MUUGLines

The Manitoba UNIX User Group Newsletter

Volume 28 No. 9, May 2016 Editor: Kevin McGregor

Next Meeting: May 10th, 2016, 7:30pm

Main Topic: Now In 3-D!

In this case we mean that this meeting will feature a Double Daemon Dash (DDD) instead of one DD and one regular presentation:

In Wyatt's Daemon-Dash segment he will be presenting the Squid caching proxy server daemon. Squid is a full featured proxy daemon with support for http, https, and ftp. Wyatt will show the quickest way to have Squid up and running on a new server, and how to keep it secure.

Paul Sierks will present on DavMail.

Where to Find the Meeting

New Room! 1M28 Manitoba Hall

New room! This month's meeting is on the east side of campus in room 1M28, 1st floor, Manitoba Hall, close to Balmoral Street. Look for a sign on the door. Doors are usually open by 7:00 pm with the meeting starting at 7:30 pm. Parking is available on the surrounding streets. Please see `http://www.uwinnipeg.ca/maps` for further information about parking and access to the campus.

The latest meeting details are always at:

`https://www.muug.mb.ca/meetings/`

Money and Stupid

Because if you've got lots of money, why not be stupid with it? Here's a fellow who apparently owns a Tesla (yes, the electric car) and is a fan of Gentoo ("Oh dear God no", I know you're thinking already) so he's managed to install Gentoo on his Tesla.

It's not the whole car, mind you, just the big touch-screen display unit. It's a Gentoo ARM system,

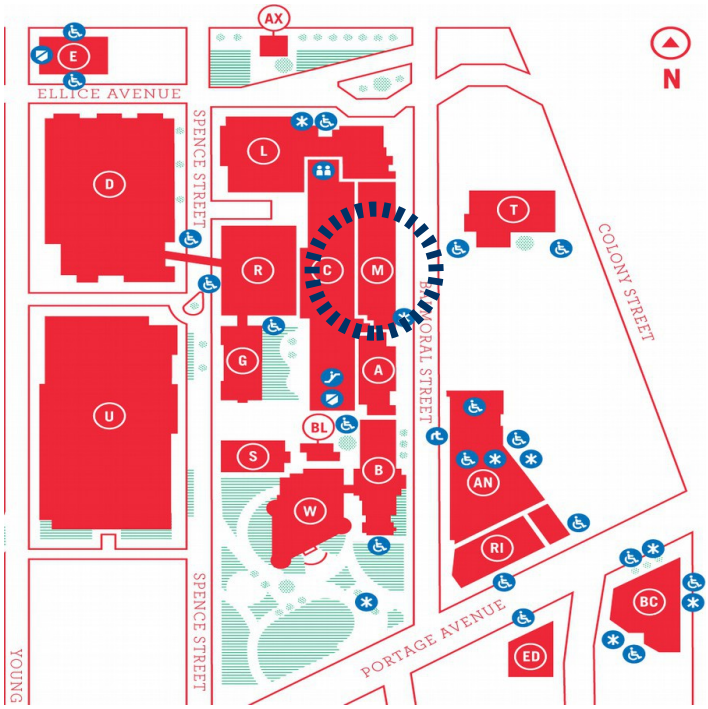

cross-compiled using a qemu-user chroot environment. Gory details at `http://www.sutesla.space/`.

The Web is Doom

Recall that Doom is a multi-level first person shooter that ships with an advanced 3D rendering engine and multiple levels, each comprised of maps, sprites and sound effects. According to a recent analysis, the average web page size is now roughly the same size as the entire game of Doom – 2.25 MB!

`https://mobiforge.com/research-analysis/the-web-is-doom`

Ubuntu 16.04 LTS Xenial Xerus

Canonical has released the 6th LTS release of Ubuntu. Its most notable features are a new app distribution format called "snaps"; the pure-container hypervisor LXD; and the included ZFS-on-Linux, which is sure

to give open-source proponents headaches and lawyers smiles.

Get it at the usual location, http://www.ubuntu.com/download. As a side note, Ubuntu 16.10's code name will be Yakkety Yak.

