendlier if the application automatically scrolls to display the latest output. The application can display the latest rocket, and the previous 20 rockets (500 pixels) with this command:

```
.c xview moveto [expr { ($x-500.0) / 4000.0}]
```

The next step is to draw the rockets. The rockets are created with the `create polygon` subcommand. The `create polygon` command can accept an arbitrary number of X/Y pairs to define the nodes on the polygon. The syntax looks like:

Syntax: *`canvasName`* `create polygon` *`coord ?option value? coord`*

A list of X/Y pairs to define corners of the polygon.

*`?option value?`*

Keyword/Value pairs that define configuration options for this polygon. Options include:

`-fill` *`color`*

A color to fill the polygon.

`-outline` *`color`*

A color for the line outlining the polygon.

`-width` *`distance`*

A value for how wide to make the outline. By default this is a number of pixels, but it can also be defined in points, inches, millimeters, or centimeters.

The mainline FORTRAN code calls a Tcl procedure named `showState` to display the current lander status. The `showState` procedure is passed the time, height, speed, and remaining fuel of the rocket. The time and height define the X and Y coordinates, and the speed and remaining fuel can be scaled to define the height and width of the rocket.

With these values, the position of each node of the polygon describing the rocket can be calculated.

The command for performing arithmetic operations in Tcl is the `expr` command. The `expr` command takes an arithmetic expression as an argument and returns a numeric result. For most Tcl applications, commands this verbose are a bit unwieldy, but not a serious problem.

```
set y2 [expr {$y - ($speed / 10)}]
set wid [expr {2 + $fuel/150.0}]
set tall [expr abs($speed)/5.0]
```

However, for calculating each node on a polygon or line, the successive `expr` commands can make the application overly verbose and difficult to read.

Lars Hellstrom described an elegant solution to this on the Tclers' Wiki (http://wiki.tcl.tk/8389). Since any string can be the name of a Tcl procedure, we can define procedures named "+" and "-", to return simple arithmetic operations. The code to do this looks like this:

```
proc + {a b} {
    return [expr $a + $b]
}
proc - {a b} {
    return [expr $a - $b]
}
```

With this trick, describing the outline of the rocket looks like this:

```
set rocketID [.c create polygon $x $y \
    [+ $x $wid] [+ $y 10 ] \
    [+ $x $wid] [+ $y $tall] \
    [+ $x (5+$wid)] [+ $y ($tall+4)] \
    [+ $x (5+$wid)] [+ $y ($tall+9)] \
    [+ $x ($wid-5)] [+ $y $tall ] \
    [- $x ($wid-5)] [+ $y $tall ] \
    [- $x (5+$wid)] [+ $y ($tall+9)] \
    [- $x (5+$wid)] [+ $y ($tall+4)] \
    [- $x $wid] [+ $y $tall ] \
    [- $x $wid] [+ $y 10] \
    -fill $fill -outline black]
```

A failing of this display is that it doesn't show the exact height, speed, fuel burned at any given time. This failing can be solved with the canvas `bind` command. The `bind` command lets us put a binding for some action on a graphic item.

We can add a binding on each rocket so that when the rocket is clicked, it will display the speed, altitude, remaining fuel, and burn.

Clicking the rocket at X location 69 would generate this display:

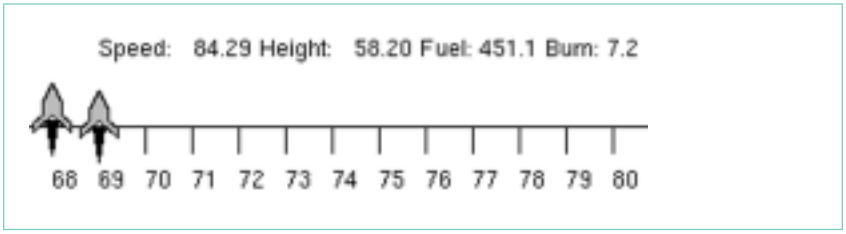


The `bind` command links an event to a widget and a script. If the event occurs while the focus is on that widget, the script associated with that event will be evaluated. Each object on the canvas can also be bound to certain actions such as having a cursor pass over the object or a button clicked while the cursor is over the object.

Syntax: *canvasName* `bind` *tagOrID eventType script*

*tagOrID*

The tag or ID number of the canvas item to have this action bound to it.

*eventType*

The event to trigger this action. Events can be defined in one of three formats:

*alphanumeric*: A single printable (alphanumeric or punctuation) character defines a KeyPress event for that character.

<<virtualEvent>>: A virtual event defined by your script with the event command.

<modifier-type-detail>: This format precisely defines any event that can occur. The fields of an event descriptor are the X windows codes (e.g., Button1 or B1).

*script*

The script to evaluate when this event occurs and the cursor is over a canvas item with this tag or ID.

The line below shows a `bind` command that links the rocket that was just drawn to a button click event. When the left mouse button is clicked over the rocket, the Tcl interpreter will evaluate the procedure `showDetails` with an X and Y location and the speed, altitude, fuel, etc. The script can be any valid Tcl script. In this case, the `showDetails` procedure is a user-provided script that clears an old information line and then creates new text on the canvas.

```
.c bind $rocketID \
"showDetails $x [- $y 30] $speed $fuel $ht $burn"
```

Notice that the script is enclosed within quotes, not curly braces. Tcl will perform substitutions on lines enclosed within quotes, while the curly braces will force the substitution to be delayed until the script is evaluated.

In this case, we want the speed, fuel, etc., to be substituted when the rocket is drawn and the binding is created. If we needed to have the substitution done when the event occurred (for instance, if the diagram were a real-time monitor for our network), we'd enclose the script in curly braces, and the variables would be substituted when the object is clicked.

The one problem with a FORTRAN mainline linked to a Tcl/Tk library is that it requires that the user have Tcl/Tk installed in order to run the FORTRAN program. The next Tclsh Spot article will explain how to use TOBE to make a single, stand-alone executable that you could ship to a customer who had never heard of Tcl/Tk.

GLEN MCCLUSKEY

# making use of c# collections

Glen McCluskey is a consultant with 20 years of experience who has focused on programming languages since 1988. He specializes in Java and C++ performance, testing, and technical documentation areas.

glenm@glenmccl.com

THE TERM "COLLECTION" IN MODERN programming languages typically refers to a group of elements, such as integers or strings, organized in a list or table of some sort. Languages such as C++ and C# provide library facilities for managing collections. In this column we'll look at C# collections and illustrate some of the techniques you can use in your applications.

## Managing Lists of Integers

To help tie down the idea of what a collection is, consider a two-part example. Suppose that you're doing some C# programming, and you need to keep a list of integer values around. Your first solution to this problem looks like so:

```
using System;
public class ListDemo1 {
  static int[] list = new int[10];
  static int currlen = 0;

  static void add(int elem) {
    if (currlen == list.Length) {
      int[] newlist = new
      int[(int)(currlen * 1.5)];
      for (int i = 0; i < currlen;
      i++)
        newlist[i] = list[i];
      list = newlist;
    }

    list[currlen++] = elem;
  }

  public static void Main(string[]
  args) {
    for (int i = 1; i <= 20; i++)
      add(i);

    Console.WriteLine("length = " +
    currlen);

    Console.Write("elements = ");
    for (int i = 0; i < currlen; i++)
      Console.Write(list[i] + " ");
    Console.WriteLine();
  }
}
```

The idiom used in this example is a common one. An array is allocated, and values are inserted into it. When the array overflows, a new array is allocated, and the contents of the old array are copied to it.

This code works, is efficient, and is easy to follow. Unfortunately, however, you may end up re-solving

this problem over and over again in your applications, possibly introducing errors along the way. For example, suppose that if instead of:

```
if (currlen == list.Length)
  ...
```

I'd said:

```
if (currlen > list.Length)
  ...
```

The code in this case might work most of the time, but occasionally it would trigger an exception caused by going off the end of the array.

An alternative is to use the C# collection class `ArrayList`, like this:

```
using System;
using System.Collections;

public class ListDemo2 {
  public static void Main(string[] args) {
    ArrayList list = new ArrayList();

    for (int i = 1; i <= 20; i++)
      list.Add(i);

    Console.WriteLine("length = " + list.Count);

    Console.Write("elements = ");
    foreach (int i in list)
      Console.Write(i + " ");
    Console.WriteLine();
  }
}
```

All the details are handled automatically for you when using `ArrayList`. The list is grown automatically, there is a standardized interface, and so on. Most of the time, such convenience is exactly what you want, but, of course, the default handling may be different from what you as a programmer might expect. For example, in the first demo above, we created the integer list with room for 10 elements, and increased its size by 50% when needed. In contrast, `ArrayList` starts with a capacity of 16, and doubles the size when the list is internally reallocated.

An object of type `ArrayList` contains an internal array, which holds the actual elements. Strictly speaking, our example does not create an `ArrayList` of integers but, rather, an `ArrayList` of references to wrapper objects, each of which contains an integer. In C#, elements of primitive type are automatically converted ("boxed") to object type as needed, and then can be converted back ("unboxed") using a cast.

Note also the use of the "foreach" statement, a convenient shorthand for iterating across the elements of a collection.

## Sorting

Suppose that we want to go further with our example and sort the list of integers into descending order. To do so, we make use of a comparer class that implements the standard interface `IComparer`. The code looks like this:

```
using System;
using System.Collections;
```

```
public class MyComparer : IComparer {
  public int Compare(object aobj, object bobj) {
    int a = (int)aobj;
    int b = (int)bobj;
    return a < b ? 1 : (a == b ? 0 : -1);
  }
}
public class SortDemo {
  public static void Main(string[] args) {
    ArrayList list = new ArrayList();
    for (int i = 1; i <= 20; i++)
      list.Add(i);
    list.Sort(new MyComparer());
    Console.WriteLine("length = " + list.Count);
    Console.Write("elements = ");
    foreach (int i in list)
      Console.Write(i + " ");
    Console.WriteLine();
  }
}
```

The `ArrayList` sort method must have some standardized way of calling back to a comparison method. If a class implements a specific interface, such as `IComparer`, then that class is guaranteed to have a method `Compare(object,object)` defined in it, and the sort routine can call this method through the passed-in `MyComparer` object.

## Making Copies of Collection Objects

We said above that an `ArrayList` of integers is really an array of references to integer wrapper objects. Such internal details matter in a case where we make a copy of a collection. For example, suppose that we have an `ArrayList`, and we add an object to it, and then copy the list, and then modify the object previously added.

We can observe what happens by running this code:

```
using System;
using System.Collections;
public class MyClass {
  private int val;
  public MyClass(int i) {
    setVal(i);
  }
  public void setVal(int i) {
    val = i;
  }
  public override string ToString() {
    return val.ToString();
  }
}
public class CloneDemo {
  public static void Main(string[] args) {
    ArrayList list1 = new ArrayList();
```

```
      MyClass obj = new MyClass(10);
      list1.Add(obj);
      Console.WriteLine(list1[0]);

      ArrayList list2 = (ArrayList)list1.Clone();
      obj.setVal(20);
      Console.WriteLine(list2[0]);
    }
  }
```

When the list is cloned, a shallow copy is made, that is, a new internal array is allocated, and the object references are copied to the new array. If an object is modified after creation, then modifying the original object will modify the copy as well, because there are multiple references to the same object.

The output of the program is:

```
10
20
```

and we observe that modifying the object does indeed have effect in the list copy.

## BitArrays and Iteration

Let's go on and look at another collection class, used to manage lists of bits. In this example we create two `BitArrays`, set some of the bits, and then perform an XOR operation:

```
using System;
using System.Collections;

public class BitDemo {
  public static void Main(string[] args) {
    BitArray ba1 = new BitArray(8);
    ba1.Set(0, true);
    ba1.Set(2, true);

    BitArray ba2 = new BitArray(8);
    ba2.Set(1, true);
    ba2.Set(3, true);

    ba1.Xor(ba2);

    for (int i = 0; i < ba1.Count; i++) {
      if (ba1.Get(i))
        Console.Write(i + " ");
    }
    Console.WriteLine();

    IEnumerator e = ba1.GetEnumerator();
    while (e.MoveNext())
      Console.Write(e.Current + " ");
    Console.WriteLine();
  }
}
```

There are a couple of ways to display the result of the XOR. The first looks at all the bits, and prints the integer index for each one that is set.

The second approach uses an enumerator, a standardized way of iterating across all the values in the collection. When we run this program, the result is:

```
0 1 2 3
True True True True False False False False
```

The enumerator code is generic, and can be used to iterate across other types of collections.

## Hashtables

A final example of C# collections is the use of hashtables. This program illustrates how a phone directory might be built up and then accessed:

```
using System;
using System.Collections;

public class HashDemo {
  public static void Main(string[] args) {
    Hashtable ht = new Hashtable();
    //ht = Hashtable.Synchronized(ht);

    ht.Add("Bill Smith", "123-4567");
    ht.Add("Mary Jones", "234-5678");
    ht.Add("John Davis", "345-6789");

    ICollection keys = ht.Keys;

    foreach (string key in keys)
      Console.WriteLine(key + "   " + ht[key]);
  }
}
```

The `Keys` property is used to obtain a list of all the keys in a hashtable. This list is enumerated and each key looked up in the table using the `[]` operator ("indexer") for the hashtable.

Note that the demo has a line:

```
//ht = Hashtable.Synchronized(ht);
```

If we uncomment this line, the effect is to put a wrapper around the hashtable reference, such that access to the hashtable is synchronized, that is, made thread-safe. By default, collections are not thread-safe, and if you are concerned about such an issue, you need to take the appropriate steps when writing C# code.

C# collections are an easy to use and powerful tool to solve many common programming problems. They are part of the standard C# library, and you can use them in place of writing your own code.

CLAIR W. GOLDSMITH AND ROB KOLSTAD

# a conversation about identity management

Clair W. Goldsmith is Associate Vice Chancellor and Chief Information Officer of the University of Texas System, and chair of the UT System Strategic Leadership Council.

CGoldsmith@utsystem.edu

Rob Kolstad is Editor of ;login:.

kolstad@usenix.org

THIS INTERVIEW WAS CONDUCTED BY Rob Kolstad with Dr. Clair W. Goldsmith on August 16, 2004.

RK: Please tell us a little about yourself and your workplace.

CG: I am the associate vice-chancellor and chief information officer of the University of Texas System. The system comprises 15 institutions, nine of which are general academic institutions and six of which are health institutions, that provide both health care and medical education. Many of the institutions are substantially involved in research as well.

The institutions employ about 81,000 and teach about 169,000 students, spanning the state of Texas in both rural and metropolitan regions.

RK: So you're associated with UT-Austin because you're located in Austin?

CG: I am part of the overall UT System Administration, an overarching organization. Our institutions are budgetarily and otherwise somewhat independent. As such, they have their own personnel offices, their own faculties, their own presidents, and so on. Their budgets are rolled up together and presented to the legislature by the folks in the office where I work. Individual institutions have roughly the same relationship to each other that Pontiac and Chevrolet do within General Motors.

RK: So you work more on overall strategy than tactical or day-to-day operations?

CG: Yes. It is largely a strategic job trying to figure out how to deal with the problems that confront the University, and then leverage the University's assets for the greater good of the greatest number.

My office has a focus that is guided by a governance structure with representatives from all the institutions. They have selected IT security, leveraging UT system-wide buying power to reduce costs, and the system-wide network as my current tasks.

RK: With almost 250,000 constituents, that must be a lot of purchasing power.

CG: Yes, it's a $7.8 billion operation. IT is, of course, a fraction of that, potentially approaching $1 billion.

RK: Let's talk about security. Any specific actions you're currently investigating?

CG: We do a lot of identification of products and services that will help our institutions perform their jobs more effectively and efficiently. We feel that one of the serious problems in the security arena is that we typically do not know precisely with whom we are dealing.

In order to deal with that, we postulate that we need a secure, scalable, standards-based interoperable iden-

tity-management infrastructure. Such a system is not simple! Being infrastructure, in many ways it is not terribly exciting to administrators who might rather purchase the latest and greatest payroll system, or maybe a course management system to meet the faculty's needs. Of course, these are valid wishes, but they all depend upon being able to manage identity.

RK: Is this identity management something like knowing exactly who is sending email (for spam management)? Or identity vis-à-vis cryptography in order to conceal communications? Or identity for purchases and contract signing? What kind of identities are you talking about?

CG: That is part of what we're talking about. We specifically define this arena as being composed of three aspects, one of which is "identity"—the name by which an individual or a service/entity is called.

RK: So you manage identities of non-human things as well?

CG: That's true; we manage the identities of various resources.

But let's simplify the discussion for the moment and talk about identity as it relates to carbon units. Those guys have things like names or social security numbers, both of which are valid identities—and you hear about those being stolen.

RK: Are we talking about the context of electronic communications?

CG: It really turns out that you can't stop there. Identity leads to the second aspect of this, which is authentication: "How certain are we that the identity we have been proffered relates to the particular individual who might be in front of us? By what credential or mechanism is that person authenticated so that the particular name is bound to that particular individual?"

This is different from signatures. It's more like going to the Notary Public and presenting your birth certificate, saying, "I am Joe Blow, and here are the papers from a trusted third party who swears that I am who I say I am." That is authentication.

The word can be used with respect to a document. The document is then "authenticated" because you somehow or other know something about the document, mainly if it is a typewritten document on a piece of paper. Whoever has read it attests in some way that it has not been altered and then signs it.

A digital signature provides assurance that the document is authentic—it hasn't been altered. Whether or not that digital signature is capable of being repudiated depends on whether or not that private key has been compromised and just how sure you are that the person who signed it is, in fact, the person they are supposed to be.

RK: I'm curious. How does one visit a Notary Public, carrying a birth certificate, for example, and then assert one's identity from some document created 15—or even 50—years ago? How can the Notary (or anyone) come to believe that the document relates to me?

CG: You've come to the word we need here: trust. At some point, you have to trust the document. In other words, the person who presented it to you has to have some ability to explain why they have the document and why they are trustworthy.

RK: Do we generally do a good job of this in USA society/culture right now? Passports, driver's licenses, etc.?

CG: Not particularly. Look at the 9/11 guys who went to the Virginia Department of Motor Vehicles and were issued driver's licenses because they had social security numbers. We do not do a good job of it.

RK: Perhaps they deserved a driver's license?

CG: I believe they did, but I recollect that their social security numbers were not valid.

RK: Do you believe that the sort of identity management you're talking about needs to be stronger than a 16-year-old going to the driver's license office and being issued a valid government ID card?

CG: I don't think any DMV ever expected their driver's licenses to be used, for example, for airplane boarding. I would like digital identity management to be at least that good as a lower bar.

RK: There seems to be a difference in impact for some of these identity issues. You or I trying to prove our age to buy a beer is a very different thing from the president of a company signing a contract that obligates his company to many different terms and exchanges millions or billions of dollars.

CG: This is the third aspect, which is "authorization." In other words, what is the identified and authenticated individual authorized to do? The more "valuable" the transactions, the more certain we must be of the identity authentication. That has to do with the level of "assurance," or how strongly we have authenticated the person with the digital credential and their "role."

RK: How do you perform identity management? Is it an embedded chip or perhaps a biometric identification? Or is it more like possession of a USB dongle with your key on it, perhaps enhanced by a password?

CG: Most of the identity management we perform is called "single factor," something you know, like a password, which turns out to be a "shared secret." The "dual factor," which is something people are moving to right now, requires both something one "has" and something one "knows," like a password and a key, e.g., a USB key, that sort of token. It contains the digital credential.

"Three factor" is something you know, something you have, and something you are, which might be a biometric like a fingerprint or a retina scan.

Higher-factor identity management is a longer-term goal right now.

RK: In the universe of almost 250,000 clients, this sounds pretty expensive.

CG: It is expensive. We have institutions that are moving to "two factor" management and have more than one institution using USB tokens. A couple more are using smart cards.

RK: Would I be doing this as a student to, for example, take an exam? Or is it more like something I'd use to cover privacy issues when I get my grades?

CG: Both of those. You might even use a token to get into and out of a dorm (in addition to a physical metal key).

RK: That raises privacy issues of monitoring locations of students, doesn't it?

CG: No, in universities just about everything is an educational record and can't be divulged without the signature of the student.

RK: But you're collecting such data?

CG: Yes.

RK: And it can't be divulged to law enforcement agencies?

CG: It can be subpoenaed through the standard subpoena process. It can't be just "coughed up."

Back to the infrastructure we're discussing, though, the digital identity, the digital credential is just one aspect of it. What do you do with it when you have it?

What you really want to do is go around and nail it to every telephone pole so that if you want to send me email, you can find my credential in the public/private crypto-key sense and encrypt the email and sign it. Then I can find your credential on some telephone pole to assure myself that the mail really is authentic. This is a little bit similar to the PGP scheme.

RK: What are the components of an identification management system that you require to put a scheme like this into practice?

CG: You need a directory service. We're looking at LDAP.

RK: What's the query to LDAP?

CG: It could be a number of different things.

RK: "I'd like a blond with blue eyes . . ."

CG: Absolutely. If she or he has agreed to release that information, we'll cough it up right away.

The second requirement is for a mechanism that can embed something like a defined (uniform) object into that directory so that those institutions who want to can share information about objects in that directory. In the case of higher education, we have chosen an object called "eduPerson," which is promulgated by EDUCAUSE and Internet2. It's a quasi-standards-based definition based on "inetOrgPerson."

RK: It's a name, address, phone number?

CG: Those things are part of inetOrgPerson; eduPerson has data that higher education might want. It has many, many attributes, like whether a person is a student or faculty member. We have fought bitterly over them for the last three years. It's quite extensible.

The eduPerson record is supposed to have the attributes that span the 3,000 higher education institutions in the USA. It might include a major, gender, age . . .

Then you can create "UTAustinPerson" with attributes that are the extension part that might be unique to UTAustin. This might be something like football tickets.

The next aspect of this is something developed by Internet2, and promulgated by Internet2 and EDUCAUSE, called "Shibboleth."

Shibboleth is a mechanism for institutions to share attributes about persons or entities in the LDAP directory. The sharing is, in the case of individuals, controllable by those individuals to some extent. In some sense, it has policy aspects to adjudicate requests.

Let me give you an example. Consider using JSTOR, the journal storage for past academic journals, a pan-university entity that stores documents. If I am a student at some university and want to access some document, I go to the "WAYF" ("Where Are You From") processor and select my institution. I contact JSTOR, who then asks my institution to get me to logon using my local credentials, login ID, and password—all from the LDAP directory.

RK: So JSTOR is somehow permitted to do all this?

CG: Yes, JSTOR is called a "resource provider" and there's a contractual arrangement with them. The next question it's going to ask is, "What is the person's role: faculty, staff, student, or maybe something more generic like 'member of community'?" Perhaps that's all the license requires. The answer might be specific (e.g., "student") or just "yes," the person is a member of the community, whatever is required to satisfy the contractual obligation. JSTOR wants to know if you meet the contractual requirements between the institution and JSTOR.

RK: Would Napster use this to authenticate downloads?

CG: This is exactly what is being used at Penn State.

RK: If UT students bought individual subscriptions, they wouldn't use this fancy management system, would they?

CG: They'd have nothing to do with it. This is for university business.

RK: Transcripts?

CG: Your transcript wouldn't be part of the system but access to the transcript might well be controlled by it.

RK: How many kilobytes in a typical student record?

CG: It's big. Maybe more than 100KB. The digital credential would be stored in the directory. Perhaps a faculty biography might be stored there. Generally, I view it more as containing items like a biography than as pointing to them. It might have group memberships like "member of English 101, section 1."

RK: Like a class schedule?

CG: No, it's memberships. A set of classes—and their members—might be derivable from the memberships, though. Strict privacy laws govern release of such data very strongly.

There's a goal here among some of the content providers and some of our business partners to reduce some of the bilateral contracts that exist.

RK: So there's a directory structure that doesn't sound so lightweight to me, and a permissions structure . . .

CG: You have an attributes release policy, which is fundamental to the system. I have a single PowerPoint slide that illustrates the Shibboleth mechanism; it takes 11 clicks to get through it.

RK: So Shibboleth is complicated?

CG: Yes, many protocol elements and policy implementers.

The identity portion is managed by the LDAP and Shibboleth pieces.

Exchanging things of value like contracts or credit card transactions or grant applications—anything that has to be signed and binds one or both parties to perform specific actions—is another matter.

The system must ascertain "authority" of the signee. Furthermore, the recipient has to trust or know with certainty that the signee has the proper authority to perform the transaction. How does the recipient trust such a thing in the case of a person they don't know?

