



# MUUGlines

The Manitoba UNIX User Group Newsletter

## May 13, 2003: Samba 3

Samba is a suite of programs that seamlessly provide SMB/CIFS file-and-print services to Windows clients. A robust and reliable implementation, it is now at version 2.2.8. What's in store for the future? In development now, and available for download for testing is version 3.0 alpha.

This version has many enhancements to the 2.x versions, including Active Directory support (this release is able to join a ADS realm as a member server and authenticate users using LDAP/Kerberos), Unicode support, better Windows 2000 printing support including publishing printer attributes in Active Directory, support for migrating from a Windows NT 4.0 domain, support for establishing trust relationships with Windows NT 4.0 domain controllers, and lots of other changes!

Kevin McGregor will show you how to get started in testing (and using!) this leading-edge technology.

Meetings are held at the IBM offices at 400 Ellice Ave. (between Edmonton and Kennedy). When you arrive, you will have to sign in at the reception desk, and then wait for someone to take you (in groups) to the meeting room. Please try to arrive by about 7:15 PM, so the meeting can start promptly at 7:30 PM. Don't be late, or you may not get in.

Limited parking is available for free on the street, or in a lot across Ellice from IBM, for \$1.00 for the evening. Indoor parking is also available nearby, at Portage Place, for \$2.00 for the evening.

## Reflections On The 25th Anniversary Of Spam

*Brad Templeton*

*[Ed. note: Brad is the founder of Clarinet, the world's first "dot com" company, and Chairman of the Board of the Electronic Frontiers Foundation (www.eff.org - join today!)]*

While many only encountered spam (junk e-mail or junk newsgroup postings) in the mid 1990s, my

research has found it goes back much further than that.

In fact, the earliest documented junk e-mailing I've uncovered was sent May 3, 1978 — 25 years ago this Saturday. (It was written May 1 but sent on May 3.) And in a surprising coincidence, just a month ago marked the 10th anniversary of March 31, 1993, the first time a USENET posting got named a spam.

I learned of that first spam through a report from Einar Stefferud [*Ed. note: a relative of mine!*] who read a history I prepared of the term "spam" (<http://www.templetons.com/brad/spamterm.html>) and how the name of the canned ham became our name for junk e-mail. I had original set out to research the history of the term, but it became impossible not to research a bit of the history of the act.

That first spam was sent by a marketer for DEC - Digital Equipment Corporation. Today, you may not know DEC, since it was bought by Compaq and is now a unit of HP, but in those days it was the leading minicomputer maker, and its computers provided the platform for the development of Unix, C and much of the Internet, to cite just a few minor events.

By 1978 the Arpanet (as the Internet was then known) had already provided network E-mail to a large number of folks at universities, government institutions and universities for over 6 years. E-mail was the biggest source of traffic on the Arpanet. A few years prior, Dave Farber had created "MsgGroup," the first network mailing list. (Though Plato and other time-sharing systems had laid the foundations for on-line community and conferencing some years before that.)

The DEC marketer, Gary Thuerk, identified only as "THUERK at DEC-MARLBORO" (There were no dots or dot-coms in those days, and the at-sign was often spelled out) decided to send a notice to everybody on the ARPANET on the west coast. In those days there was a printed directory of everybody on the Arpanet which they used as source for the list. The message trumpeted an open

house to show off new models of the Dec-20 computer, a foray into larger, almost mainframe-sized systems.

This was a spam, though the term would not be used to refer to it for another 15 years. Thuerk had his technical associate, early DEC employee Carl Gartley, send the message from his account after several edits. Alas, at first he didn't do it right. The Tops-20 mail program would only take 320 addresses, so all the other addresses overflowed into the body of the message. When they found that some customers hadn't got it, they re-sent to the rest.

As you can guess there was quite a response, with (as is typical) far more volume of debate than actual spam. It's amusing to see that one future celebrity — a young free software guru Richard Stallman — at first wondered why people were so upset about the message. He later said the mistaken placement of all the addresses into the body did bother him, but he gets the dubious honour of being perhaps the first spam defender. Of course like all of us he was 25 years younger and the problem was brand new.

In those days the Arpanet had an official “acceptable use policy” which limited it to use in support of research and education. So this message was a pretty clear violation, and the DCA, which ran the Arpanet, gave a very stern call to Thuerk's boss about the matter. The policy was well enough known over time that we would not see significant spam for many years to come after that.

See <http://www.templetons.com/brad/spam/spam25.html> for more of Brad's reflections on spam.