FreeBSD 10.3 Released

For those of you looking for an OS with ZFS support on somewhat firmer ground, take a look at the latest release of FreeBSD. A few of the latest features are support of the latest Intel Skylake chipsets, Gnome 3.16.2, improved root-on-ZFS support, and the Linux compatibility layer has been substantially improved.

https://www.freebsd.org/

OpenIndiana

For another approach to ZFS, there's always the OpenSolaris fork known as OpenIndiana. Progress continues to be made; the release known as Hipster 2016.04 is out in ISO form. It's generally a rolling release with only occasional installation ISOs, so if you'd like to get your feet wet, now would be a good time.

http://www.openindiana.org/

Nethack's Long-Awaited Update

The recent update to NetHack has been eagerly awaited by fans of that game for the last thirteen years. This shadowy group behind the update, known by fans simply as DevTeam, can be very tight-lipped about what they're up to. The community has generally viewed them with a sort of worshipful awe as they have slowly added new depth and sophistication to the game with each iteration. (A popular catchphrase is TDTTOE, or "The DevTeam Thinks of Everything.")

The release of the update seemed like a great time to talk to the developers of this beloved title, about the past and future of the game, and the devotion of the fan community that makes its ongoing development possible.

http://gamasutra.com/view/news/269726/The_story_behind_NetHacks_longawaited_updatethe_first_since_2003.php

Delete Your Linux Hosting Business

Did you hear about the owner of a web-hosting company who claimed to have erased his entire business with a single script command? It turns out he's a believer in the adage that no advertising is bad advertising – he made the whole thing up.

Marco Marsala of Italy posted a cry for help on Server Fault in April claiming he'd accidentally erased all the data on his servers including backups (which were purportedly mounted from remote storage).

"The moderators on Server Fault have been in contact with the author about this, and as you can imagine, they're not particularly amused by it," Stack Overflow said in a statement.

It's not clear yet whether there's a moral to this story.

RHEL 5 EOL Approaching

If you're a Red Hat Enterprise Linux shop, you ought to be planning your upgrade soon. On the other hand, if you have deep pockets, you've got a few years yet.

> In accordance with the Red Hat Enterprise Linux Errata Support Policy, support for Red Hat Enterprise Linux 5 will be retired on March 31, 2017, at the end of Production Phase 3. Until that date, customers will continue to receive Critical impact security patches and selected urgent priority bug fixes for RHEL 5.11 (the final RHEL 5 release).
>
> However, we recognize that some customers will wish to remain on Red Hat Enterprise Linux 5 even after the March 31, 2017 retirement date. To meet this customer requirement, Red Hat will offer customers the option to purchase the Extended Life Cycle Support (ELS) Add-On as an annually renewable subscription. This ELS Add-On provides customers with up to an additional three and a half (3.5) years of Critical impact security fixes and selected urgent priority bug fixes for RHEL 5.11. RHEL 5 ELS coverage will conclude on November 30, 2020.

The VAX Platform Is No More

This may not be news to some people, but it was announced recently that after much internal discussion, OpenBSD has officially discontinued support for the VAX architecture. It appears that 5.9 will be the last release to support these machines from the past (and the present, in some cases).

http://undeadly.org/cgi?action=article&sid=20160309192510

On the other hand...

OpenVMS 8.4-2 Released

VMS Software, Inc. (VSI) today announced the worldwide availability of VSI OpenVMS Version 8.4-2 (Maynard Release) operating system for HPE Integrity servers. The Maynard Release is the second by VSI. The new OS is compatible with HPE Integrity servers running the latest Intel Itanium 9500 series processor *[They're still making those? -Ed.]*, as well as most prior generations of the Itanium processor family. VSI also reconfirmed plans to offer OpenVMS on x86-based servers.

"This second release reaffirms our long-term commitment to the OpenVMS platform, and builds upon our highly successful first release of OpenVMS in June of 2015," said Duane P. Harris, CEO of VMS Software. "It is the first of many exciting improvements planned for OpenVMS, including future updates to the file system, TCP/IP, and other major improvements that we look forward to sharing with our customers as we work our way through the planned roadmap."

Quote Of The Month

"Our recent experience with the Linux scheduler revealed that the pressure to work around the challenging properties of modern hardware, such as non-uniform memory access latencies (NUMA), high costs of cache coherency and synchronization, and diverging CPU and memory latencies, resulted in a scheduler with an incredibly complex implementation. As a result, the very basic function of the scheduler, which is to make sure that runnable threads use idle cores, fell through the cracks."

Jean-Pierre Lozi et al

LXD 2.0 Released

LXD is a relatively new entrant in the container-management arena; the project started roughly a year and a half ago. It provides a REST-based interface to Linux containers as implemented by the LXC project. LXD made its 2.0 release on April 11, which is the first production-ready version.

At its heart, LXD is a daemon that provides a REST API to manage LXC containers. It is called a "hypervisor" for containers and seeks to replicate the experience of using virtual machines but without the overhead of hardware virtualization. LXC containers are typically "system containers" *[see also Solaris 'zones' -Ed.]* that look similar to an OS running on bare metal or a virtual machine, unlike Docker (and other) container systems that focus on "application containers". The intent is to build a more user-friendly approach to containers than what is provided by LXC.