RK: So we're talking about medical grants, for example, that move millions of dollars and require federal oversight for regulations more than we're talking about buying a CD with a credit card from a Web vendor?

CG: Yes. Let's talk about the mechanical process that has gone on in the past. Many people don't realize that a grant proposal to the National Institutes of Health includes the principal investigator's signature but also the signature of the University chief business officer or provost of research that is capable of binding the institution to the specific contractual rules and implicit laws and other context that exist between the institution and the granting agency.

RK: This is sounding very much more like legal issues than technical ones.

CG: Yes, it does. So NIH has on file, on paper, the signatures of every single person in the institutions of interest that is allowed to sign such documents. Historically, those signatures are compared manually to the signatures on the grant applications.

In the public key infrastructure (PKI) world, we have various levels of assurance. Some people want a "high assurance" digital certificate in order to sign big contracts. That technically is not the case here. The assurance here has to do with the authentication of the individual. In the federal scheme of things (and there is a mechanism and set of standards for using these tools), they have created a "bridge" in PKI terms. The bridge enables entities to establish transitive trust.

Say the Department of Education has a PKI system and the Department of Commerce also has a PKI system. How do you establish a linkage between them so that trust between them doesn't require (re-)issue certificates to members of the "other" department? What you want to do is put something in between the two PKI systems that is basically a set of policies that examines both entities and says, "For your assurance levels Red, White, and Blue, you'll require these properties of the other system," while the other system ends up with a similar statement, "For your assurance Bronze, Gold, and Platinum, you'll need . . ."—potentially a mapping between attributes of the two systems. This ultimately establishes transitive trust between the two agencies.

Interestingly, this mapping happens only at the bridge. Cross-certification happens only at the bridge. The agencies themselves continue to operate as before. That means if there's a federal bridge and NIH is cross-certified into the bridge and the UT system is cross-certified into the bridge, then a principal investigator and chief business officer at UT can digitally sign a contract and send it to NIH. NIH has an electronic process where that signature is verified by going back to source institution through the bridge, looking it up through the LDAP directory, and ensuring that the certificate is still valid and was valid at the time it was signed.

RK: These mapping institutions sound very complex, with an NxN matrix. No standards for this?

CG: We must do mappings at this point because we don't know yet what sort of standards are required. This might be more of an interim measure than a long-term solution.

RK: That would surely aid scaling. Why is this mechanism better than comparing signatures in a file cabinet?

CG: Good question. Some argue that, after 13 years of trying to implement this, perhaps this mechanism is too complex. There are a number of things, though, this provides for us.

If you're going to have an open Internet, which lots of us would like to see preserved, you're going to have to have secure transactions and trustable transactions on the network, even more than we trust credit card transactions now.

There are some types of transactions that occur at high frequency that can benefit from this type of system. Consider federal student loans. Right now, every loan is backed by a piece of paper in a file cabinet somewhere, and it's literally millions of loan applications every year.

RK: Computers have long been characterized by the accounting folks as backup for paper. Paper is "the real thing."

CG: Yes. If you'd like to get rid of paper, then you'll need an identity management system.

RK: How much will it cost to get rid of paper and save all that money?

CG: It's surely a chunk of money. But realize that right now we authenticate students, for example, hundreds of times during their tenure at an institution (e.g., library, health center, registration). It's distributed across a large number of organizations and costs a bit each time it's done. With identity management, it's conceivable that the hard part is done but once, and then authentication is simple and cheap.

There's also the legal requirements about discussing student data in email, for example. Any email that contains a social security number or student name or medical data must be encrypted. This infrastructure provides that capability, something we don't have now. It also reduces the potential for identity theft.

RK: And, of course, you can reject unsigned email and get rid of a lot of spam.

This is all hard to do, right?

CG: Not only is it hard to do, but it's all part of the infrastructure and, thus, is invisible to users and upper management. It's not as sexy as a new application, so this relates to what's really valuable to us as a society —the user doesn't want any of this right now.

RK: Would this system, in the future, displace Microsoft's Passport system or Amazon's one-click ordering by remembering your data and coordinating with those type of systems?

CG: That's certainly what Shibboleth in combination with PKI in fact can provide.

RK: The PKI has been talked about for years and years. PGP has attempted to implement some of it via machines at MIT and other places. How come we haven't seen more penetration and deployment of PKI? What are the challenges?

CG: Getting that infrastructure in place and deployed. I have one institution that's been up for about four years. They have three or four thousand certificates issued. They've been able to change passwords securely over the network. That's a big deal when your computer center is flooded with tens of thousands of gallons of water and everyone has to work at home, as happened there.

Dartmouth is issuing certificates to its incoming freshman this year

RK: All standardized and sharable?

CG: All X.509v3 and, in theory, sharable. Unfortunately, the contents of the certificates are not universally standardized. At UT, where we're trying to do the same profile across all the institutions, we're still encountering some conflict. Some institutions, for example, don't want to include a user's email address in the certificate, which is counterproductive.

RK: Do the Dartmouth freshmen get a physical token?

CG: I think the key goes on their laptop.

RK: So a stolen laptop is a stolen identity?

CG: No, you need the password, too.

RK: Ah, so it's both a stolen laptop and torture. Similar to some other systems. How does this all compare to extortion like, "Write me a check for $50,000 or I'll kill you"?

CG: About the same level.

RK: What non-technical constraints do you live under?

CG: The public isn't demanding it, and it's not on their conscience.

Logins and user IDs are proliferating. I have over a hundred. I don't put them into my browser; I write them to an encrypted file.

Of course, my browser remembers the non-critical passwords, like for the New York Times.

RK: And this identity management will be affordable?

CG: It's already affordable. LDAP is even available as "freeware."

We're doing it (though we're often doing it wrong), and we're doing it in many ways. Scalability will work.

RK: As you look forward and get the budget and start deploying this system, what do you see?

CG: There are detractors who says it's too complex and won't get off the ground. Until it's more widely deployed, it's always a possibility. My experience has been that I've been working for 13 years to get it out there. I haven't succeeded yet. But in my 13 years,

nothing has come close to displacing it or to being an alternative that is as effective as the system appears. A year from now, we'll have more progress, more Dartmouths, and more success stories.

Demonstrable savings will appear with more security and better individual privacy; we're slowly building momentum for this train.

RK: Thanks for sharing with us today!

# Thanks to USENIX Supporting Members

**Addison-Wesley/Prentice Hall PTR**

**Ajava Systems, Inc.**

**AMD**

**Aptitune Corporation**

**Asian Development Bank**

**Atos Origin BV**

**Delmar Learning**

**DoCoMo Communications Laboratories USA, Inc.**

**Electronic Frontier Foundation**

**Hewlett-Packard**

**Interhack**

**MacConnection**

**The Measurement Factory**

**Microsoft Research**

**Perfect Order**

**Portlock Software**

**Raytheon**

**Sun Microsystems, Inc.**

**Taos**

**UUNET Technologies, Inc.**

**Veritas Software**

## USENIX Supporting Member Benefits

One representative receives the benefits on behalf of the company. Join today! Send email to Catherine Allman at sales@usenix.org.

- Free subscription to *;login:*, the magazine of USENIX, both in print and online
- Online access to all Conference Proceedings, 1993–present
- All Conference Proceedings produced during the membership term can be downloaded to your institution's server, giving your students and staff full access to papers from our events
- Place one free ad in *;login:* ($1700 value) during the membership term
- Receive a 10% reduction in sponsorship and exhibit fees for USENIX-sponsored conferences, as well as premium placement on the exhibit floor
- Register up to ten staff at the member price for conferences during the membership term ($1100 value)
- Your click-through logo or company name on the USENIX Web site
- Acknowledgment in conference materials and *;login:*
- The right to vote in USENIX Association elections
- Discounts on technical registration fees for all USENIX-sponsored and co-sponsored events
- Discounts on purchasing Proceedings, CD-ROMs, and other USENIX publications
- Discounts on industry-related publications such as *Sys Admin, Linux Magazine*, and O'Reilly and No Starch Press books

RIK FARROW

# musings

Rik Farrow provides UNIX and Internet security consulting and training. He is the author of UNIX System Security and System Administrator's Guide to System V.

rik@spirit.com

IN MY PREVIOUS COLUMN, I POSITED the existence of individuals or groups that could break into computers and maintain access over a period of time without being noticed. While I have yet to actually learn of an individual being charged with crimes related to such trespass (the only people we see being prosecuted are the flamboyant Web site defacers, with a sprinkling of identity thieves), I continue to hear hints of the existence of these subtle invaders. Unlike those who sell credit card numbers for fifty cents each, these crafty groups would have access to IP (intellectual property) worth many millions of dollars or to bidding information, also worth millions.

What I have recently been reminded of are techniques that could be used to pass undetected through firewalls. These techniques permit the subtle invader persistent access to the founts of her wealth, all the while bypassing all the security products arrayed to detect any intrusion. I speak of tunnels that use two of the Internet's most common protocols: DNS and HTTP.

DNS and HTTP tunnels are nothing new. I have written of these in the past, and so have others. In the December 2003 *;login:*, Mudge wrote about techniques for detecting such tunnels (see http://www.usenix.org/publications/login/2003-12/pdfs/mudge.pdf), without actually mentioning any tools that could be used for tunneling.

You probably have heard of some forms of tunneling. Many botnets use connections to IRC channels as a method to communicate with the thousands of zombies attackers may control. IRC provides several advantages in that it can be used anonymously (through IRC relays) and that each zombie controlled through IRC connects outward through any firewall, an operation that is often permitted. If the firewall has been properly configured, outgoing connections that have not explicitly been permitted by an organization's policy will be prohibited, a restatement of an old saying by Ranum. But what of application protocols that will be permitted through a firewall by policy?

## HTTP

HTTP provides the most obvious example. Not only has the Web become the favorite hole through the firewall for attackers, HTTP can also be used for remote communications. HTTPTunnel, which has been

around for many years (http://www.nocrew.org/software/httptunnel.html), provides a simple and easy method for tunneling protocols that might otherwise be blocked by policy and the properly configured firewall. HTTPTunnel converts requests from the client side of the tunnel into conforming HTTP PUT requests (that actually contain base64 encoded data). The server side of HTTPTunnel converts these requests back into IP packets, and converts responses back into legal-appearing HTTP replies. HTTPTunnel allows an internal user to communicate with an external IMAP or SSH server, even when the firewall would otherwise prevent this communication.

A similar technique, hinted at in Mudge's document, can be used to present a command prompt to a remote user. Instead of tunneling another service, this use of HTTP relies on a local server that makes routine requests of a remote server. Each request appears to be a legitimate HTTP request but results in a command prompt being issued at the remote server. The remote user then can enter a command that gets sent as an HTTP response to the local server. The local server executes that command and packs up the results as another HTTP request to the remote server. Simple enough, and although I cannot personally point you to a working version, I have no doubt this tool exists.

DNS has also been used as a tunnel. The NSTX tool (http://nstx.dereference.de/nstx/) provides a simple and slow tunnel using DNS that is not terribly exciting. Dan Kaminski, during his talk at BlackHat 2004, demonstrated a set of much more interesting tools, ones that you can download and try out yourself (http://www.doxpara.com/).

Kaminski stirred up a bit of a fuss while at CanSecWest (a security conference in Vancouver) by explaining how public DNS servers could act as a distributed network of file servers. Kaminski calculated that his scheme, called Domaincast, would require some 20,000 servers to hold the contents of a single ISO image. Although each server would only serve up 256 bytes, it would be possible to download 700 MBs within a reasonable time using this technique. If you are wondering about finding 20,000 DNS servers, Kaminski claims to have found over 140,000 DNS servers in a single Class A network.

While Kaminski did not demonstrate Domaincast (something the overloaded network at BlackHat would very likely have prevented from succeeding), he did use a tool named droute, written entirely in Perl for portability, to demonstrate tunneling SSH over DNS. Droute converts SSH packets into legal DNS requests, and those requests get converted back into SSH at the server end. The real advantage of this technique is the ubiquity of DNS, its ability to be relayed through firewalls, while its disadvantage revolves around the same data packet size and the use of UDP. But for tunneling, DNS works as well as HTTP.

Just for grins (and Kaminski does grin a lot), he also demonstrated passing an AM talk radio feed over DNS. Using the only codex in the public domain (speex), Kaminski received and converted the talk radio feed into a recognizable audio, all over DNS. For those who didn't want to believe in the efficacy of using DNS for tunneling, this silly example was really an eye opener.

By this point, I hope I have convinced you of the reality of tunneling, with ease, through firewalls. Because of the use of approved protocols, firewalls, IDSes, and IPSes have no chance of determining through the examination of individual packets or reconstructed streams of packets that a protocol has been tunneled through HTTP or DNS. The real hope lies in analyzing traffic, as these tunnels will generate traffic that is unlike normal HTTP or DNS traffic in many respects. The HTTP shell tunnel makes outgoing requests to the same server like clockwork and may keep the connection open for a minute while waiting for a command to be entered. An HTTP shell tunnel will also reverse the usual order for packet sizes, as client requests will be larger than server responses.

Kaminski's droute would reveal the tunneling of SSH through the many requests made to the same remote DNS server. If droute were being used for a remote shell, you could detect patterns in traffic similar to those in the HTTP shell tunnel.

These subtle effects would have to be teased out of the usual barrage of network traffic. While Argus (http://www.qosient.com/argus) would be very good at collecting traffic records for analysis, I know of no OS tool that would actually perform the analysis. Mudge, on the other hand, has been working for a company (Intrusic.com) that has a product that can allegedly do this analysis.

I hope that I have, by this time, at least made you a bit uncomfortable. The notion of outgoing tunnels has never made me feel the least bit happy. And while I am on the topic of uncomfortable things, I do want to encourage you again to vote early, as that guarantees a paper record in what may be the pivotal election for the future of the US.

PHIL PENNOCK

# on IMAP service for customers

Phil is an expatriate Briton living in The Netherlands, desperately trying to live up to the job title of Senior Systems Administrator at an ISP. He uses IMAP for handling his email at home and at work.

pdp@nl.demon.net

THE ISSUE OF PROVIDING IMAP SERVICE to customers has been raised more and more frequently. This article attempts to explain the issues involved, why there's no current product at my ISP, and the need to be very careful in promising anything to customers (or potential customers).

I think it important for Sales/Product-Management to understand the issues so that they can be better informed than their competitors.

In Internet mail systems, there's one distinction that most people can happily ignore, even though it affects every message they send or receive. It's the concept of "final delivery." Email messages can be passed through many systems between the sender and recipient (though some systems will decide that 30 is too many and evidence of a mail-loop). But only one of those systems, the last, involves delivery into the user's real mailbox. That's the final delivery.

Up until final delivery, the message is in transit. If the message can't be passed on, then a delivery failure notification ("bounce") message is generated and passed back to the sender. Once it reaches the final mailbox, there will be no bounce. There's no guarantee that the mail will actually be read, but it won't be bounced because it's been delivered to "where it was supposed to go."

IMAP is designed as a system to manage mail once it has reached its final delivery point. SMTP is *before* final delivery. POP3 is a slightly weird hybrid, but not really well suited to use after final delivery. For many customers using POP3, the POP3 retrieval to their Outlook program *is* the final delivery.

Some ISPs do not provide final delivery. They pass mail on to customers. Historically, ISPs just provided SMTP; nowadays you can also find POP3. If a message is not collected within 35 (or whatever) days, the mail systems delete it. If it was read by POP3, it's automatically deleted; if it was not read by POP3, a bounce message is sent back.

Once mail reaches final delivery, it can stay in the system for as long as the user wants. There's no 35-day limit. If someone has an important message they treasure, they might want to keep it for the next 70 years or more—this is to be expected, not something exceptional.

Because ISPs just handle messages "in transit," there are rough limits on how much mail needs to be stored. ISPs engineer the systems for a different set of operating conditions than those used for providing final delivery.

For instance, systems are extremely unlikely to lose email. That doesn't mean that they *can't* lose email. An asteroid can strike the planet, destroying Western

Europe; enough disks can fail in the NetApp, all at once, to lose messages. These are both possible (and the latter is, hopefully, *more* likely). If a customer collects mail regularly, between a few hours' and a couple of days' worth of their email might be lost; if the customer doesn't collect regularly, up to 35 days' worth might be lost. The design makes this unlikely, but there's no possible design that makes losing email impossible. In the worst case, business insurance covers this.

The monetary worth of 35 days of email as opposed to 35 or 70 years of email is an interesting thing to consider . . .

Because of the issues of volume, many ISPs can't offer an IMAP product that scales to all customers. Individual small products can be offered, though, on the scale of what a company of a couple of hundred employees might implement internally—that's doable.

IMAP needs to be considered not as "another way for customers to get their email" but, instead, as "how customers manage email they've received." They each get multiple mailboxes and can move mail between them, delete mail, flag mail, add attributes and keywords, search on message content, and much more besides.

The mail-store of the IMAP system will hold immense amounts of information important to an organization. Only the organization itself knows *how* important. But any strategy needs to handle issues such as backups, replication, archives, and policies on personal mailboxes of departing staff.

But how much information can be held? How much should be held? "All email" can get to be a very large amount of data over time. Much of the information will lose relevance; some will not. From an operations point of view, it's certainly useful to me that I can look over the past three or fours years' worth of list mail, quickly searching for messages matching a few keywords on a subject that I vaguely remember coming up before; with paper memos this would not be a productive use of my time, and I'd be better off figuring out the solution from scratch (or having decent documentation).

Any "IMAP solution" offered needs to consider not just "mail for a few months." It needs to be "email useful to the customer, for the lifetime of its usefulness." It needs to provide for backups, in various forms. It needs true disaster recovery. It needs a lot that is expensive to provide—probably the reason that customers come to ISPs, after their sysadmin/consultant has told them how much it costs to do things properly.

Any IMAP offering that's cheap to implement carries a potential legal and financial nightmare—I really don't think that most ISPs offering IMAP have properly evaluated this, just as many who rushed to offer free accounts didn't evaluate their long-term business plans, either. The ISP market is still young enough that it's filled with cowboys who have much financial backing but simply don't understand what they're doing or what the long-term consequences of their products are, focusing on "get more customers now" instead of "get customers whom we'll keep and who won't be suing us into bankruptcy in three or eight years."

Email has much to offer (including searchability, mentioned above), but it also brings challenges that are typically not addressed. IMAP is an important part of the picture, but not the whole picture. I believe organizations need to understand how email messages and other forms of documentation are used, and establish policies for email retention and migration of information into more formal documentation.

For instance, a policy might resemble the following:

Email messages more than seven years old are archived onto a read-only long-life medium and deleted from the live system. After only six years, someone reviews mail to mailing lists X, Y, and Z for information that looks as though it might still be relevant and collects those messages for review by specialists.

The specialists collate those that still hold relevant important information and ensure that all the information is held in Procedures, Policies, and Guidelines or other formal documentation.

Clearly, this ties deeply into internal information management. Not losing information involves some bureaucracy (the NOC will now boo and hiss). Choosing how to handle this is not easy and may well be different for each customer. Their business processes for information management need to be designed to migrate information of any value from ephemeral communications such as email into more static forms of retaining information such as traditional documentation.

An ISP could offer some standardized services along these lines, with tools and calendaring designed to make it easy for companies to collect the information they need; with automatic burning to DVD (this year—who knows which medium in six years?) of a customer's mail every time a certain volume is accumulated or amount of time passes; with those DVDs being mailed to the customer by recorded delivery or courier. There's a lot that can be done.

But any time that you look at getting this involved in a company's internal processes, what you're actually selling is IT outsourcing services, not Internet access.

I believe that offering "proper" IMAP access is something which intrudes deeply into the market of some much larger companies; it's not something to be taken lightly, and it's certainly outside their usual areas of expertise.

Certainly, if one wished to start moving into ASP (ye olde Applications Service Provider, which we heard so much about a couple of years ago) or IT outsourcing, then that's a different matter, and IMAP is just one of the technologies which would be used to provide services. But this is not Internet access. It's not shifting information around or allowing customers to shift information to others. It's controlling how customers shift information around within their own organization, which is a fundamentally different animal.

"Just providing IMAP access to a mail drop" is a short-sighted viewpoint; unfortunately, many customers won't understand the issues here and will want "just that." Perhaps this article can be massaged into a document issued by ISPs to customers to help them think more deeply about the issues and to realize that the providers want to work with them to offer good solutions, not just take their money for whatever they can without regard to the consequences.

Some ISPs can offer a small IMAP product in the near future, but it's essential that this isn't hyped and that we are proactive in ensuring that customers have information on the limitations described here; and it's equally important that they consider this a stop-gap solution while they develop something that integrates better with their business processes.

PETER H. SALUS

# the bookworm

Peter H. Salus is a member of the ACM, the Early English Text Society, and the Trollope Society, and is a life member of the American Oriental Society. He owns neither a dog nor a cat.

peter@netpedant.com

**BOOKS REVIEWED IN THIS COLUMN**

**THE DESIGN AND IMPLEMENTATION OF THE FREEBSD OPERATING SYSTEM**
Marshall Kirk McKusick and George V. Neville-Neil
Boston, MA: Addison-Wesley, 2004.
Pp. 683. ISBN 0-201-70245-2.

**OPEN SOURCE LICENSING**
Lawrence Rosen
Boston, MA: Pearson Education, 2004.
Pp. 432. ISBN 0-131-48787-6.

**THE SUCCESS OF OPEN SOURCE**
Steven Weber
Cambridge, MA: Harvard U.P., 2004.
Pp. 312. ISBN 0-674-01292-5.

**SUCCEEDING WITH OPEN SOURCE**
Bernard Golden
Boston, MA: Addison-Wesley, 2004.
ISBN 0-321-26853-9.

**ISLANDS IN THE CLICKSTREAM**
Richard Thieme
Rockland, MA: Syngress, 2004. Pp. 336.
ISBN 1-931836-22-1.

**THE SPAM LETTERS**
Jonathan Land
San Francisco: No Starch, 2004.
Pp. 210. ISBN 1-59327-032-1.

**RULING THE ROOT**
Milton L. Mueller
Cambridge, MA: MIT Press, 2004.
Pp. 328. ISBN 0-262-63298-5.

**A PRACTICAL GUIDE TO RED HAT LINUX, 2ND ED.**
Mark G. Sobell
Boston, MA: Addison-Wesley, 2004.
Pp. 1200 + 4 CD-ROMs.
ISBN 0-13-147024-8.

This column is being written in July, and most of what I've been reading isn't "hard-core" techie stuff. But it's relevant to what we all do.

### THE BEST

McKusick and Neville-Neil have produced the very best technical book I have looked at this year. Seriously.

*The Design and Implementation of the 4.3 BSD UNIX Operating System* from 1989 contained 471 pages. The 1996 . . . *4.4 BSD* . . . tome was 577 pages. *FreeBSD* has put on another 100 pages. Not at all an excessive rate of growth. But there has been a great growth in the information and its importance.

"But I don't run FreeBSD," I hear you moaning. Well, I think you need to understand what McKusick and Neville-Neil have to say if you're running any version of UNIX or Linux. This is especially true if you're running OS X, as Darwin, which is based on FreeBSD, lies at the Apple's core.

If you need to understand just how a kernel works, you need this book. McKusick and Neville-Neil have done the community a favor. It truly deserves to be a "best seller."

### FREE/OPEN SOURCE

Rosen has brought us the first guide for managers and lawyers to the law of open source. It is definitely both thorough and interesting. But it is quite dry, and if you aren't a determined reader, you may well abandon this book partway through. Where information is concerned, you will have missed out. But this is not an easy read.

The book concludes with a series of appendices that contain the texts of nearly a dozen licenses: BSD, MIT, Apache, Artistic, GPL, LGPL, MPL, CPL, OSL, and AFL. I wish there weren't so many variants.

Weber's book is of genuine importance. Dealing with the socio-political and economic bases of open source, Weber moves to a first-rate analysis of the business models of the companies that have participated in the movement.

If you are among the crowd reading Groklaw, you should read Rosen and study Weber. Though I've mentioned the former is dry, the latter is a really good read.

I didn't like Golden's book. In brief, it outlines the "Open Source Maturity Model," which is a "formalized method" for "assessing open source software."

### SOME ESSAYS

For nearly a decade, Thieme has been writing brief pieces on creativity and imagination, the world of cyberspace, and networks. Some of the essays are very good; others leave me cold. But all have some merit. It's nice to see them put together—and as the essays are relatively brief, the book is well worth dipping into.