## OpenBSD 3.3

We “pre-announced” this last month, but it's worth repeating. As of May 1, OpenBSD 3.3 is current. Todd Miller ([Todd.Miller@courtesan.com](mailto:Todd.Miller@courtesan.com)) writes:

OpenBSD 3.3 has been released and is now available from select ftp mirrors. We believe this is our best release ever with cool new features including ProPolice stack protection in the compiler (<http://www.trl.ibm.com/projects/security/ssp/>), queue support for pf (traffic shaping), an hppa (pa-risc) port, W^X for some platforms (sparc, sparc64, alpha and hppa), privilege separation in the X server and xconsole, fewer setuid binaries, a new

spam deferral daemon and much, much more.

This is our 13th release on CD-ROM (and 14th via FTP). We remain proud of OpenBSD's record of seven years with only a single remote hole in the default install. As in our previous releases, 3.3 provides significant improvements, including new features, in nearly all areas of the system.

See <http://lwn.net/Articles/30858/> for the full announcement, and for even more detail on the architecture-independent changes from 3.2 (such as “smtpd(8) has left the building”), see <http://www.openbsd.org/plus33.html>.

## A Very Brief Introduction to Squid Configuration

*By Kevin McGregor*

I recently set up a PC at work to perform Web caching using Squid ([www.squid-cache.org](http://www.squid-cache.org)), and it was so easy I thought I'd share a few details. I will assume, though, that you know enough to successfully install a common Linux or BSD distribution, and locate and install the Squid package. This is about configuration, after all!

The necessity that urged this was the fact that not all PCs where I work had direct Internet access, which made downloading and installing updates more tedious than I thought it should be. Since I knew of one available IP address which was allowed this access, I thought setting up a PC using that address and running a Web proxy would help me out significantly. And it has!

First, a word about partitioning. This project was limited in scope to Web proxy and caching only, with no other services provided by this system. Therefore, my decisions were based on that assumption. I only had a 6 GB hard drive to play with, so I made a few guesses to work with. I used Red Hat Linux, and used the “server” installation selection. Taking most of the defaults left me with a 100 MB “boot” partition, a modest-sized swap partition to suit the 128 MB of RAM, and the rest was in “/”, the root partition. The set up I went with allowed several gigabytes that could be used for web caching, but I am currently working with a cache of 1 GB.

In my case, the configuration file, `squid.conf`, is found in `/etc/squid`. A default configuration file (`squid.conf`) is provided, with a copy named `squid.conf.default` for reference. This default file, like those of all quality open-source software,

contains many examples and explanations of the settings. Combined with the usual readme files, a basic set up is a snap.

First, use `top` or `cat /proc/meminfo` to find out how much free RAM you have before starting Squid. Squid uses some RAM for caching, and a good rule of thumb (according to the documentation) is 1/3 of available RAM. In my case there was about 75 MB free. So:

```
cache_mem 25 MB
```

Next, set the caching directory. The default location is in `/var/spool/squid`. For more sophisticated environments, you might consider mounting a separate partition as `/var/spool/squid`. The only change you need make to the following line is the first number, which is the total disk space (in KB) to dedicate to caching. As I mentioned above, I'm using approximately 1 GB:

```
cache_dir ufs /var/spool/squid 1000 16 256
```

There are several ways to limit which hosts are allowed to use this proxy; I'm allowing all hosts with IP addresses in the 192.168.7.0 and 192.168.8.0 networks to use it:

```
acl allowed_hosts src 192.168.7.0/255.255.255.0
192.168.8.0/255.255.255.0
```

"Acl" is an abbreviation for "Access Control List". "Allowed\_hosts" is a label you make up. "Src" determines that the source IP address is checked, and the remainder of the line is a space-separated list of networks (network-address/network-mask) or IP addresses to which you want to grant Internet access. The ACL by itself does nothing (despite the very obvious label you gave it). You still have to specify what type of access applies to this ACL:

```
http_access allow allowed_hosts
```

As a safety precaution, they suggest using the following in case you have web applications running on the proxy server:

```
http_access deny to_localhost
```

Lastly, to set a special hostname (other than that returned by `gethostname()`) for use in error messages and the like:

```
visible_hostname ProxyServer
```

That's it! Now enable it with

```
chkconfig squid on
```

and start it with

```
/etc/init.d/squid start
```

If all is well, set your web browser to use `<proxy-server-IP-address>` as the proxy and 3128 as the proxy port, and happy accelerated web browsing!

## Even Fonts Can Have Bugs

Announced recently at [GnomeDesktop.org](http://GnomeDesktop.org): "This is the Bitstream Vera family of fonts, version 1.10, the first version available under a free license. Please discard the beta test fonts.