The easiest ways to get started with LXD are all based on Ubuntu systems, which is not surprising given that Canonical is the main sponsor of the project. There are provisions for other distributions (Gentoo, presently) and for building the Go code from source, though. There is also an online demo that can be used try out LXD from a web browser.

The LXD daemon uses a number of kernel technologies to make the containers it runs more secure. For example, it uses namespaces and, in particular, user namespaces to separate the container users from those of the system at large. As outlined in lead developer Stéphane Graber's Introduction to LXD one of the core design principles "was to make it as safe as possible while allowing modern Linux distributions to run inside it unmodified".

A container in LXD consists of a handful of different pieces. It has a root filesystem, some profiles that contain configuration information (e.g. resource limits), devices (e.g. disks, network interfaces), and some properties (e.g. name, architecture). The root filesystems are all image-based, which is something of a departure from the template-based filesystems that LXC uses. The difference is that instead of building the filesystem from a template when the container is launched (and possibly storing the result), LXD uses a pre-built filesystem image that

typically comes from a remote image server (and is then cached locally).

These images are generally similar to fresh distribution images like those used for VMs. LXD is preconfigured with three remote image servers (for Ubuntu stable, Ubuntu daily builds, and a community-run server that has other Linux distributions). The images themselves are identified with an SHA-256 hash, so a specific image or simply the latest Ubuntu stable or daily build can be requested. Users can add their own remote LXD image servers (either public or private) as well.

Profiles provide a way to customize the container configuration and devices. A container can use multiple profiles, which are applied in order, with later profiles potentially overwriting earlier configuration entries. In addition, local container configuration is applied last for configuration entries that only apply to a single container, so they do not belong in the profiles. By default, LXD comes with two profiles, one that simply defines an "eth0" network device and a second that is suitable for running Docker images.

LXD uses Checkpoint/Restore In Userspace (CRIU) to allow snapshotting containers, either to restore them later on the same host or to migrate them elsewhere to be restored. These container snapshots look much the same as regular containers, but they are immutable and contain some extra state information that CRIU needs to restore the running state of the container.

LXD needs its own storage back-end for containers and images. Given the announcement that Canonical will be shipping ZFS with Ubuntu 16.04, it will not come as a surprise that the recommended filesystem for LXD is ZFS. But other options are possible, as described in another post in the series. In particular, Btrfs and the logical volume manager (LVM) can be used.

LXD can scale beyond just a single system running multiple containers; it can also be used to handle multiple systems each running LXD. But for huge deployments, with many systems and thousands of containers, there is an OpenStack plugin (nova-lxd) that provides a way for the OpenStack Nova compute-resource manager to treat LXD containers like VMs. That way, LXD can be integrated into OpenStack deployments.

As befits the "production-ready" nature of the release, LXD 2.0 has a stable API. Until June 2021, all of the existing interfaces will be maintained; any additions will be done using extensions that clients can discover. In addition, there will be frequent bug-fix releases, with backports from the current development tree.

There is a fair amount of competition in the container-management (or orchestration) world these days.Kubernetes, Docker Swarm, Apache Mesos, and others are all solving similar problems. LXD looks like it could have a place in that ecosystem, especially given the strong support it is receiving from Canonical and Ubuntu. For those looking for a container manager, taking a peek at LXD 2.0 may be well worth the time.

Door Prize

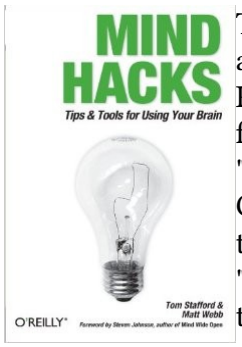

This month we will be giving away as a door prize either "Mind Performance Hacks: Tips & Tools for Overclocking Your Brain" or "Using csh & tcsh", both on O'Reilly. The former is the sequel to the previous door prize book "Mind Hacks", and is geared more towards improving your mental capabilities. The latter is an oldie-but-goodie for us fans of tcsh; no newer book has been written since, and not much has changed in the tcsh world anyhow!

MUUG Survey

Attention all members! The MUUG board wants your opinion. Do you enjoy the presentations? What are we not presenting enough of? Do you like the location? Hate the pop? We want to know! So to help you answer these questions we've put together an online survey and we'd love it if every member would take the time to give us some feedback. The link for survey is located below, and if you're reading this in the paper newsletter, don't worry, we've sent a link to your inbox too.

http://fluidsurveys.com/surveys/matcho-matics/muug-survey/