### THE RETURN OF LAZLO TOTH

Those who remember Don Novello (a.k.a. Father Guido Sarducci; a.k.a. Lazlo Toth) will be delighted with Land's *Spam Letters*. *The Lazlo Toth Letters* have delighted me since 1977, and *The Spam Letters* will still be read in 27 years, too. Several made me laugh aloud. And it's under $15!

### REAPPEARANCES

Mueller's *Ruling the Root*, about ICANN and Internet governance, is now out in paperback. While I find it intellectually dishonest, it is still the only attempt at a history of a shameful activity (IMO).

Sobell's *Practical Guide to Red Hat Linux* is out in a second edition. It comes with a full 4-CD set of Fedora Core and Enterprise Linux.

**PLANET BROADBAND**

Rouzbeh Yassini et al.

Indianapolis, IN: Cisco Press, 2004.
Pp. 168. ISBN 1587200902

**REVIEWED BY JOEL E. NATT**

"Planet Broadband" is not a Star Trek episode or some other sci-fi title. It is the title of a book that answers the questions, Where did the term "broadband" originate and why is it used when discussing cable Internet services? When I first received the book, my initial thoughts focused on the cable Internet environment, but this short book also presents the DSL side, as well as explaining how the concept of broadband communication evolved and where it might go in the future.

The book is essentially a history of the growth of the Internet and how the concept of broadband was developed. I can foresee future generations of information technology students reading this book in college courses.

A detail-oriented reader, I read the book in about eight hours; that may seem slow but, considering the topic and the events occurring within the world of telecommunications, was better than I expected. As I read, I paused from time to time to think about how the points made by Yassini are either already occurring or could be implemented. Additionally, I found enlightenment in the knowledge that both DSL and cable Internet are actually considered broadband. While the cable companies have a tendency to advertise themselves as broadband services, in reality they are only one of several kinds of broadband. Some of the services presented include virtual training, which many universities are now offering, and virtual meetings between individuals spread around the world. While *Planet Broadband* does point out that everything is not immediately on-demand right now, in time we can witness more and more services going that way.

Yassini discusses potential uses of the technology, such as checking what's in your refrigerator at home from work, or your washing machine placing a service call to the manufacturer before you know there is a problem. These are features of a broadband planet, and within time our children and their children will not know a world where modems connected to the Internet or downloads took hours. But as you read, you realize that while today constitutes the birth of broadband, this book serves more as an introduction to the Internet of tomorrow and, as such, is a must-read.

To prove the point that *Planet Broadband* is an introduction, my baby boomer mother, who works in the cable industry in customer service, read it cover to cover in one day and then told me, "This is where we are going." If a 50+-year-old parent can see it, surely the world will be there some day.

Near the end of the book, Yassini focuses on telecommuting, comparing it to working at the office. While he clearly points out that telecommuting is not for everyone, or always appropriate, it has productivity advantages and will become increasingly common. I think one of the best examples he cites is the MCI commercial of a woman working at home and changing a presentation in real time for clients and co-workers halfway across the country. That image is an example of how the world has changed thanks to the advances of broadband and the ability to telecommute.

*Planet Broadband* is not designed for the IT world so much as the general public, though I would recommend it as an excellent addition to anyone's library. For management it is an excellent resource to justify a telecommuting policy or practice, providing a good argument to upgrade from the modem and narrow band to the world of high-speed Internet.

# USENIX notes

## USENIX MEMBER BENEFITS

Members of the USENIX Association receive the following benefits:

FREE SUBSCRIPTION to *;login:*, the Association's magazine, published six times a year, featuring technical articles, system administration articles, tips and techniques, practical columns on such topics as security, Tcl, Perl, Java, and operating systems, book reviews, and summaries of sessions at USENIX conferences.

ACCESS TO ;LOGIN: online from October 1997 to last month: www.usenix.org/publications/login/.

ACCESS TO PAPERS from USENIX conferences online, starting with 1993: www.usenix.org/publications/library/proceedings/

THE RIGHT TO VOTE on matters affecting the Association, its bylaws, and election of its directors and officers.

DISCOUNTS on registration fees for all USENIX conferences.

DISCOUNTS on the purchase of proceedings and CD-ROMs from USENIX conferences.

SPECIAL DISCOUNTS on a variety of products, books, software, and periodicals. For details, see www.usenix.org/membership/specialdisc.htm.

FOR MORE INFORMATION regarding membership or benefits, please see

www.usenix.org/membership/

or contact

office@usenix.org

Phone: 510 528 8649

## USENIX BOARD OF DIRECTORS

Communicate directly with the USENIX Board of Directors by writing to *board@usenix.org*.

**PRESIDENT**

Michael B. Jones, *mike@usenix.org*

**VICE PRESIDENT**

Clem Cole, *clem@usenix.org*

**SECRETARY**

Alva Couch, *alva@usenix.org*

**TREASURER**

Theodore Ts'o, *ted@usenix.org*

**DIRECTORS**

Matt Blaze, *matt@usenix.orh*

Jon "maddog" Hall, *maddog@usenix.org*

Geoff Halprin, *geoff@usenix.org*

Marshall Kirk McKusick, *kirk@usenix.org*

**EXECUTIVE DIRECTOR**

Ellie Young, *ellie@usenix.org*

**EVENT SALES & MARKETING**

Cat Allman, *cat@usenix.org*

Anne Dickison, *anne@usenix.org*

## Summary of the USENIX Board of Directors Meetings

by Tara Mulligan

The following is a summary of the actions taken by the USENIX Board of Directors from March 3, 2004, through June 27, 2004.

**FINANCES**

Funds were allocated to engage a consultant for market research on regional training-only events.

USENIX agreed to collect donations on behalf of the Open AFS Council of Elders and disburse them according to the wishes of the Council.

USENIX will be a bronze-level sponsor of the Fifth Grace Hopper Celebration of Women in Computing Conference (October 2004).

The Board encouraged the staff to come up with a proposal and plan to create a resource for the most important legal issues in IT that system administrators should be aware of and a best practices for dealing with those issues.

**CONFERENCES**

Early Bird registration fees for nonmembers for the 2004 LISA Technical Sessions will be as follows: 1 day: $400; 2 days: $700; 3 days: $795.

It was agreed that the student nonmember registration fee for OSDI '04 would be $290.

**NEXT BOARD MEETING**

The next meeting will be held on Tuesday, November 16, 2004, at the LISA conference in Atlanta, GA.

## TWENTY YEARS AGO IN ;LOGIN:

By Peter H. Salus

In September 1984, USENIX held one of its first workshops: Distributed UNIX. It was in Newport, Rhode Island. About 80 "UNIX hackers" attended. A report on the meeting appeared in the November 1984 *;login:*.

The report was by "Veigh S. Meer." I have asked many of "the usual suspects" but have not ascertained whose alias this was. If anyone is willing to admit to it, please write me. The statute of limitations is long past.

The report was divided into three parts: a chronological narrative of the presentations, a summary of the state of the art, and "my personal view of the future."

As you might expect, this last is what has driven my wish to give credit to the author. The "personal view" is really interesting, particularly from a 20-year perspective:

> What do we need in the future? I doubt either the NASA and Livermore experience (multiple Crays) or the UCSD or Toronto configuration (hordes of students and almost no hardware) is typical of the future. Suppose the standard 5M workstation (1 Mips CPU, 1 Mpixel display, 1 Mbit network, 1 Mbyte main memory, and 1 mouse) becomes routinely available as an individual terminal. In that case, distribution for load sharing is not going to be interesting; remote files will be interesting. More important, however, is that putting several CPUs in the station will be fairly cheap. What should they do? Well, they could handle communications, so maybe efficient protocols will become less important.
>
> They could also handle caching, a favorite technique to improve performance. But most important, we need a way to use multiple processors in local environments, to improve the performance of individual jobs.
>
> UNIX was designed as a uniprocessor operating system; C was designed as a single-thread procedural language; and I think we need some research that starts from other premises. Otherwise something else is going to be the operating system of the future.

Wow! Twenty years ago. C'mon, Meer! Come out and let us know who you are.

## IN MEMORIAM: CHARLES "CHUCK" YERKES

Chuck Yerkes
December 31, 1963–
August 26, 2004

USENIX/SAGE
member #30530

"Shoes, shirt, sober . . . pick two": a typically irreverent quip from the ever-ebullient Chuck to a former manager.

Chuck Yerkes was funny, brilliant, and sometimes eccentric, but when it came to computers, he was known for his intelligence, great knowledge, and utter dedication to technology "done right." Chuck never hesitated to offer his professional opinion, usually served up with a humorous anecdote and with the intended result of someone gaining information they needed.

Chuck worked for nearly 20 years in the computer industry, including stints as an audio/visual technician (read "sound guy"), as a core member of the Internet team at J.P. Morgan in the early 1990s, for several years as a technical consultant for Sendmail, Inc., and most recently as a consultant at PeopleSoft. Chuck was passionate about open source and UNIX, and his enthusiasm for reliable computing caught on with those he worked with. He, along with his long-time partner, Valerie Acton, had been a USENIX and SAGE member since 1993. He attended many USENIX conferences, and in early August helped staff the USENIX membership booth at LinuxWorld. A resident of Berkeley, CA, on August 26 he was involved in an accident while on his way home from work on his motorcycle, and did not survive.

Those who knew him personally, professionally, via the SAGE-members mailing list, or through other electronic mailing lists will remember Chuck for his generosity, kindness, and spirit. Our sincerest sympathies go to Val and to Chuck's family and friends. Chuck brought a bright spark to the world. We will miss him.

[This tribute was contributed by Tara Mulligan, a former co-worker of Chuck's.]

## MEMBERSHIP NEWS

By Tara Mulligan
Member Services Manager

Thank you, USENIX and SAGE members, for your ongoing support! With your help, USENIX continues to offer the highest level of conferences and publications to the advanced computing systems community. We have a terrific slate of conferences lined up for the remainder of 2004 and for 2005, and we look forward to seeing you at one or more. See our Conference Calendar: http://www.usenix.org/events/.

The current membership of USENIX is 6,200, with 2,770 USENIX-only members, 3,050 USENIX/SAGE dual members, and 380 SAGE-only members. For the first time in a few years, our membership is increasing.

### SAGE MEMBERSHIP

With the recent change in governance, we would like to take this opportunity to assure members that SAGE benefits and renewals will continue as they were. Some members of the SAGE community are exploring the possibility of creating

SAGE as a separate entity from USENIX. We will keep you posted on any new developments.

### RENEWALS

In January 2004, we implemented an electronic renewal system. Under the current system, two months prior to the expiration of your account, we send you an electronic notice of impending expiration, with a link to a page with your information pre-loaded for renewal. If you do not respond to that first electronic notice, in the month prior to your account expiration we send a traditional paper renewal notice. Should we still not hear from you, in the month after your account expiration we send an electronic past-due notice. We hope that does the trick, but if not, we'll leave you alone thereafter.

Members are encouraged to renew at any time, using their former membership login. USENIX continues to send you *;login:* for a three-month grace period after expiration. Try online renewal at http://www.usenix.org/membership/.

### MEMBERSHIP CARDS

Since we are sending membership information electronically and offering purchase of membership renewals and new memberships online, USENIX is ceasing to send printed membership cards. Instead, we provide an online membership card pre-filled with your information, which you can print on your own whenever you wish.

As always, please feel free to email any questions you may have to us at membership@usenix.org.

## COMING SOON: ;LOGIN: SURVEY

USENIX will soon be conducting
an online survey about ;login:.

Watch your mailbox in October for the URL.

Spend just a few minutes
to help us make ;login: even better.

**PROFESSORS, CAMPUS STAFF, AND STUDENTS—**

**DO YOU HAVE A USENIX REPRESENTATIVE ON YOUR CAMPUS?**

**IF NOT, USENIX IS INTERESTED IN HAVING ONE**

**AT YOUR UNIVERSITY!**

The USENIX University Outreach Program is a network of representatives at campuses around the world who provide Association information to students, and encourage student involvement in USENIX. This is a volunteer program, for which USENIX is always looking for academics to participate. The program is designed for faculty who directly interact with students. We fund one representative from a campus at a time. In return for service as a campus representative, we offer a complimentary membership and other benefits.

A liaison's responsibilities include:

- Maintaining a library (online and in print) of USENIX publications at your university for student use
- Distributing calls for papers and upcoming event brochures, and re-distributing informational emails from USENIX
- Encouraging students to apply for travel stipends to conferences
- Providing students who wish to join USENIX with information and applications
- Helping students to submit research papers to relevant USENIX conferences
- Providing USENIX with feedback and suggestions on how the organization can better serve students

In return for being our "eyes and ears" on campus, liaisons receive a complimentary membership in USENIX with all membership benefits (except voting rights), and a free conference registration once a year (after one full year of service as a campus liaison).

To qualify as a campus representative, you must:

- Be full-time faculty or staff at a four year accredited university
- Have been a dues- paying member of USENIX for at least one full year in the past

For more information about our Student Programs, see
http://www.usenix.org/students

USENIX contact: Tara Mulligan, Scholastic Programs Manager, tara@usenix.org

# conference reports

- This issue's reports focus on the USENIX Annual Technical Conference (USENIX '04), held in Boston, Massachusetts, June 27–July 2, 2004.
- Our thanks to the scribe coordinator:

  Rik Farrow
- Our thanks to the summarizers:

  Bill Bogstad
  Ming Chow
  Brian Cornell
  Richard S. Cox
  Todd Deshane
  Patty Jablonski
  Rob Martin
  Martin Michlmay
  Adam S. Moskowitz
  Peter Nilsson
  G. Jason Peng
  Calicrates Policroniades
  David Reveman
  Matt Salter
  Swaroop Sridhar
  Sudarshan Srinivasan
  Matus Telgarsky
  Wanghong Yuan
  Ningning Zhu

# USENIX ANNUAL TECHNICAL CONFERENCE (USENIX '04)

*Boston, Massachusetts*
*June 27–July 2, 2004*

## PLENARY SESSION

*Summarized by Richard S. Cox*

### *Open Source and Proprietary Software: A Blending of Cultures*

*Alan Nugent, Novell*

Alan Nugent opened the USENIX Annual Technical Conference with his plenary session addressing the integration of open source software and procedures at Novell.

Many people believe that open source will destroy the software industry; that it is developed by hackers without discipline; that it is a fad; or that there is no money in open source. Seeking to debunk these myths, Alan first suggested that, rather than wrecking the industry, open source has increased diversity and thus has created opportunities. Second, open source software can be of very high quality, since a majority of open source contributors are professional developers working on projects that interest them. The community is growing daily, and contributors are quick to realize important initiatives. While open source software is free, there is a market for selling the support and maintenance contracts that large customers require before they are willing to build mission-critical systems using a package.

The adoption of open source has allowed Novell to work with customers to build solutions that more closely match their needs and infrastructure. Novell does this by providing complementary packages (open or closed) that interact with those developed by the open source community. By focusing on complementing existing projects, rather than providing substitutes, they avoid competing with open source developers, an arrangement that benefits all involved.

At Novell, this has required reworking the legal framework under which licenses are sold, expending significant effort in convincing customers to accept solutions combining proprietary and open components, and changing the focus of the organization.

Greg Mitchell asked how the sociology of the company changed as more open source developers were brought in. Alan responded that, while some employees were upset and a few even left, the acquisition of open source teams has been very successful and brought more energy throughout the company. Novell was able to do very well retaining employees from acquired companies.

## GENERAL SESSION PAPERS: INSTRUMENTATION AND DEBUGGING

*Summarized by Swaroop Sridhar*

### *Making the "Box" Transparent: System Call Performance as a First-Class Result*

*Yaoping Ruan and Vivek Pai, Princeton University*

Mr. Yaoping Ruan presented the "DeBox"ing technique for debugging OS-intensive applications. He began the talk with a motivating example of monitoring system call performance on a server running the SpecWeb99 benchmark. He pointed out that system call profile as measured from user space sometimes indicated anomalous kernel behavior. He identified the trade-off between speed, completeness, and accuracy among various profiling tools. Later, Ruan presented the

design of the DeBox system. The key idea is to make the system call performance a first-class result and return it in-band (like errno). Proposing a split between the measurement policy and mechanism, Ruan said that the applications should be able to interactively profile interesting events.

Later, Ruan gave details about the implementation of DeBox. He gave the details of profiling primitives added to the kernel and the interface available to the applications. He also provided details about the various kinds of information that the system offered, the amount of change that had to be done to the kernel and applications, and so on. Ruan went on to present a case study on Flash Web server performance. He presented various optimizations with a step-by-step performance analysis.

Ruan concluded by stating that DeBox is very effective on OS-intensive applications and complex workloads. He also claimed that the results showed that the system was portable. During the Q&A session, Ruan said that they were investigating the use of DeBox on other OS-intensive applications such as database systems, but the results were not yet available. More information about DeBox can be found at http://www.cs.princeton.edu/~yruan/debox, or by contacting {yruan, vivek}@cs.princeton.edu.

### *Dynamic Instrumentation of Production Systems*

*Bryan M. Cantrill, Michael W. Shapiro, and Adam H. Leventhal, Sun Microsystems*

In introducing Bryan Cantrill, session chair Val Henson—also from Sun Microsystems—said that she could definitely confirm Sun's use of DTrace in production. Cantrill began his power-packed speech by stating that all of today's tools were targeted at development and not production. As a result, the systems are incapable of dealing with systemic problems. Cantrill asserted that for a tool to be used in production, the necessary constraints are that there should be *zero* probe effect when disabled, and the system must be absolutely safe. To have systemic scope, both kernel and applications must be instrumentable, and the system must be able to prune and coalesce the enormous amount of data into useful information.

Later, Cantrill introduced the various concepts and features of DTrace: dynamic-only instrumentation, unified instrumentation, arbitrary context kernel instrumentation, high-level control language, predicate and arbitrary action specification, data-integrity constraints, facility for user-defined variables, data aggregation, speculative tracing, scripting capacity, boot-time tracing, virtualized consumers, etc. Next, Cantrill elaborated on the D language: syntax and use, D intermediate form, probes, providers and actions, aggregations and scalability of the architecture. Cantrill also shared some experiences with DTrace and gave some examples of D scripts and analyzed their results. Finally, using the example of a bug in gtik2 applet2—a stock ticker for GNOME desktop—he showed how a small programmer error could cause widespread damage in a production system such as SunRay server. Cantrill challenged the idea that no other existing tool could trace this problem to its root cause, and that a trace was possible only by the extensive use of aggregation functions and thread local variables provided by DTrace.

During the Q&A session, Jonathan Shapiro said that he believed that the gtik2 applet2 problem should be attributed to the fundamental problems in monolithic kernel design, and asked the speaker to comment on the use of DTrace for debugging kernel bugs. Cantrill did not totally agree with Shapiro's views, but only asserted that DTrace was effective in tracing kernel-level bugs. Answering another question, Cantrill said that there was no extra effort required to use this tool with third-party kernel-level modules. When asked whether there were any plans to port their system to Linux or any other operating system, Cantrill answered in the negative and quipped, "Use the best OS available!" The authors can be contacted at dtrace-core@kiowa.eng.sum.com.

### *Flashback: A Lightweight Extension for Rollback and Deterministic Replay for Software Debugging*

*Sudarshan M. Srinivasan, Srikanth Kandula, Christopher R. Andrews, and Yuanyuan Zhou, University of Illinois, Urbana-Champaign*

With the increase in volume and complexity of software development, there is a proportional increase in software bugs, their effects, and the difficulty in tracing or even reproducing them. Various checkpointing and logging mechanisms and their applications have received a lot of research attention in the last decade. Mr. Sudarshan Srinivasan presented Flashback, a lightweight OS extension to facilitate rollback and replay, as applied to software debugging.

After providing a brief general background and motivation for lightweight checkpointing, Srinivasan went straight into the main idea of Flashback. Flashback achieves checkpointing by forking a shadow process, thus replicating the in-memory state of the process. The processes' interactions with the system are logged so that, during replay from a checkpoint, the (shadow) process gets an execution environment similar to the original run. Srinivasan presented some challenges posed due to multi-threading, memory-mapped files, and shared memory and signals. He also presented the approaches adopted in Flashback toward solving these problems.

Srinivasan went on to present some details of the present Linux imple-

mentation regarding modifications to the kernel, changes to GDB, etc. Srinivasan identified incorporating replay support for multi-threaded applications as an area for future work. Later, responding to Val Henson's question regarding possible applications of Flashback other than debugging, Srinivasan said they were investigating uses of Flashback in other avenues, such as lightweight transaction models. The source code for Flashback can be obtained at http://carmen.cs.uiuc.edu/.

## ADVANCED SYSTEM ADMINISTRATION SIG: AUTOMATING SYSTEM AND STORAGE CONFIGURATION

*Summarized by Rob Martin*

### *The CHAMPS System: A Schedule-Optimized Change Manager*

*Alexander Keller, IBM T.J. Watson Research Center*

Dr. Alexander Keller began by describing CHAMPS (Change Management with Planning and Scheduling) as "a schedule optimized change management system" that is not yet a product. "It's a research prototype [providing] change management with planning and scheduling." Its end product is the schedule: "all the things that are going to be carried out on which machines [and] concrete systems that are going to carry out these tasks." Keller described this as "a change plan."

Keller described CHAMPS within the larger context of change management as "trying to assess the impact of a change and figure out what the dependencies between different tasks are and creating a change plan. . . . We are specifically not concerned with actually implementing or rolling out a change, because there are deployment systems that can do this." Later in the talk, Keller gave examples of such systems: cfengine and Tivoli Intelligent Orchestrator for Service Optimization.

The CHAMPS system consists of two subcomponents: the Task Graph Builder and the Planner and Scheduler. The end product of the system is a change plan depicted in a standard workflow language (BPEL4WS). This, in turn, is fed into an "off-the-shelf" deployment system which "rolls out the changes and provides feedback status information back into the [CHAMPS] system for summary planning." The workflow engine executes the plan and monitors whether each activity has completed or failed.

A key goal of CHAMPS is optimizing the schedule based on dependencies to carry out tasks in parallel wherever possible. The information used to figure out which tasks are going to be carried out in sequence and which in parallel are "product dependency descriptions." "The availability of authoritative dependency information [from package developers] is very important." Once the dependencies are put into the Task Graph Builder, the system generates the Task Graph.

"The Task Graph tells us which tasks are going to be carried out, in what order . . . , and whether they must be in sequence or can be in parallel." The Task Graph is used as input to the Planner and Scheduler. "The Planning system may make decisions such as 'we must take away a machine from customer X and give it to customer Y'; in order to do that [the system] must be aware of the service level agreements and policies that the data center has. . . . It is up to the planning system to bind the existing Task Graph to the complete system to generate concrete system names, times, and dates.

"We put in declarative information about the relationships between tasks, [and the CHAMPS system] automatically generates this schedule and allows the administrator to apply modifications to the schedule." Estimating individual task duration is crucial, and CHAMPS uses "past deployments" to calculate future durations for individual tasks.

Multiple task graphs, each representing a single change, are input into the Planner and Scheduler, which then binds the changes to services and resources and optimizes a schedule for all of the changes. "We are treating this problem as an optimization problem." The optimization is done by "fifty pages of Java . . . not visible to the administrator. We support a very general level of objective functions [for] minimizing penalties, maximizing profits. The administrator selects from push-button options that provide choices like 'maximize profits,' 'minimize downtime,' 'maximize throughput,' 'minimize costs.' By selecting one [or a combination] of these choices the optimization parameters are automatically set." The CHAMPS system then calculates the optimum schedule, if necessary deciding that certain changes cannot be accomplished given the overall set of changes requested.

Keller concluded by listing the areas that require more work in the future, including "tooling for deployment descriptors," reusing change plans (storing them in an XML library, for example), knowing when a plan is running behind schedule, carrying configuration information along with the workflows, and identifying parameters that flow out of one task and are required for other downstream tasks.

During the Q&A session, there was a lively exchange on the "sad state of dependencies in software packages." Is there a standard for describing dependencies? Work done by the Grid Forum on defining a standard, and the use of dependency sniffing tools were mentioned.

### *Autonomics in System Configuration*

*Paul Anderson, University of Edinburgh*

What is system configuration? Paul Anderson, professor and researcher at the University of Edinburgh, starts out with some background on the general subject. When you want to build a new site, you start off with three things: the hardware (empty disks and bare metal), the software, and specifications and policies about how you want the final system to run. The core of the configuration problem is to take those three things and put them together to get some sort of computer system that performs to the specification. Anderson refers to the final site as a "fabric," a term he borrows from the recent work in grid computing.