Note that these fonts have tighter spacing (leading) than the beta test version of Vera, and that 00 and 11 are clearly distinguishable in the mono family, as well as many other fixes and adjustments. The leading adjustments required redesigning the accents, and took extra time to complete."

Get the new versions at <http://www.gnome.org/fonts/>.

## Novell Forge Launched

Having contributed for years to open source community projects such as Apache and OpenLDAP, Novell is today creating a more direct path to its open source initiatives with the launch of the Novell® Forge Web site (<http://forge.novell.com>), an open source developer resource. As part of the announcement of Novell Forge, the company is releasing the source code of the Novell Nsure™ UDDI Server, which makes Web services registries more secure and easier to manage by adding identity management capabilities to the Universal Description, Discovery and Integration (UDDI) standard.

"Novell believes strongly that the use of open source technology must be reciprocated by giving back to the community; and for years we've quietly contributed to a number of open source projects," said Chris Stone, Novell vice chairman - Office of the CEO. "The launch of NetWare® 6.5, which makes significant use of open source components like Apache, Tomcat, MySQL\* and PHP/Perl, is a milestone in Novell's commitment to open source - so it only makes sense that we also increase our commitment to giving back, with the launch of Novell Forge and the release of the first industrial-strength, open source UDDI server."

Through the Novell Forge site, developers will

be able to download, modify, exchange and update open source code released by Novell; share projects and ideas with the broader Novell development community; and participate in vertical market and technology communities. In addition to the Novell Nsure UDDI Server, the site will host Novell's existing open source projects, such as the company's DSML support for eDirectory™ that was recently published through OpenLDAP. Other open source additions to the site will be announced later this year.

"We welcome Novell's participation in open source development," said open source consultant and advocate Bruce Perens. "Recently, Novell contracted with MySQL AB to provide a commercial MySQL license to its users with every NetWare system. Since MySQL uses the same code base for its commercial-licensed and GPL versions, Novell's license payments support the continued development of MySQL under the GPL license.

"Now, Novell is taking the next step, by producing its own free software under OSI-approved licenses, including a UDDI implementation and other facilities," Perens said. "They have taken some care to act as good citizens of the open source developer community."

The Novell Forge Web site and Novell Nsure UDDI Server are available today to open source developers worldwide. For more information, go to <http://forge.novell.com>. Developers can also test a live version of the Novell Nsure UDDI Server by going to <http://uddi.novell.com>.

## Chandler 0.1 Released

The first release of Chandler (the personal information manager system being developed by Mitch Kapor's Open Source Applications Foundation) is now available. It is an early (v0.1) release, aimed more at letting people look at the code than providing a useful application.

It gives me great pleasure to announce the first release of Chandler, release 0.1! While we are still very early in the design and implementation process, we intend for this 0.1 release to make us a more fully open project. We have made the release available for download, opened up our bug tracking database, and opened our source code repository. We have also spent quite a bit of time in the past few weeks focusing on improving our code and documentation.

We are focusing first on architectural issues, not end-user features, yet there are some interesting things to look at in this release. In particular, the outlines of the Chandler peer-to-peer sharing framework are visible. There are also the very crude beginnings of a calendar and contact manager to play with.

We haven't begun to focus on a polished user interface, the current user interface is a placeholder. So rest assured that future releases will bring dramatic improvement in this area as well.

The 0.1 release does have a fair amount of the framework for writing "viewer parcels" — plug-ins that can be added to Chandler extremely easily. In fact, two of our sample viewer parcels were written by OSAF team members who were not hired as programmers and who were completely unfamiliar with the tools.

In future releases, we will improve both the architectural framework and user experience significantly, with a short-term emphasis on novel features over robustness or well-understood features. We need to incorporate our original features as early as possible so that we can explore their implications, and we need to make it as easy as possible to implement those original features.

0.1 Release info: (links to all end user and developer documents) at [http://www.osafoundation.org/Rel\\_0.1\\_Info.htm](http://www.osafoundation.org/Rel_0.1_Info.htm)

## Sending Us E-Mail?

Due to the amount of e-mail MUUG receives, we've set up an auto-reply to give you immediate feedback, and redirect some of the e-mail to the appropriate places. Why not look at <http://www.muug.mb.ca/about.html#contacts> first?

## Share Your Thoughts

E-mail us with your comments on the newsletter, whether it's criticisms or commendations, and continue to send in articles or ideas for same.

If you have a How-To or other idea, and aren't ready to give a presentation at MUUG, an article is a great alternative! If you can write better than the editor, that's terrific; if you can't, submit it anyway and we'll get it into shape for publication. We know that many of you have some great ideas and lots of knowledge.

Why not share? Mailto: [editor@muug.mb.ca](mailto:editor@muug.mb.ca).