Anderson points out that the "main thing to notice is the big distinction between the software and the configuration policies. The pile of software you start off with has no configuration and in theory can all be put on all the machines that you've got. It's the specifications and the configuration policies that differentiate the individual machines."

Configuration starts with the base layer of internal services inside your fabric at a lower level than the applications you want to end up with. DNS, NFS, DHCP, and like services form a base layer you have to get going before you build anything on top of it.

The idea of autonomics is to "take some of the low-level decision making away from the system administrator and have a lot of things happen automatically, so the system administrator can move up a level and think of higher-level policies and planning." As with a compiler, you trust the autonomic system to place low-level data and decide which bits go where.

After the initial configuration, as change occurs due to load balancing, software or hardware failure, and day-to-day use, the autonomic system adjusts the fabric and configuration of the system so that it comes back into alignment with the original specifications and policies. The feedback from the autonomic system does not make changes to the original specification; rather, it brings the fabric back to providing the original services and policies specified.

"Autonomics is not new. Cfengine and lcfg are examples of tools that provide this sort of automatic fixing up of configuration files when something goes wrong at the host level. There are inter-host autonomic systems like fault-tolerance systems, RAID, and load balancing that will adjust systems. What is new is trying to think of this in a uniform way and integrating it into the configuration system."

Anderson described the major issues under consideration in researching autonomic solutions. He described "a declarative specification of what the system behavior should look like. Some kind of logical statement that is true about the system rather than a recipe about how to get there. If you don't have a good declarative statement to start with, then you don't know what to do. . . . The language you need to describe the configuration is not a programming language: We are not talking about a process, we are talking about the description of the configuration data and the way that that system actually is." Anderson gives an example of a declarative statement. "'Host X uses host M as a mail server.' In most configuration systems you don't see statements like these. Rather, you see lower-level details, like a script setting parameters for sendmail.cf."

The goal is to use the declarative language to describe the system and "let the autonomic system juggle the details" to make sure the specifications remain true.

An example declaration: "Make sure we have two DHCP servers on each network segment." This expresses a high-level policy rather than details like "make this machine configured as a DHCP server." The final goal of an autonomic system is to take these declarative statements and generate the details. "System administrators will specify those criteria that are important for the job without specifying the details. The important point is not to specify too much detail, because you need to give the autonomic system room to move." If something breaks, the autonomic system needs flexibility in order to fix the problem.

Autonomic systems require a lot of trust in the system. The system automatically makes some serious decisions for you. "System administrators are not normally happy giving that kind of freedom to the system. You want the system to decide things for you but you want to be able to review them and adjust them to make sure they are right." Autonomic systems will need to provide feedback as to why something has happened. "You should be able to ask the system, 'Why have you put that there?'" Some mechanism for reviewing system actions and tuning the policy implementation for future actions needs to be provided.

What Anderson is seeking is a compromise between the two extremes of, on the one hand, a complete expert system that can be given high-level policy goals and perform all reasoning and logic decisions, generating all individual assignments for all machines and services, and a solution that is based on hand-crafting (or scripting) the low-level specifications for machines and services required to deliver the specified policy. What we want is some autonomics but not a completely unpredictable system.

"The autonomic system has to be able to change all aspects of a system configuration dynamically. UNIX was never designed to be re-

configured on the fly." UNIX has all sorts of config files in all sorts of formats; services may need to be stopped and re-started in order to make certain changes. "This is a big problem in incorporating autonomics into system configuration."

Anderson concluded by reviewing the lcfg system, analyzing where it has useful autonomic capabilities and where it falls short. He pointed to the http://www.lcfg.org Web site and the LISA '03 Gridweaver paper for those who want to explore the complete details.

## GENERAL SESSION PAPERS: SWIMMING IN A SEA OF DATA

*Summarized by G. Jason Peng and Wanghong Yuan*

### *Email Prioritization: Reducing Delays on Legitimate Mail Caused by Junk Mail*

*Dan Twining, Matthew M. Williamson, Miranda J.F. Mowbray, and Maher Rahmouni, Hewlett-Packard Labs*

Matthew Williamson discussed the motivation for this paper. In particular, he described the delay problem caused by junk emails, the distribution of junk mails, and the source of junk mails. Dan Twining then presented the proposed approach, which combines pre-acceptance (header scanning) and post-acceptance (content scanning) to predict the next message type based on sending history. The pre-acceptance method maintains the number of good and total messages and tells if a server is good based on the ratio. The system is implemented in a lightweight manner and shows good results on a real system.

### *Redundancy Elimination Within Large Collections of Files*

*Purushottam Kulkarni, University of Massachusetts; Fred Douglis, Jason LaVoie, and John M. Tracey, IBM T.J. Watson Research Center*

Storage needs keep growing as per-byte cost gets cheaper. The goal in storage is to increase efficiency by reducing redundancy. Current techniques (compression, duplicate block-and-chunk suppression, and resemblance detection) have shortcomings. Purushottam Kulkarni proposed a technique called Redundancy Elimination at Block Level (REBL), which first detects duplicate chunks and encodes blocks using the resemblance technique. This paper also evaluates five techniques to quantify the effectiveness of REBL.

### *Alternatives for Detecting Redundancy in Storage Systems Data*

*Calicrates Policroniades and Ian Pratt, Cambridge University*

Calicrates Policroniades introduced the benefits of redundancy elimination and previous techniques for redundancy elimination, and then compared three frequently used techniques: whole-file content hashing (WF), fixed-size blocking (FSB), and content-defined chunks (CDC). The results show that in terms of compression ratio, CDC is the best, FSB is almost as good, and WF is the worst. But when compression, processing overhead, and storage overhead are considered, however, no one solution wins.

## ADVANCED SYSTEM ADMINISTRATION SIG: SYSTEM ADMINISTRATION: THE BIG PICTURE

*Summarized by Rob Martin*

### *The Technical Big Picture*

*Alva Couch, Tufts University*

Each fall at Tufts University, Professor Alva Couch presents a talk to his students on the Big Picture in system administration. In Couch's words, "Where are we going? What are we going to do? How is it going to work? What is going to be the benefit?" This year at USENIX Tech '04, Professor Couch let us in on a preview of his "technical briefing on next year's big picture" talk for his students.

According to Couch, the future of system administration is about "cost models" and "supporting the enterprise mission." Couch summarized it this way: "Based upon a cost model, we can re-define good system administrating. That idea is rather cosmic, because what we are doing right now, what we consider 'good' right now, I would claim does not make sense with respect to any cost model. . . . Looking at things from a broader perspective of lifecycle costs, we get a better idea of whether we are doing our jobs."

Professor Couch refers to traditional SA thinking and practice as "micro-scale reasoning" and "the bottom-up approach." He includes "adhering to practices, process maturity, six nines at the server, closing tickets quickly, reducing troubleshooting costs" as examples of micro-scale thinking. The opposite of this is the "top-down approach": "In a top-down approach we start at the organization and mission and work down. It turns out that starting at that point and thinking out the whole nature of the profession leads to different conclusions and that's the subject of this talk today."

Professor Couch lists some observations drawn from "macro-scale thinking": "System administration enables particular things. It enables missions. It supports plans. It manages resources. It enforces policies. There is a very high level at which, Burgess says, 'the system administrator manages human computer ecologies.'" Professor Couch says this macro-scale thinking will lead to some sacrilegious ideas. "Six nines for the mission does not require six nines at the servers. . . . There is a fundamental idea that we build six nine infrastructures upon six nine servers and six nine foundations. That is actually not true. Meta-stability is enough. Perceived stability is enough. Security that

compromises mission can be inappropriate and counter-productive. Security is not an end unto itself. It is part of a larger mission picture and availability and security have to be balanced."

Three recent published papers that "got a lot of flak in the last LISA and LISAs before" got Professor Couch thinking this way. The papers were on the cost of downtime (Patterson); the timing of security patches (Beattie et al.); and. resource management without quotas for specific users (Burgess).

Patterson says downtime can be quantified. Professor Couch "cannot believe the resistance that this idea got. Nobody ever talked about cost models before. And maybe because of that this was a very controversial idea."

Couch then talked about the paper that "caused a riot." "Beattie et al. showed that if uptime was important because downtime was expensive, then waiting a couple of weeks to apply a new security patch was probably optimal. . . . The thinking was about a cost model; it was not about simple religion, about just applying things because you are supposed to, but about understanding that patches of a problem in the very first case . . . have problems themselves and that waiting a couple of weeks to apply all of them was a better enterprise strategy.

"Finally, another very bizarre idea: protect useful work instead of limiting people. Burgess actually proposes a game theoretical approach for quotas. The idea of the game is extremely simple. You have a very simple strategy for deleting pigs and you counter every strategy the user could use to defeat you. It uses random scheduling for cleanup to beat the wily user.

"These three examples have common attributes. They consider the broader picture of enabling work and mission rather than fixing systems. They consider costs and values rather than just considering the cost of implementing a change. They also include the cost of not doing things, as well as the cost of doing things."

Couch proposes that we have to change the way we are doing SA: "In pursuing micro-scale perfection we are pursuing a private game, and nobody except us cares. We have to play a different game. . . . Micro-scale system administration as we know it is doomed. The idea of pursuing six nines at the server will be a solved problem in 10 years." He referred to the previous USENIX Tech session on autonomics and said, "We are beginning to understand how to automate installs, automate deployments, and automate monitoring and recovery. In the near future much of this will be automated. Unfortunately, system administration is subject to outsourcing. . . . But we can't . . . automate macro-scale thinking. That's the future of system administration. Being able to take these boxes with six nines and make them talk to each other to support an enterprise mission.

"The future will be about understanding cost and value and how they relate. It's going to be about interacting with middleware. It's not about supporting users; it's about supporting missions."

Professor Couch says the new challenge is to "take the human mission and turn it into something the machine can understand." To do this we will need to research new areas, such as economic models to describe "making day-to-day cost-value decisions." Professor Couch suggests we need to ask, "How does one best quantify the value of mission support? How does SA work impede or aid mission? Is anything you are doing getting in the way of mission and how can you stop?"

He concluded by reminding us what can't be automated and outsourced in SA: "We are here at this conference because one can't outsource community. Outsourcing affects and limits many things. One cannot outsource the value of people being together in one place and thinking about a common problem. That will remain a factor in system administration I think for as long as we live."

## GENERAL SESSION PAPERS: NETWORK PERFORMANCE

*Summarized by Sudarshan Srinivasan*

### *Monkey See, Monkey Do: A Tool for TCP Tracing and Replaying*

*Yu-Chung Cheng, Stefan Savage, and Geoffrey M. Voelker, University of California, San Diego; Urs Hölzle and Neal Cardwell, Google*

The authors describe Monkey See, Monkey Do (MS-MD), a tool that generates realistic client requests in order to test changes to the back end of Google search engines. Currently, changes to servers are tested by either using synthetic (and consequently unrealistic) workloads or real users, making it risky and effort-consuming. The motivation for developing MS-MD is to overcome the shortcomings of existing approaches of testing.

The tool has two phases of operation—the Monkey See phase, where it observes real network connections, measuring network traffic parameters along the way, and the Monkey Do phase, where it generates realistic workloads based on the previously observed metrics. The recorded parameters include the HTTP header, query parameters, delay ACK policy, and other measurable quantities (such as response time). All tracing is done in front of the server farm, and the authors assume that congestion—i.e., queueing of requests—happens, if at all, only along the data path. They also assume that the Web servers themselves are well provisioned and that there is no congestion in the intranet. Caching

behavior is also recorded and replayed by MS-MD.

They evaluate the tool with respect to two questions: how accurately it reproduces the workload, and how accurately it predicts server performance with changes effected. Results show that the measured times without changes to the kernel match up more or less between the original run and the simulation using MS-MD. The tool is more accurate when the RTTs are small; they ascribe this behavior to the fact that the client emulators are on Linux systems, which have a more aggressive ACK policy than traditional Windows clients. Experiments also show that the tool accurately predicts changes in network behavior when services are changed. The tool works for Google, and the authors contend that it will also be usable in other domains.

### *A Transport Layer Approach for Improving End-to-End Performance and Robustness Using Redundant Paths*

*Ming Zhang, Junwen Lai, Larry Peterson, and Randolph Wang, Princeton University; Arvind Krishnamurthy, Yale University*

Ming described mTCP, a transport-level network protocol developed by the authors for aggregating the bandwidth of multiple heterogeneous paths between two hosts. Bandwidth aggregation provides the benefits of improved performance compared to individual network connections, while also improving the resilience of the aggregate connection. The main challenges for providing effective bandwidth aggregation are congestion control, congestion sharing, recovery from failed paths, and selecting which paths to use for packets dynamically.

mTCP uses a single send/receive buffer for all connections, along with per-path congestion control. Packets get striped across the various possible links. This leads to a greater chance of packets getting reordered, though along each channel packets are still in order. This generates too many DUP ACKS. The problem of packet reordering is solved by using SACK TCP. mTCP uses an extended scoreboard algorithm to figure out which packets have been received and which are outstanding. Packets are sent in the order in which they are queued, and the choice of channel is based on proportional scheduling.

For handling shared congestion, mTCP drops one or more of the shared connections in the presence of congestion, so that single-channel connections do not suffer at the expense of aggregated connections. Shared connections are detected by studying the correlations between the different fast retransmissions—closely related fast retransmissions between two links point to a shared connection. For path selection, overlay networks are used to create candidate paths from which a subset of paths is selected greedily with the minimum common links between them. The greedy algorithm chooses paths that are most disjoint so that there is minimum interference between the paths in terms of performance impact and the effect of failed links.

Performance measurements show that the throughput of mTCP is more or less cumulative of the individual network throughputs, as it should ideally be. Separate per-path congestion control provides better throughput than combined control. The failure-detection and recovery mechanisms adopted by mTCP work effectively, allowing the network to recover seamlessly from the failure of one or more of the links. Finally, the throughput of the mTCP system is significantly better than individual paths in the presence of congestion.

### *Multihoming Performance Benefits: An Experimental Evaluation of Practical Enterprise Strategies*

*Aditya Akella and Srinivasan Seshan, Carnegie Mellon University; Anees Shaikh, IBM T.J. Watson Research Center*

Multihoming is a frequently adopted strategy in enterprise networks for improving performance as well as availability of the network. Route controllers are used to provide the required performance and availability characteristics. Aditya presented the results of experiments conducted by the authors to evaluate the performance benefits achievable from commercially available multihoming network solutions. This work extends an earlier study, which showed potential performance improvements of up to 40% compared to no route control for a three-ISP network under ideal route control. Aditya also presented a simple near-optimal greedy route control algorithm.

The route controller needs to monitor performance on the ISP links and choose the best link to send packets through based on the performance of the ISPs at that time. It also needs a way to direct traffic once this choice is made. The authors use EWMA (exponential weighted moving average) to track average performance of each of the ISPs and decide on the best path. While redirecting outbound traffic along a particular ISP's network is trivial, redirecting incoming traffic is a little more difficult. In-bound traffic from within the enterprise is redirected using NAT, and externally originated traffic is handled by modifying the DNS entries.

While monitoring the ISP links' performance, only the traffic to the top Web servers is observed, since they account for most of the network traffic. The actual monitoring can be either passive or active. In passive measurement, the route controller periodically measures the turnaround time (time between the last HTTP request and the first response); this is an estimate of the RTT of the network link. Route controllers doing active monitoring initiate out-of-band probes to get

network performance metrics. Sliding windows and frequency counts are two methods used to decide the frequency of monitoring.

The performance measurements show that even simple passive monitoring of the network connections offers significant performance gains compared to no route control. Active monitoring outperforms passive monitoring, but only by a small margin. The use of history in EWMA does not offer any performance benefit; the current performance of the network connections is a good indicator of future performance. With regard to the effect of the frequency of sampling on performance, results show that a very small sampling interval is harmful to the performance because of frequent changes to NAT and DNS entries, while a too-large sampling interval may lead to stale values being used for the network metrics.

## ADVANCED SYSTEM ADMINISTRATION SIG: LARGE STORAGE

*Summarized by Adam S. Moskowitz*

### *Autonomic Policy-Based Storage Management*

*Kaladhar Voruganti, IBM Almaden Research*

Kaladhar Voruganti discussed his work on the SMaestro Network Storage Planner, which is part of IBM's autonomic computing effort. Despite advances in technology, storage resources tend to be poorly utilized, it is difficult to map application needs to storage systems based on their capabilities, and the number of system administrators needed to manage storage has not decreased (despite claims of easier configuration and management). Customers are moving to SAN and NAS solutions, but these problems are not going away. Voruganti believes that there are both process and technological aspects to the problems (and their solutions).

There is currently virtually no automated solution to analyzing a customer's environment (with respect to storage needs), no way to translate high-level business goals into technology requirements and policies, no way to plan a storage infrastructure based on these needs and requirements, and no way to monitor the solution to ensure that it does not violate the stated goals and requirements. Voruganti's work is limited to a tool to helping plan the storage infrastructure. He was also careful to note that this tool is intended to make things easier for the architect, not replace him. There are many daunting tasks facing an architect when planning a storage infrastructure: interoperability concerns, best practices to follow, gathering data from multiple sources (all of which use different reporting mechanisms), performing complicated "what-if" analysis, and validating that everything was done correctly. SMaestro is intended to make at least some of this easier.

Typical application requirements include I/Os per second (both average and peak), bandwidth, percentages of random/sequential reads/writes, MTBF, maximum acceptable outage times, encryption, integrity, wrote-once support, recoverability, retention times, scalability, and (of course) cost. All this must be balanced against typical storage system capabilities such as latency, I/O rates, availability, points of failure, "hot upgrade," reliability, and (here it is again) cost. SMaestro uses templates for both categories as well as for virtualization software, backup/restore software, and SAN file systems. Policies are written in plain English and then translated, using data models, into something that can be used with the templates to suggest a storage architecture and to verify that the proposed architecture meets the requirements.

### *Experiences with Large Storage Environments*

*Andrew Hume, AT&T Research*

In his talk, subtitled "Inside an inode, no one can hear you scream," Andrew Hume discussed his experiences trying to actually deal with large quantities of data. How large? In one case his project collected between 200GB and 700GB of new data every day. Hume was quick to point out that "20TB just doesn't go as far as it used to." Hume also mentioned improvement in compression algorithms, some of which study the data and then routinely deliver ratios well above 90%. [Summarizer's note: Some of these compression programs were written by Dave Korn; when I asked him if they would be released to the public, he said, "We want to, we're trying to, but we're not there yet."]

To process this data, Hume uses no RAID, no SAN or NAS, no distributed filesystems, and no special hardware save for a high-speed low-latency network and suitable host adapters on each node. Instead, data is broken into chunks, shipped to the next available node in the cluster, and then processed locally. In doing this sort of thing, Hume learned that networks and disk controllers are not as reliable as previously thought. To handle these unreported errors, all data is checksummed before and after each transfer, sometimes more than once. Part of Hume's work is trying to be able to prove that files are "correct" to a standard high enough to be accepted in court. In part this is being driven by legislation such as the Sarbanes-Oxley Act.

Hume claims that backup systems are getting worse; they now have to buy new backup hardware about every three years. His project currently uses AIT-2, which he finds unacceptably slow. By comparison, computers get faster and more reliable every time he buys them.

In the Q&A session, as a comment on a question from an audience member (to Voruganti), Andrew claimed to be from "the Ken Thompson school of thought on expert systems: there's table look-up, fraud, and grand fraud."

## PLENARY SESSION

*Summarized by Richard S. Cox*

### *Network Complexity: How Do I Manage All of This?*

*Eliot Lear, Cisco Systems*

Eliot presented an appeal for work on network management. While providers were previously concerned with only a few very expensive devices, we now have a huge number and variety of devices ranging from routers and switches to laptops and phones, all of which need to be monitored and managed in order to support critical network applications.

The first step in managing is to know what is connected to the network. Unfortunately, today's discovery mechanisms won't work for the millions of devices on future networks. Instead, devices will need to "call home" or send notification when their status changes. This presents issues from what to name the device to determining where to send the notifications. In some cases, there may be multiple interested parties: for example, an ISP, a VPN provider, a voice service provider, as well as the customer, may all be interested in updates from an IP phone; managing this while respecting privacy and security is a major challenge.

Having found the network components, we next need to determine their current status. As a bright point, standards for monitoring and state retrieval from individual devices, such as SNMPv3 and syslog, are maturing and investment is being made in tools. Even simple devices now support an SNMP interface. However, dealing with the large amounts of data generated by all these devices is still an unsolved problem. Correlating reports from multiple sources to identify a fault or anomaly is important, both to limit the amount of information presented to a human and to avoid burdening the infrastructure with excessive management traffic. Some support from the networking hardware—for example, programmable aggregation in the routers—might be useful here.

Having determined the state of the network, closing the loop requires the ability to control the devices. Here, standards are less advanced, though a certain amount of that can be attributed to the cutting-edge nature of the field. The NETCONF protocol is an effort to provide a common transport protocol and syntax for configuration, but vendor-independent configuration schemas are still unspecified.

## GENERAL SESSION PAPERS: OVERLAYS IN PRACTICE

*Summarized by Ningning Zhu*

### *Handling Churn in a DHT*

*Sean Rhea, Dennis Geels, and John Kubiatowicz, University of California, Berkeley; Timothy Roscoe, Intel Research*

■ *Awarded Best Paper*

According to statistics from real P2P networks such as Kazaa, churn ("the continuous process of node arrival and departure") is prevalent in real life. Through experiment on ModelNet—an emulated network—the authors show that several important distributed hash table (DHT) variances (e.g., Tapestry, Chord, and Pastry) all failed to handle churn very efficiently.

This work identifies and explores three factors affecting DHT performance under churn: reactive versus periodic failure recovery, message timeout calculation, and proximity neighbor selection. Results show that periodic recovery is better than reactive recovery; TCP-style timeout calculations outperform those based on a virtual coordinate scheme; and simple global sampling is as good as other much more sophisticated schemes in neighbor selection.

They've used the DHT implementation Bamboo, which is derived from Pastry but has been enhanced with the above techniques to handle churn. The code and additional information can be found at http://bamboo-dht.org.

Q: Is TCP-style timeout doing well because of absence of background traffic?

A: To some degree, the answer is probably yes, although the benchmark itself also creates some load imbalance. In any case, background traffic would probably hurt performance of a virtual coordinate scheme as well.

Q: How did you have a good timeout for the return path, and how did you measure latency when the return path was different from the lookup path?

A: There is an ACK for each lookup message on every hop to get latencies, and a conservative timeout is used for the return path. There is an average of five hops in the lookup and only one return hop, so this conservative estimate isn't too far off. We could not have explored the possibilities of using virtual coordinates for only the return hop, or have the return path be along the lookup path.

Q: In this paper, what is the dimension of virtual coordinate space?

A: It is a 2.5 dimensions space with $x$ and $y$ in a plane and $z$ above the plane. The distance between ($x1$, $y1$, $z1$) and ($x2$, $y2$, $z2$) is Z1 plus Z2 plus the square root of $((x1–x2)^2 + (y1–y2)^2)$. The latest Vivaldi work seems to indicate that this is a good metric.

Q: Why does Tapestry work fine in simulation but not so well in a more realistic network emulation?

A: Because Tapestry has no leaf set; Tapestry really needs to recover quickly from routing table neighbor failures. This problem leads it to be built to recover reactively, therefore to suffer from the same problems that Bamboo/Pastry exhibits from reactive recovery.

Q: Have you tried some sort of periodic/reactive hybrid?

A: It is really hard to do without reverting to the policy of increasing traffic under stress. There's probably some middle ground, but we're not sure what it is yet.

***A Network Positioning System for the Internet***

*T.S. Eugene Ng, Rice University; Hui Zhang, Carnegie Mellon University*

Knowledge about network distance is essential for the performance optimization of a large distributed system. For an $n$-node system, directly computing the delay between each pair of nodes takes $O(n^2)$ time. The author proposes to solve the scalability issue by building a network positioning system (NPS) using a Euclidean space model. Each node infers its network coordinates by measuring their distance to several reference points, and network distance between any pair can then be calculated by their network coordinates.

In building such a network positioning system, there are many practical issues, including system bootstrap, how to support a large number of hosts, how to select reference points, how to maintain position consistency, how to adapt to Internet dynamics, and how to maintain position stability. The system is evaluated on PlanetLab with 127 nodes using an 8D Euclidean model; results showed that position accuracy was fully maintained through the 20-hour testing period.

Q: Is there any technique from NTP that NPS can also benefit from, considering that NTP and NPS have some common issues to deal with?

A: I'm not very familiar with NTP and therefore have no clear answer.

Q: Why did you choose to use 8D Euclidean space? Is there anything particularly important about the number of dimensions?

A: From my experience, 6 to 8 dimensions would all be fine. A too-low number would not yield the desired accuracy, and a too-high number increases the calculation complexity.

Q: Has the system been tested on a practical Internet domain other than PlanetLab, which has a more artificial environment? If yes, what was the result?

A: No. But I expect the scheme to work well on practical Internet domains, too.

Q: Why did you choose Euclidean space instead of another model?

A: Several other geometry models were explored; there wasn't much difference in the result.

***Early Experience with an Internet Broadcast System Based on Overlay Multicast***

*Yang-hua Chu, Aditya Ganjam, Sanjay G. Rao, Kunwadee Sripanidkulchai, Jibin Zhan, and Hui Zhang, Carnegie Mellon University; T.S. Eugene Ng, Rice University*

Internet multicast has been studied for many years; the protocol design and evaluation were mostly based on static analysis and simulation. ESM (End System Multicast) is the first mature and deployed system to use Overlay Multicast for broadcasting video and audio streams. The system has had four publishers and 4000 users, providing unique experiences on Internet broadcasting.

ESM requires no change to network infrastructure; the end hosts in an ESM system are programmable to support application-specific customizations. ESM is a distributed and self-improving protocol optimized for end-to-end bandwidth. It supports heterogeneous receivers, NATs, and firewalls, and has user-friendly managing tools. The results indicate that the overlay tree built by ESM is quite optimized and well structured, and 90% of the users get 90% of the bandwidth. In this paper, Kay Sripanidkulchai presented a concept called "Resource Index" to quantify the bandwidth utilization in the system. When Resource Index indicates the system resource utilization is high, some users experience degradation of video quality. One important lesson learned from ESM is that there are a lot of NATs in the system (70%), which ties up resources and causes poor performance.

The ESM system can be accessed on line at http://esm.cs.cmu.edu.

Q: There have been several recent proposals, such as CoopNet and SplitStream, to use multiple trees for streaming. From your experience do you think multiple tree approaches are necessary?

A: From what we've seen, single trees seem to do fairly well, though for larger-scale groups, multiple trees may become useful.

Q: How do you deal with the data proxy server in your system?

A: It is counted as NAT behind firewall.

Q: How do you avoid "free rides," i.e., when a user lies about the resource that he or she is going to contribute to the system?

A: ESM uses actual measurement instead of trusting a user's claim.

## GENERAL SESSION PAPERS: SECURE SERVICE

*Summarized by Wanghong Yuan*

### *Reliability and Security in the CoDeeN Content Distribution Network*

*Limin Wang, KyoungSoo Park, Ruoming Pang, Vivek Pai, and Larry Peterson, Princeton University*

KyoungSoo Park first introduced CoDeeN, an academic content-distribution network on PlanetLab, and its security problems. The root of these problems is that CoDeeN has no end-to-end authentication. KyoungSoo then described their approach to security, which includes multi-level rate limiting and privilege separation. They achieve reliability by using active local and peer monitoring. In addition, he discussed DNS problems solved via mapping objects in the same proxy and CoDNS. Finally, he summarized the lessons and future work, including robot detection, CoDeploy and CoDNS. More information is available at http://codeen.cs.princeton.edu.

Q: What causes the DNS problems?

A: Local DNS server overload.

Q: What are other solutions for accessing local content?

A: More efficient approaches, with more information; currently the privilege separation is simple.

### *Building Secure High-Performance Web Services with OKWS*

*Max Krohn, MIT*

The motivation story is SparkMatch version 1, which crashed with 500,000 signups. Version 2 solved some problems in the database, but there were too many connections. Version 3 in 2002 distributed database but the development cycle was too long. Max summarized the desired Web service features: thin fast server, smart gzip support, small number of database connections, memory reclamation, and an easy and safe way to run C/C++ code, and mentioned that the major problem is dynamic content. Currently, Apache/PHP does not work well, due to the poor isolation.

Max then described their system, called the OK Web server (OKWS). Its design is that of a multi-service Web site, and its isolation strategy is the least-privilege principle. Max also gave an example of how to build a Web service with OKWS and illustrated how the isolation works in OKWS in detail. OKWS is implemented in C++ with SFS libraries, database translation libraries, and Perl-like tools. The key point is one process and one thread for one service, without synchronization. SparkMatchv4 is built using OKWS, which compared favorably to Apache, Haboob, and Flash. The source code is available at http://www.okws.org.

Q: Is there an advantage for not maintaining the database pool?

A: Yes, based on observation of Apache.

Q: Why not replace script? Why develop a new Web server?

A: Security problems in Apache.

Q: How do you support two requests sharing the same service?

A: It's possible to do this; the paper has more detail on how.

### *REX: Secure, Extensible Remote Execution*

*Michael Kaminsky, Eric Peterson, M. Frans Kaashoek, and Kevin Fu, MIT; Daniel B. Giffin, David Mazières, New York University*

The motivation is that remote execution is important but there are features not widely available in current tools. The major problem addressed in REX is locating the simplest abstraction that can support all of these features. Michael described how to establish a session, run a program, pass file descriptors, use X Window system forwarding, connect through NAT and dynamic IP address, and forward restricted credentials. The evaluation tries to answer two questions: Is REX reliable and are there any architecture benefits? More information is available at http://www.fs.net.

Q: Are there problems in file access?

A: Only in some access, such as write and read.

Q: What is the relationship between TCP and channel?

A: There is one TCP and there are multiple channels.

## USEBSD SIG

*Summarized by Adam S. Moskovitz*

### *The NetBSD Update System*

*Alistair Crooks, The NetBSD Project*

Alistair Crooks described a system for downloading and installing binary patches, similar in many ways to Microsoft's Windows Update Facility, but for use on any number of platforms. Crooks' system runs on a variety of platforms, including the *BSD variants, Mac OS X, Linux, and Solaris. Like the Microsoft system, NetBSD Update is easy to use and gives the user three options for automatic behavior: inform the user; inform and download appropriate packages; inform, download, and install the packages/patches. Crooks uses a file on the update server to list packages for which vulnerabilities exist and a program that runs on the target system to say which of those packages is present and should be updated.

NetBSD Update includes other important features, the most significant (in my mind) being the ability to digitally sign update packages and the user's ability to accept or reject those updates based on the validity of the signature(s). Unfortunately, like so many other systems, the lack of a widely accepted public key infrastructure means this feature is still a bit more cum-

bersome than it ought to be (which, of course, should not be considered a shortcoming of Crooks' work).

Another feature is the ability to undo the effects of an update. NetBSD Update automatically preserves all files that will be overwritten and stores them for the user in case the update causes more problems than it solves, or if worse vulnerabilities are found in the update than existed in the unpatched system. [Summarizer's note: Of course, this sort of thing has never happened—the phrase "the patch for the latest jumbo patch" is merely a joke among system administrators.]

### *A Software Approach to Distributing Requests for DNS Service Using GNU Zebra, ISC BIND 9, and FreeBSD*

*Joe Abley, Internet Systems Consortium, Inc.*

[Summarizer's note: This talk/system deals with certain aspects of routing that go beyond my general knowledge of the subject; if something you read doesn't make sense, the error is almost certainly mine and not the speaker's.]

Joe Abley described a system for distributing DNS requests across multiple hosts without the use of dedicated load balancers, by using a "service address" (sometimes called a "virtual IP address"), anycast, and the Equal Cost Multi-Path (ECMP) feature of the OSPF routing protocol. This system is currently in use for the F root name server (run by ISC—the Internet Systems Consortium), which also provides slave service for 30–40 ccTLDs. Currently, 24 nodes across California and in New York, Tokyo, and Stockholm make up what appears to be a single root name server.

Individual nodes in the cluster are configured with unique unicast IP addresses, and with the service address on the loopback interface. Hosts inform the routers of their readiness to answer requests via OSPF Link State Advertisements; simple wrapper informs the router when an instance of "named" fails (internal assertion failures are set to dump core and exit).

For UDP-based DNS queries, it doesn't matter which node in a cluster provides the answer. For TCP-based operations like zone transfers, all packets in the transaction must be routed to a single node; Abley uses a combination of "flow hashing" (via Cisco Express Forwarding or the "load-balance per-packet" feature on Juniper routers), avoiding ECMP routes for stateful transactions, and using BGP.

## GENERAL SESSION PAPERS: THE NETWORK-APPLICATION INTERFACE

*Summarized by Calicrates Policroniades*

### *Network Subsystems Reloaded: A High-Performance, Defensible Network Subsystem*

*Anshumal Sinha, Sandeep Sarat, and Jonathan S. Shapiro, Johns Hopkins University*

Anshumal Sinha introduced his talk with several issues observed in in-kernel monolithic network systems: security (they represent a single point of failure), maintainability (robustness-critical code is large and difficult to maintain and debug), and flexibility (lack of support for the simultaneous existence of multiple protocols, complexity to do application-specific optimizations). He stressed that monolithic network systems are only used because of their performance benefits. The author mentioned that previous user-level implementations have failed to deliver sufficient throughput, noting, however, their hypothesis that earlier systems failed to provide an appropriate solution to a key problem: performance degradation resulting from data copying from one address space to another in order to provide protection domain boundaries efficiently. Failing to properly manipulate data from different protection domains degrades the performance of the system.

The author then presented the methodology to evaluate their solution. Based on the EROS micro-kernel to support domain factoring, they built two network subsystems to evaluate the costs due to the user-level implementation, the domain factoring, and the microkernel performance. In particular, they built an EROS-based monolithic network subsystem and an EROS-based domain factored network subsystem. Anshumal explained in detail the implementation of both systems and their distinctive features. When presenting experimental results and evaluation, he included a conventional Linux in-kernel network stack as a reference baseline for performance comparison. A detailed explanation of their results in terms of latency and throughput showed that the performance exhibited by the domain factored network subsystem was comparable or close to the other two strategies (EROS monolithic approach and conventional Linux network stack).

Anshumal concluded his talk by remarking that domain factoring is more feasible than previously assumed, that the instruction cache plays a significant role in the performance of the system, and that factoring provides the basis for defensible systems.

### *accept()able Strategies for Improving Web Server Performance*

*Tim Brecht, David Pariag, and Louay Gammo, University of Waterloo*

Tim Brecht discussed how particular strategies to handle connection requests affect the performance of Web servers. He began by mentioning how to improve Web servers' performance by modifying their corresponding accept strategies. He mentioned that an adequate solu-

tion should not only improve Web servers' peak performance but also be able to maintain it even under overload conditions with a large number of connections. He presented throughput results for three architecturally different Web servers: the event-driven, user-mode micro-server (39–71% improvement); the multi-threaded, user-mode Knot (0–32%); and the kernel-mode TUX (19–36%).

Tim stressed that current multi-accept servers overemphasize acceptance of new connections and ignore the processing of existing connections. Second, he mentioned that his work aimed to reduce the gap in performance typically seen between kernel-mode and user-mode servers. Next, he described in detail the three architectures that were analyzed in the paper and explained how the accept-limit, which defines an upper limit on the number of connections accepted consecutively, affects each technique's performance.

The presenter made a careful analysis of the impact of varying the accept-limit on each of the servers based on their experimental results. In his experiments they used two workloads: a one-packet workload and a SpecWeb99-like workload that uses httperf to generate overload conditions. The experimental results presented by the author were organized by server performance, queue drop rates, and latencies observed under different accept-limit policies. Tim mentioned that it is necessary to ensure that Web servers accept connections at sufficiently high rates so that a balance between accepting and working times can be adequately established. Finally, he said they were able to demonstrate that performance improvements can be obtained by modifying the accept strategies used in Web servers.

### *Lazy Asynchronous I/O for Event-Driven Servers*

*Khaled Elmeleegy, Anupam Chanda, and Alan L. Cox, Rice University; Willy Zwaenepoel, EPFL, Lausanne*
*Presenter: Khaled Elmeleegy*

Khaled Elmeleegy began his presentation explaining why event-driven architectures are used and the difficulties that programmers experience when developing high-performance servers using existing I/O libraries. He remarked that current I/O libraries have an incomplete coverage or leave to applications the burden of state maintenance. In contrast, Lazy Asynchronous I/O (LAIO) is a good option to develop high-performance, event-driven servers with less programming effort, because it covers all possible blocking I/O calls, creates a continuation only when an operation actually blocks, and notifies applications only when a blocking operation has been wholly completed. Next, Khaled explained the way in which event-driven servers generally work and the role that event-handlers play in their operation; he also mentioned how blocking operations degrade server throughput. He presented LAIO as a solution to the typical blocking problems seen in event-driven servers and proceeded to describe the LAIO API, their functions, and implementation.

After introducing the two different workloads used in their experiments, Khaled compared LAIO's throughput and performance against other I/O libraries. With this purpose, the authors modified the networking and disk I/O strategies of the Flash Web server and measured the results obtained for different versions of the server. In situations where performance was comparable, the complexity of the programs decreased because it is not necessary to write handlers or to maintain state as in conventional non-blocking I/O, which finally leads to a reduction of the amount of code that needs to be written.

Khaled finished his talk by highlighting LAIO's generality (covering all the I/O calls) and simplicity (requiring fewer lines of code without handlers or the need to maintain state). In terms of throughput, LAIO meets or exceeds the performance of other methods. During Q&A, the audience mainly focused on comparing LAIO with multi-threaded servers; Khaled mentioned, however, they had decided to focus the study on event-driven strategies.

## USEBSD SIG

*Summarized by Matus Telgarsky*

### *Building a Secure Digital Cinema Server Using FreeBSD*

*Nate Lawson, Cryptography Research*

Nate's dense, practical, and often anecdotal talk went beyond general crypto issues and concepts to also discuss the troubles and travails afflicting construction of a digital cinema server and problems with its wide acceptance.

After providing a quick overview of Cryptography Research Inc., Nate presented an extremely cogent diagram that proved one of the strongest bromides of the security circuit—you are only as strong as your weakest link. A graph pitted probability of compromise against effort/cost of attack. The perfect scenario features a deep curve predicting massive effort for even the slightest crack; however, usual circumstances produce an almost inverted curve (symmetric against $y=x$), where bribing employees, using common scripted attacks, and abusing operating system holes is often easy and constitutes the majority of compromises, rather than a crack of an encryption key, a shield many often deem fully sufficient.

Nate continued to describe three fundamental underpinnings of any security paradigm: strength (which encryption provides),

assurance (pragmatic assessment of attack methods, especially easier entrances), and renewability (post-mortem reconstruction).

Traditional analogue cinema (obviously) uses film cameras, is transferred to a digital format for post-processing (effects, editing, etc.), is printed thousands of times (at $3,000 a copy, which degrades after a week of use anyway), and is played on $30,000 projectors. Digital cinema is directly captured to hard drives and obviously omits any tedious conversion; however, distribution and projection standards do not exist—costs are prohibitive, and the market is in flux. Not only does refitting cost around $100,000, but there are even questions whether the theater is responsible for said expenditure. In 2003, only 90 theaters in the U.S. were digital (30 in 2000).

Digi-Flicks enlisted Cryptography Research to design a new security system and, hopefully, extend digital cinema acceptance. Planned design goals were transport independence, thorough use of strong crypto algorithms, multi-factor authentication (i.e., simultaneously utilized smart card, pass code, and key file), flexible authorization policies, reliable playback over imperfect media, and, of course, a rapid development cycle. Amusingly enough, for the 300 target theaters, shipping hard drives was found to be the most cost-effective distribution method.

The sample extant hardware was rather unpleasant—a simple UNIX-like OS over a 33MHz PowerPC with 64MB RAM, too many ASICs, and no documentation. Serial and even Ethernet proved too slow for data transmission, so the SCSI interface was selected to actually transfer the films (FreeBSD was selected partially due to the ease with which the disc access code could be modified to play nice with this chaotic configuration). The encrypted drive would be decoded with a smart card and by dialing in to an auth server. Nate spent a moment covering common pitfalls in encryption selection, including an amusing pair of images presenting a still identifiable pattern in an encrypted image due to naively small block selection.

A recurring suggestion was careful thread-model analysis in order to realistically and sensibly determine circumstances and identify weak points, and then assign design parameters accordingly. Nate detailed many possible scenarios (one-time read-only access, repeated read-only access, one-time read-write access, and repeated read-write access) and the specific dangers within each. Similarly, a good security structure depends on clear top-down study and careful perusal of all possibilities.

***Panel: The State of the BSD Projects***

*Chair: Marshall Kirk McKusick*

*The FreeBSD Project: Robert Watson, Core Team Member, The FreeBSD Project*

*The NetBSD Project: Christos Zoulas, President, NetBSD Foundation*

*The DragonFly BSD Project: Matt Dillon, Project Leader, The DragonFly BSD Project*

All attending BSD parties were able to give impromptu presentations detailing past work, current revelations, and future plans. FreeBSD, NetBSD, and DragonFly BSD were represented; rumblings abounded regarding Darwin and OpenBSD, supporters of which were unfortunately concurrently occupied with other tasks.

Robert Watson's speedy flight through FreeBSD 4 (stable) and 5 (development) built a good measure of excitement and confidence in the impressive list of features stabilizing in the new code. 4.10 has survived healthily with minor security updates and has enjoyed hardened security. 5.X has been in continuous development for five years, with hard work solidifying a new SMP model and threading core, among numerous other improvements. 5.3 features include gcc 3.4, PCI and ACPI work, X.org X server, more fine-grained locking work (including heavy "giant" lock removal in many subsystems), SMP thread scheduler tweaks, and the flexible pf packet filter, among others. SMPng is satisfactorily transitioning from bug and correctness checking to performance tuning.

Next came Alistair Crooks presenting NetBSD, alive since 1993 and eagerly awaiting a 2.0 release. Nascent features include SMP work, scheduler activations, kqueues, wireless drivers, and many other features. A primary goal of NetBSD is to function on many architectures—the netbsd.org sidebar presents an impressive list of 54 disparate devices. Time was also spent describing the package distribution system, pkgsrc, which is easily configurable, consistent, supports multiple versions of installed programs, and currently boasts over 4500 packages.

Matt Dillon discussed his ambitious DragonFly BSD project, which is steadily advancing upon a 1.0 release. Already a veteran hacker (contributor to Linux and FreeBSD, among many other projects), Matt's aggressively progressive plan for restructuring BSD revolves around a message-passing core with a lightweight IPC model. Much of the talk and current focus is extremely low level—for instance, much focus is going into restructuring to optimize cache usage in all subsystems, from the ground up. One corollary goal is minimizing use of fine-grained mutexes. DragonFly is a continuation of FreeBSD 4.X, though a pleasant rapport exists with the parent project, and indeed updates still filter through.

This enthusiastic one and a half hour panel showed the BSD projects to be in excellent condition,

with a dedicated group of hackers backing each—in fact, one of the few times I witnessed the BSD hackers leaving the laptop room (and their diligent hackery) was to attend this informative panel. FreeBSD and NetBSD pointed out their foundations, facilitating donations through cheerily tax-deductible exchanges. DragonFly has not yet formed a foundation, though Matt Dillon would certainly enjoy frequent surreptitious anonymous gifts.

## PLENARY SESSION

*Summarized by Swaroop Sridhar*

### *Thinking Sensibly About Security in an Uncertain World*

*Bruce Schneier, Counterpane Internet Security, Inc.*

Mr. Bruce Schneier delivered his thought-provoking and entertaining talk without any kind of visual aid. He began by saying that we are living in an interesting era called "silly security season." Introducing the notion of all of us being "security consumers," he said that we need to step back and analyze whether this security is really worth it. Is it worth the billions of dollars and the loss of convenience, anonymity, performance, or freedom? Elaborating on security trade-offs, he said that it is well known that trade-offs are ubiquitous, but there is a fundamental security trade-off paradox: We, who claim to be the "most intelligent" species on the face of this earth, always make the wrong security trade-offs. Much of Schneier's talk was based on why this is so and how it could be fixed.

Schneier said that the way to do good security trade-offs is to slow down and have a basis for rational discussion, rather than making quick decisions based on emotions. However, he warned that this could be quite complicated because the meaning of factors such as inconvenience, risk, and privacy are often subjective.

Next, Schneier brought up the topic of risk assessment. He warned that people are always bothered about "spectacular" risks (e.g., risks while flying in a plane) and downplay "pedestrian" or "under control" risks (e.g., risks while driving), which matter much more in our lives. Schneier identified technology and media as two main culprits causing this problem. News, by definition, means that which does not happen every day. The media only show the uncommon happenings, replaying them over and over to create a feeling that they are very common. Technology contributes its bit, obscuring risks by hiding operational details from users.

Again, bringing up the topic of why extreme trade-offs (such as National ID cards) are taken for little gain, Schneier said that security decisions are usually made for non-security reasons. This leads to a notion of an "agenda" among all the "players" of the bigger system, of which security is a part. For example, closing national highways is good according to a police agenda, but bad according to a public agenda. Schneier proposed a model of "Security Utilitarianism"—which leads to the greatest security for the greatest number of people.

Schneier stated that one of the fundamental problems is that we often have no control over the security policies that are implemented. The right kind of security should be worked out by means of negotiations and deliberations. He cautioned, however, that the negotiations should be held at the right time, with the right people. Arguing with a security guard at an airport gate, for example, would be a bad idea. Schneier identified four factors that can effect a change in security norms—government rules and laws, market forces (e.g., refusing to use an insecure OS), technology, and social norms. He said that by turning the above four knobs, we must be able to work out the right kind of security.

Later, Schneier said that we need to accept the risks as real, and try to reduce them. He also proposed that one possible solution is to put the person who can best mitigate the risk in charge of the risk. He illustrated this point with the example of a supermarket cash register, where the customer is "used" to guard against cashier malpractices. Schneier concluded by saying that as individuals we have very little power, but as an aggregate we can achieve a lot toward our collective good.

## GENERAL SESSION PAPERS: UNPLUGGED

*Summarized by Matus Telgarsky*

### *Energy Efficient Prefetching and Caching*

*Athanasios E. Papathanasiou and Michael L. Scott, University of Rochester*

■ *Awarded Best Paper*

Modern operating system design prescribes a plethora of heuristics for caching and prefetching to aid in disk performance, with nary a word about energy savings. Studies attribute 9–32% of laptop energy expenditure to hard disk use; hence, heavy savings in that realm will result in drastic overall improvements.

Traditional prefetching techniques aim to reduce disk access latency by attempting to maintain the working set of an application's disk data cached by replacing unused cache elements with simplistically determined prefetch targets. Unfortunately, this does not preclude the possibility of an application reading and writing at arbitrary times, obviating the possibility of simply tacking on any sort of basic power management scheme. Indeed, many of the tests on a stock Linux kernel

showed 100% of the disk idle times to fall beneath one second, not nearly long enough to enter a disk's power-saving state without incurring a net power efficiency loss owing to energy required for transition.

The presentation detailed the design and implementation of Bursty, a mechanism providing highly speculative prefetching, a kernel interface to hint disk access patterns, and a daemon both to monitor and manage the system. Applications may hint improperly, or lack hints entirely, so the monitor must both generate and judge extant hints. The prefetching is also self-aware—the success rate of the algorithm is constantly measured to determine whether further adjustments are required. Additionally, per-application idle times are irrelevant if they are not in phase between applications; hence, the daemon also attempts to organize these patterns to allow for consistent disk avoidance between all applications. Once the predicted idle period is estimated to beyond the intersection of regular drive use and idle use combined with transition expenditure, the drive is powered down into an appropriate low-power mode. A variety of tests with different applications using a variety of workloads and disk access patterns (and memory configurations—Bursty is hungry!) found 60–80% energy savings, with negligible losses in efficiency.

### *Time-Based Fairness Improves Performance in Multi-Rate WLANs*

*Godfrey Tan and John Guttag, MIT*

Modern Wireless networks theoretically maintain decent throughput when congested, though in practice the common utilization of rate diversity as an automatic signal strengthening scheme causes standard throughput-based fairness schemes to result in unexpectedly poor performance. This paper presents an overview of the problem, design and implementation of a solution—termed "time-based fairness"—and experimental verification.

802.11b was used as the test case. Though traditionally known to sport 11Mbps, the standard also defines three other rates: 1, 2, and 5.5. Vendors use these speeds when packet transmission failure becomes a problem, eventually bumping the speed to the slowest rate, which features the highest signal resilience. Current channel proportioning and access point downlink scheduling techniques result in throughput-based fairness, meaning a slower rate receives a larger portion of channel time, ostensibly aiding the feeble companion but causing the hare to be tied to the turtle.

Time-based fairness apportions channel use equally by time, resulting in much higher possible throughput. The average time for network tasks to complete is also reduced, obviously a benefit to many mobile users and definitely to anyone who would rather have things to do while a laggy transmission completes. The implementation is flexible enough to function properly on extant access points and does not need extensive modification on clients; adoption is easy and backwards-compatible.

### *EmStar: A Software Environment for Developing and Deploying Wireless Sensor Networks*

*Lewis Girod, Jeremy Elson, Alberto Cerpa, Thanos Stathopoulos, Nithya Ramanathan, and Deborah Estrin, UCLA*

The burgeoning study, experimentation, and deployment of wireless sensor network applications is generating a need for full-fledged development suites. EmStar provides just that for 32-bit embedded MicroServer platforms: tools and libraries providing simulation, visualization, and emulation. Other functionality aids in development and testing of IPC and network communication.

Due to time constraints, a large portion of the talk focused on FUSD (Framework for User-Space Devices), which is a kernel module proxy to device file events. FUSD, though new, is already used by numerous applications to simplify communication with device nodes. It is essentially a micro-kernel extension to Linux. At the moment, performance is sufficient but unsatisfactory; read throughput, for instance, can be 3 to 17 times slower than analogous read performance without the FUSD proxy overhead.

EmSim and EmCee are simulation tools; these in turn are modular to allow for easy extension and minimized footprint. EmRun starts up, maintains, and shuts down an EmStar system according to a policy in a configuration file; it features process respawn, in-memory logging, fast startup, and graceful shutdown. All components (more than are listed here) are written with modularity in mind, and code is heavily reused. It has already proved useful in numerous projects at the CENS labs working with a variegated set of hardware.

## SECURITY SIG

*Summarized by Ming Chow*

### *Panel: The Politicization of Security*

*Moderator: Avi Rubin, Johns Hopkins University*

*Panelists: Ed Felten, Princeton University; Jeff Grove, ACM; Gary McGraw, Cigital*

The common theme of this panel was how politicized security, especially that relating to technology, has become. Professor Avi Rubin spoke of his experiences working on the issue of electronic voting (eVoting). He spoke about dealing with policy issues, and about how eVoting has become a partisan, politically charged issue and, as

such, is targeted for abuse. An example is that companies producing eVoting technologies and equipment have strong political ties. The goal from each political party is to "not have the other guy win." Professor Rubin has been on major news sources (e.g., CNN) speaking about technical issues of eVoting, and has received numerous telephone calls from both Democrats and Republicans. Professor Rubin recounted being called to testify in front of Congress about eVoting, and recalled the amount of fighting and bickering on both political sides dealing with the issue. He summed up the current state of politics by saying that "partisanship has never been worse."

Gary McGraw spoke of the long history of politicization of scientific research and development and the degree to which current scientific research and development are influenced by politics (like Galileo and Darwin centuries ago). He stated that security and terrorism are sensitive subjects, and that "we should understand the problem, having worked in an asymmetric situation for years in computer security." McGraw also said that too often "individual rights can be trumped in the name of security" (e.g., DMCA and the Patriot Act).

Jeff Grove has worked with the government on Capitol Hill, and expressed his dissatisfaction on the number of bad laws being implemented, including the DMCA, and the regulation of P2P networks. Grove outlined how the Senate can address and jump on issues and make dumb laws. The problem persists because of bad conclusions, bad assumptions, and lack of basic understanding about technologies. In addition, there is a small handful of powerful players who are effective in influencing the government to create laws fitting their agendas. Bad laws expose developers to liabilities, even when there's no infringement, and provide civil enforcement by encouraging legal actions by the entertainment industry.

Professor Ed Felten spoke about the Digital Millennium Copyright Act (DMCA) and his work, which made national headlines several years ago. Professor Felten stated that the DMCA was created by negotiations in which computer scientists were not involved. His work with advisee John Halderman was discussed—the weak DRM technology created by SunnComm could be bypassed on Windows computers by holding down the shift key. The government was cracking down on Professor Felten's research and he was threatened by the RIAA under the DMCA. Professor Felten settled with both Princeton University and the government by creating educational packets for the government on security research. Professor Felten recalled testifying in Congress about a bill to limit developing tools on decoding technologies, and summarized the atmosphere in one word: "theater." Finally, he gave out his Web site: http://www.freedom-to-tinker.com.

The theme from all of the panel speakers was clear: "We (the computing and scientific communities) need to step up to the plate and educate people on technological issues." The goal can be accomplished by being more involved, by being partisan, and by talking to anyone who is curious. Openness and debate are encouraged and are healthy. It is critical to tell the truth and to convince people about what's really going on. Gary McGraw also said that attacking systems is a necessary part of security and that outlawing attacks makes little sense. Finally, media and politics are great investments: the "Slashdot effect" helps ridicule bad laws, and working even with your local government is a 10–15-year investment.

## FREENIX OPENING REMARKS AND AWARDS

*Summarized by Martin Michlmayr*

*Bart Massey, Portland State University; Keith Packard, Hewlett-Packard Cambridge Research Lab*

Bart Massey and Keith Packard opened the FREENIX track, a forum devoted to free and open source software, by giving a brief summary of papers that were submitted this year. Out of 61 papers submitted, 15 were accepted. The organizers were happy to see that among the accepted papers, seven were from students, and seven were non-US papers. They said that the quality of all submitted papers was very high and that the review process was more formal than in the last few years, adding three external reviewers to the program committee. They also thanked DoCoMo for sponsoring student travel for the conference.

In this opening speech, two awards to papers in the FREENIX track were given. The Best Paper award went to "Wayback: A User-level Versioning File System for Linux," and the Best Student Paper was "Design and Implementation of Netdude, a Framework for Packet Trace Manipulation."

There will be another FREENIX track at USENIX '05 in Anaheim, California. Since future USENIX conferences will take place around April, the deadline for FREENIX submissions is October 22, 2004, rather than in December. More information on the next FREENIX track and Call for Papers can be found at http://www.usenix.org/events/usenix05/cfp/freenix.html.

## FREENIX INVITED TALK

*Summarized by Martin Michlmayr*

### *The Technical Changes in Qt Version 4*

*Matthias Ettrich, Trolltech*
*Linux/Open Source*

Matthias Ettrich, founder of the KDE project and a main developer on Qt, gave an overview of the next generation of Qt, a cross-platform C++ GUI toolkit. Qt supports X11, Microsoft Windows, Mac OS X, and embedded Linux, and offers native look and feel on each of these platforms. Qt provides single-source compatibility: one source code compiles on all target platforms. While Qt mainly offered GUI functions in the past, it is much more than a GUI library these days: It also supports I/O, printing, networking, SQL, process handling, and threading. One aim of Qt is to provide an excellent programming experience.

Qt introduced the signals-and-slot concept in order to allow different GUI components to communicate. You can connect any signal to any number of slots in any module, and communication is done at run time. The sender and receiver don't need to know each other. In version 4, connections can be either synchronous or asynchronous ("equal connections"); this will allow thread communication. Arthur is Qt's paint subsystem, and version 4 will offer several new features: linear gradient brushes, alpha-blended drawing, anti-aliased lines, painter paths, and an OpenGL backend.

Interview is a model/view framework for tree views, lists views, and tables. In the Model-View-Controller (MVC) paradigm, all these components are separated from each other. The model contains data, the view renders data, and the controller transforms interaction with the view into actions to be performed on the model. Qt 4 will introduce semi-transparent windows and flicker-free painting, and it will also implicitly provide double-buffering for all built-in and custom widgets: this is transparent, and no code needs to be rewritten. In addition, it will allow large windows (even modern window systems limit a widget's coordinate system to 16 bit, but Qt 4 won't have this limitation), and there will be improvements in size and performance. Qt 3 was originally designed for desktop computers (with fast CPUs with FPUs, lots of RAM and disk space). On the other hand, Qtopia was designed for embedded systems. Qt 4 aims at merging the benefits of both product lines into one.

In summary, Qt 4 will provide a number of new features that will offer new possibilities for cross-platform development. Ettrich hopes that a first technology preview will be made really soon, with another one following in Q3 2004. A beta of Qt 4 should be released in Q4 2004, with the final version following in Q1 2005.

## SECURITY SIG

*Summarized by Ming Chow*

### *Panel: Wireless Devices and Consumer Privacy*

*Organizers: Ari Juels, RSA Laboratories; Richard Smith, Consultant*

*Panelists: Markus Jakobsson, RSA Laboratories; Frank Schroth, uLocate; Matthew Gray, Newbury Networks*

The panel talked about wireless technologies, including GPS and RFID, and privacy issues concerning them. Frank Schroth discussed uLocate's wireless technology, which enables small business users to view the location of phones in their account, including maps and routes. uLocate's technology is based on GPS on the Nextel Network. The interface on cellular phones is a Java-based application that transmits data to server via UDP. Permissioned users can view information on their phone or via Web. The benefits of the technologies to small business include convenience, efficiency, and safety. The security of the uLocate service consists of two layers: a carrier level and an application level. Mr. Schroth also discussed privacy concerns about the technology, namely, addressable IP addresses, privacy, and leadership and ownership of risks.

Matthew Gray of Newbury Networks discussed his concerns about location-tracking technologies. The goal at Newbury Networks is to see what people are accessing without interfering with larger networks (e.g., Starbucks) and to eliminate such false positives to enhance security and privacy. Mr. Gray noted that security and privacy are in opposition to each other. He said that consumers and regulators must understand the risks of location-tracking technologies.

Marcus Jakobsson of RSA Laboratories listed three ways in which location privacy can be violated: active attacks (keep asking a device), passive attacks (listen to communication from other devices), and remote attacks (infer location from public information). He said that legislation for location privacy is necessary and meaningful. However, such law will be difficult to enforce because detecting abuse by institutions is hard, and it is even harder to detect abuse by individuals. He noted that countermeasures have been proposed but not deployed. In conclusion, Mr. Jakobsson listed several things that must be done immediately: the threats of location privacy must be studied and understood, legislation must be enacted, and countermeasures must be implemented.

Finally, Ari Juels and Richard Smith discussed radio frequency identification (RFID) tags and privacy concerns about the technology. Mr. Juels presented a brief tutorial of the RFID technology: an RFID tag

uses a chip (IC) antenna slightly larger than a quarter. Currently, many people have tools and gadgets that have RFID tags, such as E-ZPass, Mobil SpeedPass, and physical-access cards. RFID tags are seen as next-generation barcodes; Mr. Juels listed the benefits of RFID tags over barcodes (they're fast, efficient, mobile, can uniquely specify objects, and require little computational power). What this means is that the world will consist of billions of $0.05 computers. The major privacy problem concerning RFID technology is that it can be used to profile a person incredibly easily and quickly, providing detailed information, for example, on artificial body parts and other details of a person. Mr. Juels noted that approximately 42% of Goggle hits on a search for RFID contain the word "privacy." The solution to the privacy problem is to kill RFID tags. However, RFID tags are too useful. Mr. Juels concluded his talk by saying that there is serious danger to privacy if the technology is deployed naively, but the danger can be mitigated to strike a technical balance with society.

At the end of the talk, the panel discussed what must be done now to mitigate privacy concerns. One question to the panel was whether policy legislation hurts or helps technical development. A member of the panel suggested that a policy of saving data for 90 days would be sufficient. There was also a discussion about disclosure of information to consumers. Mr. Schroth responded that it has been startling to him how people do not care about disclosing information about themselves: people are willing to give lots of information to companies, including passwords.

## FREENIX SESSION: SERVER

*Summarized by Matus Telgarsky*

### *Migrating an MVS Mainframe Application to a PC*

*Glenn S. Fowler, Andrew G. Hume, David G. Korn, and Kiem-Phong Vo, AT&T Labs*

Rotting at the hearts of many old institutions' organizational frameworks are mainframes and their respective applications. Though the software may be relatively dependable, operational costs are prohibitive (the task emulated within this paper is estimated at $20,000 per month just for mainframe use), and the code consists of thousands, even millions of lines of ancient COBOL and JCL, on a system without a hierarchical file system. Emulating the process is a feasible and cost-effective alternative.

The MVS was to handle a mammoth of data, so a variety of tools were written to efficiently compress it prior to transmission. The OpenCOBOL compiler was extended to handle a few language extensions and different character sets, parse compressed data directly, and also receive a few performance enhancements. An extended sort program was built to enable MVS features, and a flexible JCL interpreter was built with handy features such as ksh script generation. An unsophisticated scheduler was developed to emulate MVS handling of processes.

David Korn took a moment to quip that a 25-year-old tip indicated that sort is optimally performed on a UNIX machine by transferring it to tape, performing it on a mainframe, and transporting it back—yet today the situation is reversed. Two 2.8GHz Pentium 4 machines were used to emulate the mainframes, at under $4,000 total. The 60-hour MVS task took 19 hours on the shiny new silicon. Data transmission ballooned surprisingly to nearly 24 hours due to tapes acting—predictably—unpredictably fussily.

### *C-JDBC: Flexible Database Clustering Middleware*

*Emmanuel Cecchet, INRIA; Julie Marguerite, ObjectWeb; Willy Zwaenepoel, EPFL*

The general trend in modern high-power computing is toward clusters of commodity machines, achieving a significantly superior price-power ratio over more traditional and expensive many-CPU SMP machines. Though many tiers of server applications have extended to utilize this trend, in general RDBMS installations have lagged behind. Limited support has come from Oracle and IBM, but open source databases either relied on simplistic master-slave replication services or other similar compromises.

Clustered JDBC (C-JDBC) is an open source database middleware which abstracts pools of databases into a virtual database, complete with load balancing, query caching, logging, checkpointing, scheduling, authentication, and other features. JDBC is used to connect to virtually any database (as they basically all provide JDBC drivers), and allows for seamless integration of heterogeneous database farms into single resources. Performance has also been considered deeply: For most workloads, increasing nodes results in linear benchmark improvements, meaning superbly minor overhead and excellent scalability.

Fault tolerance and redundancy are not only accounted for with a flexible load balancer, but C-JDBC controllers themselves can be stacked horizontally to virtualize the same databases and seamlessly provide redundancy. Arbitrary trees may be constructed by attaching C-JDBC controllers as client databases to other C-JDBC controllers. Though only 10 months have passed since its initial beta release, C-JDBC has already been downloaded more than 15,000 times.

***Wayback: A User-Level Versioning File System for Linux***

*Brian Cornell, Peter A. Dinda, and Fabian E. Bustamente, Northwestern University*

■ ***Awarded Best Paper***

Modern file systems and operating systems, though very tolerant of naughty applications thanks to caching, journaling, prefetching, locking, and a whole pile of other nuances, are still rather unhelpful to the common and simple user errors of accidental file deletion and overwrites. Some efforts have been undertaken to provide generalized undelete operations, and code management repositories provide versioning, but for the most part these sorts of mistakes must be protected against at the application level, if at all.

Wayback provides a versioning history through a hidden undo log for any directory remounted with it. It depends on the underlying file system to provide the storage, hence alleviating complexities arising from partitioning and supporting the disparate menagerie of extant file systems. The undo log is precisely that—it records data, allowing it to return to the previous state of a file. Any write operation causes a new entry to be placed in the log. File-system calls are observed by using the FUSE proxy, greatly facilitating development but unfortunately hampering speed slightly. Even so, Wayback is quite quick, usually not too far regressed from the file system inside FUSE.

Future plans potentially include compression, redo logs to provide bi-directional revision traversal, and even hierarchical version storage. The system is absolutely transparent and provides a simple usage paradigm to return to old versions. Since all changes are stored, a foggy memory and a dim spark of energy are enough to recover practically anything.

## SECURITY SIG

*Summarized by Ming Chow*

***Debate: Is an Operating System Monoculture a Threat to Security?***

*Dan Geer, Verdasys, Inc.; Scott Charney, Microsoft*

*Moderated by Avi Rubin, Johns Hopkins University*

In one of the most anticipated events of the conference, Dan Geer debated Scott Charney on whether an operating system monoculture is a threat to society. Dan Geer—an instrumental contributor in the MIT Athena Project, former CTO at @Stake, Inc., and the former president of USENIX—argued that the problem is avoidable and mitigable, but difficult. He compared the current situation to the natural world and biology, and how a diverse gene set mitigates predation and disease. In a computing context, a diverse series of operating systems mitigates the onslaught of attacks and security breaches. Dr. Geer said that "if there is a monoculture, then the bigger the species, the juicier and more attractive." He was critical of the current state of computing and Microsoft's dominance: There are too many gadgets in Windows, leading to confusion as the operating system goes beyond its threshold. In addition, he stated that there are more serious security problems that are not publicly known, and virus writers are two steps ahead of antivirus writers. Dr. Geer concluded his arguments by reflecting on history, that "lessons learned in the real world apply to the computing world: All monocultures live on borrowed time like cotton and potatoes, and we are subject to the laws of nature."

Scott Charney spent years working in the public sector, most notably combating cybercrimes as chief of the Computer Crime and Intellectual Property Section (CCIPS) in the Criminal Division of the US Department of Justice. Mr. Charney stated that diversity of operating systems is not the issue, nor does it matter; the real issue is "how to catch the bad guys." He explained that there really isn't a monoculture of operating systems—not all Windows systems are alike. In fact, Mr. Charney argued that monocultures may be beneficial. He gave the example of Southwest Airlines and how all the planes are Boeing 737s. Southwest Airlines has accepted the risks that all their planes are Boeing 737s, and the major benefit is low-cost maintenance (e.g., pilots can operate any plane without having to learn each one individually). Mr. Charney explained that a computer hacker's goals are to compromise confidentiality and the integrity and configuration of accounts and systems. Finally, Mr. Charney compared the digital age to the Industrial Revolution: makers do not want security because the benefits of technology outweigh security. He believes that it is unfortunate that our current situation reflects that of the Industrial Revolution, and we need to look at security holistically.

During the Q&A session, Dr. Geer was asked about recommendations regarding the fact that there are only a small handful of operating systems available. He responded by saying that "we do not have enough alternatives" and "standards that matter must be platform independent." Mr. Charney was asked about the insecurity of Microsoft Internet Explorer. He responded that Microsoft is releasing better APIs for its products as part of its antitrust settlement. He added that "the issue [security] is in large applications."

Both experts gave their closing remarks after all questions. Mr. Charney concluded by saying that "we have dug ourselves into a deep hole and we need to understand computer security issues holistically to dig out of the hole." Dr. Geer concluded by ascribing the public sickness to bowing to one operating system, and pointed out

that virus writes attack one culture. He reminded the audience of nature's lessons, and that those with the most to lose are those who are the most interdependent.

## PLENARY SESSION

*Summarized by Patty Jablonski and Todd Deshane*

### *Cheap Hardware + Fault Tolerance = Web Site*

*Rob Pike, Google, Inc.*

As we were waiting for the presentation to begin, a continuous stream of words and phases in many different languages scrolled down the open terminal window that showed on the projector screen. This stream of words and phrases contained people's names, Web sites, and places. Rob Pike opened the presentation by saying that this is an example of unfiltered search queries that people submit to the Google search engine and that this is what Google looks like "from Google's perspective."

To demonstrate Google's global presence, Rob showed a map of the world that depicted queries/square degree/hour for a selected day. Spots of light represented the Google queries received from a given area of the world at that time. He noted that there was a place in Tokyo, Japan, that never goes dark. He chose this particular day of data because on that day, back in August, the northeastern United States had experienced a power outage. This area was seen as a dark area on the map that would otherwise have been lighted. He jokingly said, "Where there's electricity, there are Google users."

So, what is Google.com made out of? The search engine consists of a crawler, an indexer and a query handler. The crawler collects documents and makes "copies of the Internet," which now contains over 4 billion Web documents and over 880 million images. The indexer processes and represents the data, and the query handler processes user queries. These three components require Web servers to take the queries, index servers to store the names of the documents, document servers to store all of the Web documents, and ad servers to determine what advertisements to show based on auction money.

There are four main sources of failure that could occur in this type of Web server system. The hardware (e.g., disks), software, the network, or power could fail at any given time. Google deals with each of these types of failures through redundancy and replication.

At Google they expect hardware, especially disks, to fail on a daily basis. Rob Pike stated that the there is a "mean time to failure of three years for one machine, so for 1,000 computers, expect to lose one per day!" So, at Google, they expect to lose many more than one computer per day. With such a high loss rate and the need of multiple computers for storing the 10s of terabytes of the Internet, Google needs a lot of machines and disk space. It makes sense that to save money on the large number of computer purchases, Google chose to buy relatively cheap hardware. In the beginning of Google, when it was located at google.stanford.edu, it consisted of a few cheap PCs in a Stanford University computer lab. Shortly after, it moved to someone's garage, until it finally reached its current location in more sophisticated computer racks contained within multiple data centers across the world. Google's success story is unlike many others during the dot-com boom, during which time many startup companies spent all of their money on expensive hardware with "gleaming racks," while Google held their entire systems together with Velcro!

The PCs at Google are unreliable, cheap, and fast. In order to make these computers reliable, Google must use fault-tolerant software. Reliability, scalability, and load balancing are achieved by breaking resources into pieces and replicating everything. The software is aware of the redundancy structure and spreads things around to avoid a single point of failure. Google's PageRank system (a ranking that is based on the number of links to a given Web page) is used for ordering document names in the index. The index gets broken into "shards" based on its PageRank score. Higher-ranked pages will be replicated more than lower-ranked pages, so that the higher-ranked pages are less likely to be lost in a time of failure. Replication is done at the index level, the document level, and across data centers. When a query is sent, DNS resolves to 1 of *n* data centers (presumably the one closest to the sender). The load balancer then chooses one of the Web servers, which then sends the query to one replica of each index shard. Since the index is read-only and replicated across index servers, the search is done in parallel. This entire process takes approximately a quarter of a second. Google's underlying file system is abbreviated GFS. This file system is large and distributed and contains chunks of files on chunk servers (there are multiple chunks per chunk server). The chunks are replicated, often three times but more often if they are heavily used. The chunk servers act as a master and support automated fail-over. (For more information on GFS, see the GoogleFS paper in SOSP '03.) As seen here, reliable software is more important than expensive hardware in Google's case.

As mentioned, network and power outages can also occur. Rob noted that, unlike their reliance on cheap hardware, Google cannot operate effectively with cheap networking equipment, such as switches and routers. It is important to note that network or power failures only reduce capacity (the query request will merely time out and just needs to be reissued). When an entire

rack of servers was accidentally wiped clean, no data was actually lost—just the terabytes of storage capacity. The same automated software-recovery mechanisms are in place if there is an unreachable network or a power failure.

Throughout Google's existence, they have learned many valuable lessons. The best lesson is "failures will happen, plan for it and survive." When things break, they break too often for humans to fix, so there needs to be an automated, "self-healing" system in place. Another important lesson when dealing with commodity (cheap) hardware and components is that you need to use better software and be careful not to cut corners too much (i.e., two machines per power supply may seem good on paper, but "it is a false economy," since system restoration requires you to power down both computers when only one needs to be off). It is also necessary to make sure that there is adequate cooling for all of the computers. And, finally, Google needs to continue to improve by adapting their fault-tolerant techniques, software, and algorithms to newer and faster hardware. They continue to look for new architectures for redundancy, improve automated failure and recovery, and develop their new services (e.g., Gmail) around these principles.

## FREENIX SESSION: FREE DESKTOP

*Summarized by Brian Cornell*

### *Glitz: Hardware Accelerated Image Compositing Using OpenGL*

*Peter Nilsson and David Reveman, Umeå University*

Desktop computer users have many more demands today than they used to. Users expect features such as translucency, shadows, and transformations to be prevalent. Hardware manufactures have met these demands and provide graphics processors capable of performing these tasks at high speeds. Peter Nilsson and David Reveman introduced Glitz, which brings this hardware power to the user by providing an interface between OpenGL and the graphics library Cairo, where it can be used easily by developers.

The goals of Glitz's design are to create a system with efficiency, quality, and consistency. Glitz implements hardware features, including native off-screen drawing, image transformation, polygon rendering, clipping, gradients, and convolution filtering. Glitz also allows for seamless integration between 2D and 3D environments.

Both accuracy and performance in Glitz were compared against other rendering engines. Glitz was found to operate anywhere from 3 to 200 times faster than the XRender extension to X servers. Using Glitz, Peter and David have shown that Cairo can be very fast. Glitz is at http://glitz.freedesktop.org.

### *High Performance X Servers in the Kdrive Architecture*

*Eric Anholt, LinuxFund*

Eric began by introducing Kdrive, a small X server written by Keith Packard. Since Kdrive is a smaller server, it is easier to change. Eric then went into some background about the capabilities of modern hardware and of current extensions to the X server. Finally, he introduced the Kdrive Acceleration Architecture (KAA), an acceleration extension to the Kdrive X server.

The KAA offers a few improvements, including compositing, blending, and an improved off-screen memory manager. Whereas other implementations are limited to off-screen memory of the same color depth and image width as the screen, the memory manager in KAA allows off-screen memory to be of any type and size. This memory manager also determines what should be kept in which type of memory using a system that keeps track of scores based on the use of buffers.

In his initial implementation, Eric was able to render anti-aliased text five times faster than previously. However, text using composite alpha for subpixel anti-aliasing was five times slower. Eric would like to implement X video, support for GLX and a fix for the composite alpha speeds in future versions. His work can be found at http://pdx.freedesktop.org/~anholt/freenix2004.

### *How Xlib Is Implemented (and What We're Doing About It)*

*Jamey Sharp, Portland State University*

Jamey began his presentation by explaining how Xlib works. He introduced an abstraction between three layers in Xlib: transport, protocol, and utilities. The transport layer handles communication between the client and server, the protocol layer constructs requests for the server, and the utilities layer does everything else. The transport and protocol layers, though very important, are not a big part of the Xlib implementation. Xlib was designed long ago for the systems that were available then and has been added onto many times since, creating a system that is not well designed.

Jamey introduced a solution to this: XCB. XCB is an implementation of an X client library that is simpler and smaller than Xlib. XCB focuses mainly on the transport and protocol layers. The problem with XCB is that most X programs are written to use Xlib, and rewriting them to use XCB would be a major task. Thus Jamey needed a migration path to make this process easier.

Jamey first tried to implement this migration by reimplementing the Xlib API using XCB. The problem with that is that the Xlib API is enormous. After making little progress into the immense repository of Xlib, Jamey started over, this time going from bottom up. He began with a complete Xlib imple-

mentation, and gradually replaced the more crucial components with XCB equivalents. The result is that current Xlib programs can easily by migrated to use these smaller client libraries, and there is no noticeable change in performance. Jamey's work can be found at http://xcb.freedesktop.org.

## USELINUX SIG

*Summarized by Patty Jablonski and Todd Deshane*

### *The FlightGear Flight Simulator*

*Alexander R. Perry, P.A. Murray*

The FlightGear Flight Simulator is a graphics project that simulates flying aircraft in reality. Perry says, "It is not a game." This is an open source project that has been released under the GNU General Public License (GPL) for Mac, Win32, IRIX, and Linux 32 and 64 bit. The simulator is portable, modular, platform neutral, and uses advanced algorithms: "It uses models, not just guesses."

The FlightGear project was started in 1996 by David Murr. Today its worldwide developer community, which includes Perry, consists of 89 people and is still growing. This group is inclusive (goes beyond just software engineers) and is multi-disciplinary (includes both technical and nontechnical people). Perry states that beginners are welcome to join in the project's development as well. For more information on the version releases of FlightGear, visit http://www.flightgear.org/version.html.

FlightGear has many features that make the simulation as realistic as possible. The flight simulator has 3D aircraft and scenery as seen from the pilot's perspective in the cockpit. The lighting changes are realistic and the aircraft is complex. Various things affect the overall flight experience, such as an open window on the aircraft, weather conditions (clouds and smog), and temperature. Intentional impairments have been employed selectively, which makes the simulator more difficult to use but accurately depicts behavior and views. Real-life problems and subtle interactions that can distract or confuse the pilot have been accurately modeled. Many of these features, including instrumentation problems and forces of nature, are often reported as bugs, since people who use Microsoft Flight Simulator are not as familiar with the situations that real pilots face as seen in FlightGear.

The FlightGear implementation is modular and quite complex: "It takes a lot of code to make things behave badly." FlightGear uses networking for remote access and allows a flight instructor to adjust the pilot's settings without the pilot knowing. It uses XML to allow changes to the flight environment. Additionally, there is a property database that stores all of the scenery for the entire planet. A large amount of storage space is needed for this. The 3D graphics and audio are made with OpenGC and OpenAL. Third-party extensions to the flight simulator are generally done in Python.

To download the latest version of FlightGear, see http://www.flightgear.org. Related projects and links:FlightGear Aviation Training Device: http://fgatd.sourceforge.net
OpenAL: http://www.openal.org
OpenGC: http://www.opengc.org
PLIB: http://plib.sourceforge.net

### *Making RCU Safe for Deep Sub-Millisecond Response Real-Time Applications*

*Dipankar Sarma and Paul E. McKenney, IBM*

RCU, or Read-Copy Update, is a reader-writer synchronization mechanism for the 2.6 Linux kernel. RCU is best for "read-mostly" data structures. Although this works well for most situations, real-time applications in the Linux environment are becoming more popular and are in need of quicker response times. RCU callbacks at the end of grace periods (between context switches) cause too much latency in these real-time applications.

First, it is important to distinguish between readers and writers. Readers can access old versions of files independently of subsequent writers. In this case, garbage collection is needed to remove old or invalid copies of the files. Writers, on the other hand, create new files and delete old ones atomically. Because of this, readers have little to no overhead, while writers have a substantial amount of overhead.

Real-time latencies are 800 microseconds measured under load. Measurements were taken with Andrew Morton's "amlat" tool. This 800 microsecond latency is too long for such applications as engine controls, where there is a need to have three degrees of control when measuring revolutions per minute (rpm). There are three ways in which Sarma and McKenney are trying to solve this latency problem.

One possible solution for the latency problem described here is "Per-CPU Daemon." The primary advantage of this approach is that it is transparent to users of "call_rcu( )." Disadvantages include proliferation of kernel daemons and tuning parameters.

Another potential solution proposed by Sarma and McKenney is called "Direct Invocation of RCU Callbacks." Advantages to this option are that there are no kernel daemons and no tuning parameters, and it eliminates "softirqs" for callbacks. Disadvantages are that it is not transparent to the user and it can cause problems if used incorrectly by the user.

Finally, the third option presented by Sarma and McKenney on dealing with real-time latency is "Throttling of RCU Callbacks." The main advantage of this option is that, like "Per-CPU Daemon,"

it is transparent to users of "call_rcu( )." Disadvantages of this method are tuning parameters and that the current implementation is based on iterations and not time.

The initial performance results show that all three approaches have similar, significant performance increases and little complexity. Sarma and McKenney conclude with their argument that RCU can be made safe for real time. Finally, they found that up to this point, "Throttling of RCU Callbacks" is the less intrusive of the two transparent choices. They note that other performance issues exist for real-time applications and that tools to identify problems other than latency are currently under development.

***Making Hardware Just Work***

*Robert Love, Ximian*

The problem with the Linux desktop is that Linux lags behind its competition in terms of hardware management. Instead of having to su to root and mount drives for a simple plug-and-play device, we want our desktop to be able to figure out what to do for us. "Think MacOS X simplicity," Robert Love says.

Robert Love is one of the main developers of "Project Utopia," an umbrella project focused on using the 2.6 Linux kernel's new features, sysfs and hotplug, along with HAL (hardware abstraction layer), udev (a user-level device management system), D-BUS (a message bus), and the GNOME Volume Manager. The goal of this project is to be clean and elegant without the use of kernel hacks, like "supermount." It is very important to "do things right."

The 2.6 Linux kernel does not provide central management as is. It also does not have a platform-agnostic daemon to take advantage of sysfs' device database and hotplug's ability to provide notification of when a new device is added. The additional components of Project Utopia are needed to provide this support.

HAL is a central repository of device information. HAL has a platform-independent interface and persistent key/value pairs, provides asynchronous notifications of changes to devices, and handles Universal Resource Identifiers (URIs) to access devices. The application programmer's interface (API) to HAL is "libhal." Using HAL instead of legacy code reduces the amount of code needed to handle devices significantly (thousands of lines of photo application code can be reduced to less than 10 lines of code with HAL).

On most current Linux distributions, the contents of /dev contains about 18,000 device nodes. Realistically, you only want to see a list of the devices that you currently have. The clean, elegant user-space solution to this is called "udev." /udev only lists devices that you actually have on your system and can be renamed for your convenience.

In order to tie all of this together, Project Utopia needed a message-passing system called D-BUS. The kernel can send out D-BUS messages of new devices. The kernel/D-BUS layer is used by the GNOME Volume Manager.

The GNOME Volume Manager, a manager of disk and other media volumes, allows you to automatically manage volumes, automounting or autoplaying new media/devices, like automatically playing a CD or DVD. It can also create desktop icons based on the type of device or media attached to your system.

Linux is made up of a lot of projects, which often lack integration. The Utopia Project developers hope to bring some unification and integration to the Linux desktop.

## FREENIX SESSION: SECURITY

*Summarized by Matt Salter*

***Design and Implementation of Netdude, a Framework for Packet Trace Manipulation***

*Christian Kreibich, University of Cambridge, UK*

■ *Awarded Best Student Paper*

Solving a problem that involves manipulating network traffic often requires complex filtering, fine-grained and large-scale editing, and visualization. Finding well-maintained tools with the desired functionality is often a hassle and not always possible. Another approach is to write your own solution. While tools that allow editing of captured network traffic have been created, they are often not reusable at the API level, since their functionality is only available in standalone executables.

The network dump data displayer and editor, Netdude, is a framework for packet inspection and manipulation. Netdude has GUI and command line usage paradigms. It allows for scaling of trace sizes and is reusable at all levels, as well as being extensible.

A bottom-up view of Netdude's layered architecture is as follows: libpcap handles elementary trace file operations. libpcacpnav is a wrapper around libpcap that allows one to jump to arbitrary points in the trace file, identified by timestamps or offsets. libpcacpnav uses heuristics to get in sync with the packet stream. Above libpcacpnav is libnetdude, the core of the framework which makes the editing of large traces transparent. libnetdude is extensible through two kinds of plugins: feature and protocol. It provides per-packet TCP dump output, as well as an observer/observee API to inform the user of updates. The GUI is GTK-based and extensible through the same kinds of plugins as libnetdude, and

updates itself through libnetdude's observer API.

Handling of big trace files always involves limiting the number of packets in memory. Since it is not possible to simply use mmap() for inserting and deleting packets, trace files are edited at the granularity of trace areas, which are bounded by timestamps or fractional offsets. Modified trace areas become trace parts, which are flattened onto the original trace file when the file is saved.

Netdude has served as a mechanism to conveniently access suspicious network activity, create traces for network performance evaluation, edit honeypot traffic, and generate IDS signatures. There is much work left to do, including packet resizing and support of scripting environments. Help is welcome!

### *Trusted Path Execution for the Linux 2.6 Kernel as a Linux Security Module*

*Niki A. Rahimi, IBM*

Trusted Path Execution (TPE) was originally a kernel patch to OpenBSD 2.4 created by Mike Schiffman. It was later modified for OpenBSD 2.8 and 2.9 by the Stephanie project.

Currently, TPE is implemented for Linux 2.5/2.6 as a Linux Security Module (LSM). TPE's notion of a trusted path is not to be confused with the more common concept of trusted path in a network context. In TPE, a trusted path is root owned and neither group nor world writable. A trusted user, root by default, is any user on the access control list (ACL) determined by the system administrator.

The TPE LSM performs a check upon execution of a file by utilizing the `tpe_bprm_set_security` hook in the LSM framework. Upon execution of a file, the module verifies whether the user and the path are trusted. The TPE ACL is modified via a sysfs pseudo file-system approach. A directory called tpefs is created containing two files, add and del. Users are added and deleted by writing their UID to the add and del files, respectively. TPE enhances security by preventing execution of untrusted code on the system. The check of path and user occurs exactly before execution is allowed, and if the user and the path are both untrusted, execution is denied.

In addition to trusting code in root owned directories, TPE LSM trusts code in directories of trusted users. TPE is part of the LSM patch as of 2.5.70. It is open to improvements as it is released under a dual BSD/GPL license. LSM, accepted as the current method to introduce security to the Linux kernel by the kernel community, is a small project with lots of potential, for which many more modules are needed.

### *Modular Construction of DTE Policies*

*Serge E. Hallyn, IBM Linux Technology Center; Phil Kearns, College of William and Mary*

Domain Type Enforcement (DTE) is a mandatory access control system introduced in the 1970s by Honeywell, TIS. It assigns types to files and domains to processes. A domain is structured as a list of sets. One of these is the entry type set, which specifies through which types the domain may be entered. Another is the type access set, which specifies which types the domain can access. The signal access set specifies which domains the domain can signal, and the transitions set specifies which domains the domain can transition to. There are two types of domain transitions, auto and exec. When a process under some domain executes a file which is an entry point to another domain, either it must switch to the new domain (auto) or exercise the default option of keeping its domain (exec).

A DTE policy contains lots of domains, types, and defaults. The policy module files presented in this paper are a collection of domains, types, and group definitions, as well as the access rules pertaining to them. Type definitions consist of both path assignments and access grants. Both type and domain definitions may contain assert statements that are used for maintenance of policy constraints, which are interpreted and enforced by a policy consistency class. Since domains and types can declare conflicting access rules, priorities for the access rules are defined. These priorities are determined by placing specific rules over general rules, inbound access rules over outbound access rules, and use of the "absolute" keyword. Group definitions facilitate more generic modules. These are achieved through either the keyword "all," binding a group name to a set of domains or types, or namespace globbing. Groups are expanded only at time of reference and can be dynamically extended.

Modules interact with the system by obtaining system-specific data. They can also be moved between systems and shipped with software. Modules are loaded into the DTE LSM module through a configuration file which is generated by a script that takes a list of module files as input. Work related to this project includes DTEX by Chuck Fox, Fedora, the IBM research project Goyko, and Tresys. This project still needs to be applied to SELinux, which would require object classes and fine-grained permissions, and the modules need to be distributed. Future work also includes possible improvement of the priority specification.

## USELINUX SIG

*Summarized by Martin Michlmayr*

### *Building and Maintaining an International Volunteer Linux Community*

*Jenn Vesperman, author and consultant; Val Henson, Sun Microsystems*

Val Henson shared her insights about creating a volunteer community based on her experience with the LinuxChix project. LinuxChix is an international community whose focus is to create a friendly, predominantly female Linux community. The project was founded in 1999 by Deb Richardson, who created a Web site, mailing lists, and a logo. By 2001, Deb was burning out as the LinuxChix project struggled to remain active. At that time, Jenn Vesperman was chosen as the new coordinator.

Henson summarized the lessons she learned about building and running an international volunteer community, grouping them into three categories: First, the social category, where she suggests that you have to build a sense of community; second, there are organizational aspects, which boils down to delegation; the third category is the technical one, and Henson emphasized the importance of using technology that distributes well.

Henson suggested that building a sense of community was fairly easy for LinuxChix, because building community was the goal of the project. Women are quite excited to find other women who share their interests, and they create bonds fairly easily. One interesting observation Henson made was that being friendly and nice does not attract incompetent people. While LinuxChix' explicit goal is being friendly to its members, there is some hostility in other projects to new members, possibly in order to have a high barrier to joining in order to keep incompetent contributors away. According to Henson, this strategy does not work; being friendly is the better approach. To help create community, LinuxChix uses specialized mailing lists, some of them focused and technical but others allowing completely off-topic discussions. A member-only posting policy on some lists increases the sense of community.

Henson emphasized the importance of delegation as part of organization. She suggested that the first coordinator of the project burned out because she took on too many tasks herself; Henson therefore jokingly said that the number one rule is to do nothing yourself. Instead, delegate to other people; once a task has been delegated, let go and don't interfere. It is also important to give credit. One task of the main coordinator is to monitor the health of other volunteers, and to act accordingly—for example, by sending an overworked volunteer on vacation and by finding more volunteers to help out. Finally, she also suggested that rules should be kept to a minimum—rules will drive good people away, and trolls won't abide by them anyway.

### *Indexing Arbitrary Data with SWISH-E*

*Josh Rabinowitz, SkateboardDirectory.com*

Josh Rabinowitz introduced SWISH-E, a simple Web-indexing system for humans. He summarized the tool as being a fast, powerful, flexible, free, and easy-to-use system for indexing collections of Web pages, and suggested that the definition would not be complete if any of these adjectives were removed. SWISH-E is based on Kevin Hughes' SWISH project from 1994 and is now maintained by Bill Moseley. The tool is written in C and creates binary indexes. There are C, Perl, and PHP interfaces.

There are several alternatives to SWISH-E, such as htdig and MySQL, but Rabinowitz argued that SWISH-E has several attractive features. It is fast and robust, and the tool undergoes steady development and bug fixes. There is good and extensive documentation, as well as a lively and informative mailing list. Finally, SWISH-E offers a "bulk insert" method which doesn't presume to know how or what you want to index. SWISH-E currently handles XML, HTML, and text, but there are two methods for dealing with arbitrary files. First, an external program can be written that converts a file to a format SWISH-E understands. Second, a FileFilter for each given file type can be created. This approach is more modular, but it's slower, since it invokes a child process for each file.

Rabinowitz showed in some examples the ease of creating indexes with SWISH-E and introduced SMAN (http://www.joshr.com/src/sman/), a tool to search UNIX man pages that is based on SWISH-E. Finally, Rabinowitz summarized some future development. The 2GB size limit should be removed soon, and UTF-8 support is a major feature that is being worked on. The ranking system should be rewritten, and the main developer of SWISH-E is interested in working together with a graduate student who would like to pursue such a project.

## FREENIX SESSION

### *Demonstration: Croquet, a Networked Collaborative 3D Immersive Environment*

*Dave Reed, HP Labs*

This demonstration of a 3D networked, collaborative Croquet environment is considered a "research project, not a product," says Dave Reed, project developer. It demonstrates a peer-to-peer networked application that supports collaborative computing and scalable computation.

The demonstration used two PCs networked together on the same switch (the network is meant to

provide different views of the same world), but Reed claimed that the environment can support many more peers. There were some networking problems during the demonstration which made one of the views update too slowly or inaccurately show the world. Because of this, Reed presented the environment in one perspective only.

The environment showed clouds on a 3D plane with windows or "portals" on it. These portals acted as hyperlinks or mirrors to other worlds. Dave went from portal to portal showing us what each had inside. One portal showed a recursive pyramid and another had water that rippled when the mouse was moved over it. Still other portals showed images of people, a chess board, and a flag with a spring and mast to show how the flag changes when moved with the mouse.

He entered into one portal that brought us to an underwater world. He explained that with this scenario, you can "change the laws of physics" by making heavy objects float or have whispers only be heard by certain people across the room. In this underwater world, Reed was represented as a fish. He used a Paint program to draw a new character on the fly and then added this character to the world. The funny part was that there were large help signs in the water explaining "how to make a fish" or "how to make your fish swim" for the beginners. He then left this world and showed a vast world with waterfalls and trees in it that he called a "traditional video game world."

This 3D Croquet environment was implemented using OpenAL and OpenGL for its audiovisual features. It works with communicating objects whose messages are replicated or "cloned"; most messages do not go over the network (almost all computation is local). Objects are governed by mouse movements and positions so that the appropriate action is taken. This system is considered real time, and it is therefore very important that each machine in the network has the correct synchronized time.

In the future, Reed would like to have this Croquet environment implemented for small devices, such as cell phones. He also hopes to develop a security model for the system. There is an important unanswered question: What should people be allowed/not allowed to do in this environment (what is considered cheating, what should be allowed to be read/written, how much should you be allowed to see?)?

This Croquet project is approximately nine months old and is being developed in partnership with the University of Minnesota and the University of Wisconsin. The project is scheduled to be released as open source in an open and public forum.

## USELINUX SIG

*Summarized by Adam S. Moskovitz*

### *Linux and Genomics: The Two Revolutions*

*Martin Krzywinski and Yaron Butterfield, Genome Sciences Centre*

The session started off with Martin Krzywinski, from Canada's Michael Smith Genome Sciences Centre, talking about Linux and genomics, their near-parallel rapid advancements, how Linux is used in genomics, and how genomics has independently adopted many of the same "ideals" as the Linux community.

Martin started by discussing what I believe are the most significant parallels between the Linux and genomics communities, namely, openness and innovation. Just as the Linux community encourages people to build useful things and give them away, the genomics community does that not only with software tools but with the data itself. The most well-known example of this is the human genome project, where the completed genome was uploaded to the public databases pretty much as soon as it was ready. Most public genomic sequencing centers submit new data to Genbank (a public repository of sequence information) pretty much as soon as it is collected.

The genomics community, like the Linux community, actively contributes software to the public on a regular basis. A genome browser from USCS and the Ensembl browser/data miner/visualizer are both freely available. Jim Kent's genomic assembler, which was instrumental in the public effort to complete the human genome, is also freely available. Finally, Lincoln Stein has contributed numerous CPAN modules, both for genomics and such commonly used modules as the Perl interface to Tom Boutell's libgd and Stein's own CGI.pm.

Using these and other public tools, the Smith Centre was the first to publicly release the full sequence of the coronavirus believed to be responsible for SARS (Sudden Acute Respiratory Syndrome). This was accomplished in just five days, using an eight-way Linux system. Krzywinski shared some of the feedback the center received; much of it was positive but some wasn't. One person wrote:

Subject: You have to be NUTS!

My daughter doesn't think its such a good idea to have the gene sequencing for the new coronavirus on the internet. I don't either! There should have been a better way! You must be crazy!

[Summarizer's note: I suppose some people feel the same way about making the Linux source code public.]

By the way, Martin Krzywinski gets my award for most interesting slides: black drawings on a red

background and the funkiest font I've seen in a long time.

### *Thin Client Linux, a Case Presentation of Implementation*

*Martin Echt, Capital Cardiology Associates; Jordan Rosen, Lille Corp.*

In this presentation, Martin and Jordan described how Martin's medical practice decided to install a Linux-based thin client instead of Windows PCs and how that decision has worked out.

Martin started by describing their practice (with more than 200 employees total, 40+ doctors, seven offices, seven hospitals, in New York and western Massachusetts), their work load (128,000 patient visits, 800 open-heart surgeries, 380,000 services billed per year for $22 million in revenue), and their MIS needs (billing, storing test results, financial planning and analysis, payroll, medical imaging, calendar, email, word processing, and more).

Martin then proceeded to give a fairly detailed cost-benefit analysis of "thick" Windows versus thick Linux versus thin Linux. While their initial outlay was about $15,000 more for thin Linux, subsequent savings more than made up for that difference. Specifically, for Martin's practice, their initial outlay per workstation would have been about $3,000 for Windows compared to about $2,100 for thin Linux. Their first-year operating costs showed a similar savings: $2,800 vs. $1,300. With 200 workstations, the savings from choosing a thin Linux client were clear. The last savings was in significantly reduced support costs for remote sites: because almost everything was done at the central office, and because hardware maintenance was mostly reduced to swapping out bad machines (which could be done by an employee with no special skills), their remote maintenance costs were reduced to nearly nothing.

Martin also pointed out several other benefits of thin Linux, chiefly, that employees were more productive. With Windows, too many applications could be customized and employees spent too much time doing this with no real gain in production; with Linux it was easier to disallow such customizations. Another benefit was that applications could be customized to prevent user-caused "outages." The most obvious example was preventing users from closing an application without properly quitting; they simply removed the "X" button from the menu bar! The last of these savings Martin mentioned was preventing employees from using the computers for personal use (things like playing solitaire, downloading music files, and setting fancy screen savers). He estimated that at 15 minutes per day per employee, such wasted time cost his company over $110,000 each year.

Jordan then took over and presented the technical side of things. The first thing he mentioned was that some applications could not be made to work under Linux; for these the practice kept 10 "thick" Windows systems and set up a single Windows 2000 Server system for data storage; Samba was used for logons and drive mappings. Next, Jordan discussed the high and low points of the software used on the thin Linux client: OpenOffice worked quite well, but other applications (Evolution and Mozilla) had a few problems, such as tending to crash or not handling certain required functions (some Web sites, calendaring). There were some problems with low-level things such as file permissions and lack of file locking in OpenOffice.

On the whole, user acceptance of the thin Linux client was high, and the practice has been running for 300 days without a single server crash, the network has never gone down except from human error, the remote desktop (via VPN) has yet to fail, and no viruses or worms have affected their network.

### *Towards Carrier Grade Linux Platforms*

*Ibrahim Haddad, Ericsson Research*

The third talk of this session was what appears to be a refereed paper, written and presented by Ibrahim Haddad from Ericsson Research on "Carrier Grade Linux"—that is, a Linux operating system capable of being used in servers and switches in a public telecommunications network. Typically, such servers require 99.999% reliability (less than five minutes of downtime per year), and switches require 99.9999% (less than 30 seconds downtime). Obviously, no version of Linux is there yet, but Haddad's talk summarized what is needed to get there, as well as what features will be required by carriers before Linux could be used to replace existing, proprietary systems.

Haddad presented an overview of the groups (committees, working groups, associations) working on this problem: The PCI Industrial Computer Manufacturers Group (highly available hardware), the Carrier Grade Linux Working Group of the Open Source Development Labs (Linux improvements), and the Service Availability Forum (defining high-availability APIs). Haddad works with the CGL working group, who released their first public draft in May 2004.

The remainder of Haddad's presentation covered three services not found in the stock Linux release that would be required for mission-critical environments. The first service was TIPC (Transparent Inter-Process Communication), an intra- and inter-cluster protocol that provides a framework for supervising and reporting topology changes. TIPC has been used by Ericsson for several years now and has been available as open source since February 2003. The second service, DigSig (Distributed Digital Signature), is part of the larger Dis-

tributed Security Infrastructure (DSI) initiative. This service allows an administrator to embed digital signatures in ELF binaries and adds functionality to the Linux kernel that prevents unsigned or badly signed binaries from executing. DigSig has been available as open source since January 2003. The last service is a package that implements native support for asynchronous events in the Linux kernel. Carrier grade platforms must process huge numbers of events quickly and efficiently, and this package implements the first tier of such services. It was also released as open source in January 2003.

## PLENARY SESSION

*Summarized by Martin Michlmayr*

### *The State of the Spam*

*Eric Allman, Sendmail, Inc.*

Eric Allman opened his speech in a very funny way when he analyzed the way talks on spam traditionally work. They first summarize what spam is all about, mention that spam is bad, go on to say that the situation is really bad, and finally claim that their product will solve all spam problems. Allman did not go this route, and while he suggested several ways to combat spam, he also made it clear that it will take years to come up with effective solutions and that everyone has to work together.

In the beginning, Allman gave some statistics and summarized claims about spam. Apparently, there are about 90 "world class spammers" who pay US$100,000 per month for bandwidth and servers. According to SpamHaus, 200 spam operations account for 90% of all spam. Spam costs mere microcents per message, which is why spammers continue to operate even though AOL rejects 80% of all incoming mail as spam.

Allman proceeded to summarize existing and new technologies used against spam: realtime blackhole lists (RBLs), content filtering, and challenge-response. RBLs are controversial because their false-positive rate is pretty high. Content filtering comes in different classes: heuristic filters work by observing what spammers are doing and creating means to detect and counter them. Unfortunately, this leaves us in a reactive mode; they send spam, we adapt our tools, and in the meantime we suffer from spam. Fingerprinting and collaboration store a fingerprint of a spam message so other people can test the fingerprint and discard spam. Again, this method is reactive and Allman suggested that it is only effective when the fingerprint database is updated every 15 seconds! Machine learning filters let the computer figure out the interesting stuff. This method needs two piles of "training data": spam and not-spam. While this method works fairly well for individuals, this is less the case on the server level, since legitimate mail varies a lot depending on the user.

Newer methods which are currently being worked on are traffic analysis, identity authentication, and economic schemes. Traffic-based filtering observes typical traffic patterns. For example, a host that normally sends 100 messages in a month and suddenly sends millions in a few minutes is very suspicious. One possible way to use this is to greylist a host and slow down the connection significantly. Identify-based filtering almost always requires authentication. You can use allow-lists and lock-lists in order to reduce the amount of resources spent on more expensive spam checks. There are two philosophies: everything not explicitly illegal gets through (default to accept) or all not explicitly legal gets blocked (default to deny). Sender authentication is not an anti-spam solution in and of itself, but it is essential for identity-based algorithms. We already have SMTP AUTH and TLS, but both are MTA-to-MTA, not end-to-end. Per-user authentication would be possible with PGP or S/MIME. Finally, there are economic schemes which shift costs from recipient to sender. A very small cost doesn't hurt usual senders (perhaps 100/day) but does hurt bulk senders (millions/day). These systems do not necessarily have to be cash-based since the credit can come in a different form.

At the end of the talk, Allman made several predictions. First, he suggested that spam will never go away completely. Authentication will help but won't solve the problem by itself. He thinks that spam will be "manageable" within two to three years, and that legislation will scare away bit players, but not large commercial spammers.

## FREENIX SESSION: SOFTWARE ENGINEERING

*Summarized by Brian Cornell*

### *Managing Volunteer Activity in Free Software Projects*

*Martin Michlmayr, University of Melbourne*

Martin is a member of the Debian GNU/Linux team and brought his experience with free software projects to the community. The main problem he introduced was that volunteers will sometimes neglect their duties, and it is hard to figure out when they do. For small projects this can mean that the project dies because nobody finishes it. For large projects this means that the quality of the product suffers and there are delays in the release of new versions.

Martin went on to describe how Debian is organized and what they do about this problem. At Debian there isn't a hierarchical management structure, so developers aren't supervised by a manager. Therefore they have to carefully look through hundreds of developers to figure out who is neglecting to do what they should. They find these people

in many ways: for example, when there is a bug in a package that is release critical or when a newer version of the software a package provides is available. Every once in a while they compile a list of people who appear to be neglecting packages.

Once you know who is not doing what they should, you have to do something about it. Kicking someone off of a project unnecessarily is not a good idea, because they could provide a lot of help for your project, and they may not be easy to replace. Debian contacts maintainers asking them if they are still active, and gives them two to three weeks to respond. If they don't respond, they're contacted again and given more time to respond before they are eventually removed. Because these people are volunteers, Debian cannot be overly demanding of them, and therefore a polite system like this is necessary.

Martin also points out that you can try to prevent the dereliction problem by having redundancy throughout the project.

### *Creating a Portable Programming Language Using Open Source Software*

*Andreas Bauer, Technische Universität München*

Andreas Bauer's talk would have been very welcome in any class on compilers. He gave detailed information about the use of gcc to create new programming languages. Gcc, the Gnu Compiler Collection, has support for many different languages, and independently has support for many architectures. Andreas presented the capabilities of gcc through a simple expression language he called Toy.

Andreas talked about the current design of gcc's programming language interface. Gcc uses trees to express the language, and then generates an intermediate language called RTL based on these trees. Programming language interfaces are responsible for generating these trees during parsing. Unfortunately, these trees are somewhat tailored toward C and don't make concepts such as tail recursion, garbage collection, and scope representation easy. For this reason, a new system called SSA is being developed with generic trees and more optimization.

## FREENIX INVITED TALK

*Summarized by David Reveman and Peter Nilsson*

### *Current GTK+ Development*

*Mattias Clasen*

GTK+ is a multiplatform toolkit for creating graphical user interfaces with excellent internationalization support. GTK+ was initially developed for and used by the GIMP, the GNU Image Manipulation Program. Today GTK+ is used by a large number of applications and is the toolkit used by the GNU project's GNOME desktop. It can be used with a wide range of programming languages.

Mattias described the different components GTK+ is based on. Glib is a low-level core library that provides data-structure handling for C, portability wrappers, and interfaces for such runtime functionality as an event loop, threads, dynamic loading, and an object system. Pango is a library for layout and rendering of text, with an emphasis on internationalization. The ATK library provides a set of interfaces for accessibility, and the GDK library provides a layer of abstraction that sits between GTK+ and the underlying windowing system.

The talk briefly covered the additions to the 2.4 release, like the new, much improved file chooser and the new combo box. He mentioned that current maintenance of the 2.4 release is mainly directed to bug fixing and performance improvements. It is very complex to fix bugs in such a widely used toolkit without breaking backwards compatibility. Performance improvements have been made primarily in three areas: predictive exposes, reduced flicker by unsetting the background, and reduced signal emission overhead.

Current goals for GTK+ are to provide a full platform, close gaps to higher-level software layers, sanitize the GNOME library stack, keep up with evolving UI needs, and maintain binary compatibility.

The 2.6 release is planned for December 2004 and will contain solidified 2.4 add-ons as well as other smaller additions. The new file chooser will work well in 2.6; improvements include shared settings with Nautilus, automatic shortcuts, recent files, and the ability to choose file formats in the Save dialog. Some missing features will be added to the combo box, including separators, scrolling, and so-called insensitive items. New additions to 2.6 will also be made in areas of command line argument parsing, an icon list widget, a progress cell renderer, and some widgets from libgnomeui. Support for rotated text has been added to Pango, and more work will also be going into Pango.

Mattias talked a lot about what we can expect to see in the future of GTK+. Some of the planned changes are a new rendering model, support for RGBA visuals, an improved theme system, a built-in printing system, and full introspection. A big change that will happen to GTK+ is the introduction of a new rendering model. This will be accomplished by moving to the Cairo library for rendering. Cairo is a modern 2D graphics library with a PostScript-like API. It has capabilities similar to Java 2D, SVG, and PDF 1.4; alpha-compositing is a natural part of Cairo. Cairo has output back ends for X, OpenGL, local image buffers and PostScript. Support for RGBA visuals will be added to GTK+ and will make translucent windows, fade-in effects for menus, and drop shad-

ows well supported. The motivation for a new improved theme system is the desire to remove GTK+ dependencies, fully support Cairo's rendering model, and include layout access in the theme system. A full theme system specification should be made available and will most likely use a standard syntax like XML's CSS. The built-in printing system will include appropriate printing dialogs and will be based on Cairo, with back ends for CUPS, lpr, and GDI. Introspection is useful for language bindings, documentation, and IDEs. GTK+ already supports introspection of type hierarchy, properties, and signals, but not yet of virtual functions in class structs and library functions, which will be added in the future.

## EXTREME LINUX SIG

*Summarized by Matt Salter*

### *A New Distributed Security Model for Linux Clusters*

*Makan Pourzandi, Open Systems Lab, Ericsson Research*

The target applications for distributed security are large distributed applications with a large software base that provide around-the-clock service and require high availability (99–99.999% uptime). The model presented in this paper specifically targets Linux clustered servers and is intended for servers exposed to the public, providing services to different operators, and running untrusted third-party software.

Distributed security has several requirements. One is security isolation, or compartmentalization. This is needed because exploitable vulnerabilities are probable in a large software base; without compartmentalization, a single vulnerability could expose the entire system. Runtime changes to the security context must be possible and reflected immediately, and application-layer security cannot be relied upon, since administrators must then contend with vulnerabilities in applications over which they have no direct control. The challenges of distributed security include creating a coherent implementation that does not leave any security gaps, integrating different security solutions from different vendors, and managing the system to prevent misconfigurations and inconsistencies.

Because most of the target applications have only a few users with whom everything is done, a security policy based on process is needed. At the node level, such security is achieved through mandatory access control. The model presented in this paper extends mandatory access control to the entire cluster. Processes are assigned a unique security ID (ScID), assembled from the ScID of the binary (stored in the ELF header), the ScID of the parent process, and the node security ID (SnID). To achieve compartmentalization, virtual security zones are set up inside the cluster. Security zones are groups of ScIDs and SnIDs. The distributed security policy allows for access control decisions on the process level based on the IDs of the source and target processes. Network, socket, and transition rules also exist. The architecture of the distributed access control implementation is as follows: each cluster has a single security server and each node has a security manager. The security policy is propagated from the security server to the security managers, which enforce policy at the node level via secure communication channels.

This model is not intended to replace existing security solutions, but, rather, to serve as an add-on to them. Challenges include creating a comprehensible and acceptable security policy and explicitly defining security zones in the distributed security policy.

### *Implementing Clusters for High Availability*

*James E.J. Bottomley, SteelEye Technology*

A "highly available" (HA) system is any system that takes action to increase availability beyond what would ordinarily be possible. HA clusters consist of multiple networked local machines with some type of shared storage. There are three types of HA clusters. The simplest type is a two-node-only cluster, which cannot be scaled. A second type is the quorate cluster, which is centrally controlled and will not work without a membership service. A quorate cluster is defined such that no other cluster may be formed from excluded nodes, which means it cannot be split into two clusters. If the cluster is split, the majority of the nodes survive. The final type is the resource-driven cluster, in which resources are grouped by which services they belong to. In a resource-driven cluster, a node must simply establish ownership of a group to export the service. Resource-driven clusters also allow independent subclusters to form. The simplest of these cluster types is two-node-only, followed by resource-driven, and the far more complex quorate cluster. Recovery is much faster in resource-driven clusters than in quorate clusters.

Determining availability is difficult because you need to know what the system's uptime and downtime are in your environment. While duplication of nodes allows you to determine downtime, it does not allow you to determine uptime. Uptime can only be controlled through careful implementation and deployment of the cluster. However, whether availability or downtime is significant depends on the type of service being offered.

Often, it is the application that fails instead of the server. Monitoring applications is important so that application failures can be spotted

and corrected. Local application recovery is important as well, since applications fail more often than nodes, and local recovery decreases downtime and minimizes disruption. Also, monitoring for failures in general is important, since while redundancy protects your system from the first failure, the second failure will take your system down.

Uptime can be improved by assessing cluster hardware and eliminating single points of failure (SPOFs). Clusterwide SPOFs should be eliminated entirely, while individual-node SPOFs should be evaluated to see if eliminating them would improve uptime. In a shared storage cluster, the real SPOF is storage and should be addressed through replication by making sure that the external array is configured as RAID 1. Power supply, mechanical devices, and the connection to the storage are the node SPOFs. One way to eliminate the connection to a storage SPOF is to have multiple connections to the storage from a node. This is called a multipath cluster.

The biggest Linux-specific problem faced by cluster manufacturers with binary modules is simply keeping up with the kernel patches and releases. Another problem is the dreaded "oops," which kills kernel processes and then tries to continue. If the kernel was in a critical section at oops time, the system may hang. Large Block Device support (LBD) is a Linux feature that helps clusters. It is limited to 2TB in the 2.4 kernel. Multipath solutions are different for every vendor in the 2.4 kernel, but an attempt is being made to unify the architecture on Device Mapper for 2.6.

## FREENIX SESSION: SYSTEM BUILDING

*Summarized by Brian Cornell*

### *KDE Kontact: An Application Integration Framework*

*David Faure, Ingo Klöcker, Tobias König, Daniel Molkentin, Zack Rusin, Don Sanders, and Cornelius Schumacher, KDE Project*

Cornelius Schumacher presented Kontact, a Personal Information Manager (PIM) for the K Desktop Environment (KDE). Kontact was designed to integrate individual components such as Kmail, Korganizer, Kaddressbook, Knotes, Knode, and Kpilot. The developers wanted an interface with which all of these programs could be used together, without maintaining the separation between the individual projects.

Kontact was designed with the basic goal of keeping all of the components in it as separate as possible without the user being able to tell. To satisfy this, the components had to be integrated on an application level and still be able to run alone. But to maintain the semblance of integration, the components needed an integrated UI, inter-component communication, and shared settings.

With these constraints in mind, the KDE Kontact team designed Kontact to use plugins from each application with a Kpart for the user interface. The components then communicate using DCOP. Using the Kparts—basically component versions of the applications—Kontact embeds each component into a unified Kontact user interface. The components can also use a unified configuration through Kconfig. Kontact is an ongoing project: http://www.kontact.org.

### *mGTK: An SML Binding of GTK+*

*Ken Friis Larsen and Henning Niss, IT University of Copenhagen, Denmark*

Henning Niss presented mGTK, a binding to the GTK+ graphics toolkit for the Standard ML language. The goal of this project was to provide SML access to a good general-purpose toolkit. Keeping this in mind, the developers wanted a direct binding to the C interface of GTK; they wanted it to work under any SML compiler, and they wanted compile-time type checking. Other interfaces to GTK only give errors at runtime, making it harder to fix bugs and optimize programs.

SML is a functional language with a formal definition. It is separated into two parts: the core language and the module language. There are many implementations of SML, and the mGTK developers targeted two of them, Moscow ML and MLton. They used a system of type constraints, including what are known as phantom types to enforce type checking.

Using mGTK, GTK+ classes are translated to SML signatures. Class types are represented as SML types, and methods are implemented as functions. mGTK can automatically generate the SML binding based on the gtk.defs file that comes with the GTK API. mGTK is available at http://mgtk.sourceforge.net.

### *Xen and the Art of Repeated Research*

*Bryan Clark, Todd Deshane, Eli Dow, Stephen Evanchik, Matthew Finlayson, Jason Herne, and Jeanna Neefe Matthews, Clarkson University*

Repeated research is a process often used to verify results in scientific research. Jeanna, Stephen, and Todd presented an application of this process to the world of computer science research. As a class, they tried to reproduce the tests in the paper "Xen and the Art of Visualization." They wanted to see if they could get the same results, and to apply more tests to Xen in hopes of further examining its performance.

Reproducing the environment in which the original Xen tests were run was not easy. They had to first obtain the same hardware and then install the same software used in the original tests. The Xen authors listed the benchmarks they used, making it easy to reproduce those, but assembling and running them was time-consuming. Also, some of the benchmarks used were closed benchmarks, so they had to be replaced with similar open source alternatives. The result of all of the work was that their repeated measurements were within 5% of the original measurements.

The team applied many other tests to Xen: they tested its usability as a set of virtual Web servers; they tested it on commodity hardware, rather than the server machine that the original tests had been run on; and, finally, they compared the performance of Xen to that of an IBM zServer. They learned from this that repeating research is not easy, but it is an important reality check in the development of new technologies.

## EXTREME LINUX SIG

*Summarized by Bill Bogstad*

### *Scaling Linux to Extremes: Experience with a 512-CPU Shared Memory Linux System*

*Ray Bryant, John Baron, John Hawkes, Arthur Raefsky, and Jack Steiner, Silicon Graphics, Inc.*

Ray Bryant spoke about SGI's Altix Itanium 2-based HPC (high performance computing) servers. Non-shared memory computing clusters are frequently talked about today, but SGI believes that NUMA (non-uniform memory access) shared-memory compute servers remain appropriate for many HPC applications. SGI's Altix systems are architecturally similar to their MIPS-based Origin 3000 servers. The basic building block of an Altix system is a computing brick that has two pairs of Itanium 2 CPUs. Each pair of CPUs can have up to 16GBs of dedicated local memory. Bricks are connected using SGI's NUMAlink technology, which supports cache coherency and uses specialized routers. Altix systems have achieved records on a number of HPC benchmarks, including SPEComp L2001 in June 2004. The underlying interconnect technology supports up to 2048 CPUs, but SGI currently only supports 256 (soon 512) CPUs in a single SSI (shared system image) under Linux.

SGI believes that porting of single CPU applications to an SSI system can be much easier then porting to a computing cluster since non-performance-critical code can be left as non-parallel. When SGI decided to develop a NUMA system using Itanium CPUs, a Linux port to the Itanium was already available and it was decided that it would be easier to start from this port than to move IRIX from MIPS to Itanium. The current goal is that if an application runs on a generic Itanium under RedHat AS 3.0, then it should run on an Altix system.

However, even getting the 2.4 Linux kernel to run well on the Altix hardware has been an interesting challenge. The kernel had to be taught the performance differences between local and remote memory. On a 512-CPU system this is critical, since only 0.4% of total system memory is local (i.e., fast). A new round-robin buffer cache page allocation algorithm is used to avoid having a brick fill up all of its local memory with cached pages, which would leave no local memory in which to run applications. An O(1) scheduler was added with the elimination of a global runqueue lock and a resulting sixfold improvement on some benchmarks. Elimination of system global variables in favor of per-CPU variables and value aggregation as required was needed to support very large systems. Without these changes, the system would spend all of its time pounding the cache coherency hardware just to keep system status variables updated. A number of kernel hash tables are sized at O(1%) of total system memory, which is more memory than exists at any one brick in the system. These tables are now spread out in the same way that the buffer cache is.

Changes were also made to allow system operators effective use of the system. The dplace command allows the operator predictable memory and CPU allocation to the threads in a single process. By specifying the appropriate parameters, performance can be enhanced by taking advantage of knowledge of the memory access patterns of the various threads in a process.

Looking forward, many of SGI's changes have made their way into the 2.6 Linux kernel. As a result, a generic 2.6 kernel will boot on an Altix system. As SGI moves to supporting kernels based on 2.6, they expect improved scalability and the ability to support larger systems.

### *Quantian: A Single-System Image Scientific Cluster Computing Environment*

*Dirk Eddelbuettel, Debian Project*

Quantian is a Linux distribution that is focused on cluster-based scientific computing. It was first released in March 2003 and has gone through a number of major releases since then. The latest releases can no longer fit onto a single CD and now require a bootable DVD or booting from a hard disk. Quantian's lineage can be traced back to the popular Debian distribution. The path is from Debian to Knoppix to cluster-Knoppix to Quantian. From Knoppix, it inherits read-only media-based simplicity and automatic hardware detection, along with support for persistent data on USB storage devices. clusterKnoppix adds zero-configuration OpenMosix clustering with automatic process migration along with the cluster-compatible Mosix File Sys-

tem. A single machine can be booted from Quantian media and then other machines can network-boot via the PXE protocol and form a single Mosix cluster.

Quantian extends clusterKnoppix with a large number of scientific computing applications. In particular, Beowulf-style clustering tools and libraries are included along with the statistical package R and the SNOW extensions. SNOW allows easy access to high-level parallel statistical computing. Some Knoppix packages that are not related to scientific computing or related software development have been dropped in order to make room for Quantian's scientific computing additions.

Currently, Quantian is essentially a one-man operation maintained by Dirk. He responds to requests for the addition of new packages as time and interest allow. Distribution size, network security concerns, and surveying users for their needs and configurations remain open issues for him. Even though it is primarily a repackaging of other components, Quantian deserves a look if you are interested in scientific computing. At the end of his talk, Dirk mentioned that the laptop he was using was running Quantian with a USB flash drive for persistent storage. It seems that his employer will not let him install Linux on the company-supplied laptop, so he has found another way. Let's hope Dirk keeps finding another way.

### *Cluster Computing in a Computer Major in a College of Criminal Justice*

*Boris Bondarenko and Douglas E. Salane, John Jay College of Criminal Justice*

John Jay College is a specialized liberal arts college within the City University of New York system. It offers degrees in Law and Police Science, Fire Science, and Forensic Science among others. So, you might ask, just what kind of cluster computing is needed in a College of Criminal Justice? Douglas Salane made it clear that there are a number of areas where significant computing resources can be helpful.

Current and planned projects include simulations of the fires that occurred after the attack on the World Trade Center, database analysis and data mining of the FBI's National Incident-Based Reporting System, and molecular modeling for toxicology studies.

John Jay College has a relatively small cluster-computing facility. The compute cluster consists of 12 nodes with two CPUs each. A separate database cluster has four nodes, and the computing laboratory has 30 Linux workstations. Still, they had to go through much of the same decision-making processes that larger facilities might go through. Blade/rack systems or piles of PCs? What network file system to use? What interconnect technology? How to manage and monitor the cluster? How to test the correct functioning of the cluster? A cluster-specific Linux distribution or self-configuration?

Verifying the correct functioning of the cluster was of particular concern to Douglas. This concern was strengthened when the test software that is included in the BLACS portion of the ScaLAPACK software library reported incorrect results for some of its tests. In the end, the error was traced to a faulty Gigabit Ethernet card in one of the machines. Other cluster-computing packages don't always provide those kinds of tests. On the other hand, ScaLAPACK can be difficult to use.

For a small site, just figuring out what cluster-computing software is available and how to set it up is a significant undertaking. Unfortunately, Linux distributions, like the previously mentioned Quantian, were not available when they first started working on their cluster. Support for heterogeneous clusters would also help by allowing them to expand the size of their cluster over time without sacrificing performance to the demands of optimizing software to the lowest common denominator.

;login:

PERIODICALS POSTAGE
**PAID**
AT BERKELEY, CALIFORNIA
AND ADDITIONAL OFFICES



USENIX Association
2560 Ninth Street, Suite 215
Berkeley, CA 94710



POSTMASTER
Send Address Changes to *;login:*
2560 Ninth Street, Suite 215
Berkeley, CA 94710

**********CAR-RT LOT**C000