LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY

- LIGO -

CALIFORNIA INSTITUTE OF TECHNOLOGY

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Document Type	**LIGO-T070102-00-Z**	2007/09/05

Post-S5 Operating System for the LIGO Data Grid

Paul Armor, Carsten Aulbert, Lisa Bouge, Duncan Brown, Gerald Davies, Erik Espinoza, Henning Fehrmann, Kevin Flasch, Steffen Grunewald, Ben Johnson, Tyler Petire

Distribution of this draft:

LIGO Scientific Collaboration

California Institute of Technology
LIGO Project - MS 51-33
Pasadena CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS NW22-295
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

WWW: http://www.ligo.caltech.edu/

Processed with LaTeX on 2007/09/05

Contents

1 Summary

The current reference operating system (OS) for the LIGO Data Grid (LDG) is Fedora Core 4 (FC4). This OS is no longer supported and is scheduled to be upgraded shortly after the S5 run ends in September. The authors of this document have been requested by the LSC Computing Committee to select a replacement OS for this upgrade. This document reports our recommend choices and justification of that choice.

After consideration of the factors described in detail in the next sections, the conclusion of the committee is that there are two operating systems suitable for the next LDG OS: CentOS 5 and Debian Linux. CentOS 5 is a free community-supported GNU/Linux distribution derived from the commercially supported Red Hat Enterprise Linux (RHEL) distribution. Debian is a community-supported GNU/Linux distribution managed by the Debian Project.

CentOS 5 fulfills the criteria requested for long term support for bug, fixes, security and new hardware support, due to being derived from the RHEL distributions which have five year support. It is compatible with our current middleware and end-user software and there is sufficient expertise in the LSC to make transition from our existing operating systems to CentOS 5 straightforward.

Debian is also compatible with our current middleware and end-user software, although not quite as seamlessly as CentOS at present. This is due, however, to the current prioritization of Red Hat derived operating systems in the LSC. We see no reason why Debian support could not be brought up to the current level of CentOS support, if Debian is selected as the OS of choice. The LSC Computing Committee should note that the life cycle of Debian is not fixed, but rather Debian is released when the community decides that a new version is ready. This release cycle is much longer than Fedora Core, but as a result we would expect a Debian OS installed at the end of S5 to be supported for around three or four years, rather than the five years of support expected

for CentOS 5. If five year support is required, then the free Ubuntu Server Long Term Support (LTS), a Debian derived distribution supported by Canonical Ltd. is a viable option instead of the core Debian distribution. Again, there is sufficient expertise in the LSC to make transition from our existing operating systems to Debian or Ubuntu straightforward.

The committee sees two advantages of CentOS, both of which derive from it's RHEL heritage. Red Hat have announced that they will form a collaboration with the Condor team. To this end, they have opened an office in Madison staffed by Red Hat personal and have stated that they intend to closely integrate Condor into RHEL, and back port their fixes to the Condor source. It is hard to speculate where this collaboration will lead, but Red Hat have also invested heavily in the Xen virtual machine project and a goal of the Condor team is to support checkpointing of compute jobs using virtualization. One might expect that such functionality would appear first on RHEL (and hence CentOS), given this collaboration. The second advantage of CentOS is support from Sun Microsystems for native QFS file system support. Native QFS support will allow direct mounting of the LIGO Laboratory archive file systems on Linux nodes (rather than the present indirect route via Solaris and NFS). Sun currently support SuSE Enterprise 8, 9 and 10 and RHEL 3 and 4. Support for QFS on Debian has not been announced by Sun, but might be readily available after SUN opened the source code for QFS (as announced in April 2007).

Advantages in favour of Ubuntu/Debian are that these distributions are extremely easy to install and use especially for end-users. Although being not directly the focus of this document, having a server set-up which is almost identical to scientists' systems would ease the usage of post-S5 servers. Also, Ubuntu/Debian feature and share an extremely sophisticated package management with a single packaging format. Unlike for RPM-based distribution where the sometime incompatible file formats prohibit this, almost every time one can take a package from eiher Ubuntu or Debian and easiliy port it to the other. Also, the upgrade-manager tool allows a seamless upgrade from one Ubuntu version to the other while retaining personal settings in /etc.

Based upon this superior package management, making currently available more than 20000 packages, one can use the FAI package which allows a very easy and fast path to install many servers/compute nodes with extremely small overhead within a few minutes. Even automatic install time packge selection is possible, depending of the automatically detected hardware - thus no need to pregenerate cloning images anymore. There is also the possiblity to update a full cluster "softly", i.e. without a complete re-install.

Finally, Ubuntu is currently the most successful distribution with the largest user base and is being delivered pre-installed by Dell.

We expect that, regardless of the final choice, the LSC will still have to invest in the manpower required to package certain middleware and user tools for the selected distribution (e.g. the LIGO/Virgo Frame library), but we do not see a strong argument in favor of either distribution with regard to the availability of packages. The Debian distributions have a larger package base than CentOS and so if CentOS is selected, the LSC will have to continue to maintain some additional middleware such as GSL and FFTW3. We do not see this as a significant disadvantage, however.

It is the opinion of the committee that migration from the current LDG operating systems (Fedora Core 4 and Debian) to either CentOS 5 or Debian/Ubuntu is practical and requests the LSC Computing Committee to either select one of these operating systems, or to charge this committee with specific instructions to allow a choice to be made.

2 Introduction

The authors of this document have been requested to recommend a choice for the next reference platform for the LIGO Data Grid (LDG) to the LSC Computing Committee. Our goal was to select an OS for the LDG that tries to balance, as well as possible, the various selection criteria listed below to provide as scientifically productive LDG as possible.

The authors, and their fields of expertise, as pertaining to this committee are:

- Paul Armor: UW-Milwaukee Tier 2 Cluster
- Lisa Bouge: CDS Interface
- Duncan Brown (Chair): Analysis Software
- Erik Espinoza: LIGO Tier 1 Cluster
- Kevin Flasch: LDR/Globus Middleware
- Steffen Grunewald: GEO Tier 1 and 2 Clusters
- Ben Johnson: Observatory Clusters
- Tyler Petire: Penn State Tier 2 Cluster
- Henning Fehrmann: GEO Tier 1 Cluster
- Carsten Aulbert: GEO Tier 1 Cluster
- Gerald Davies: GEO Tier 1 Cluster, Virtualization Support.

This committee was created in April 2007 and requested to report by the end of June 2007.

The operating systems considered by the committee were as follows:

1. Commercial GNU/Linux distributions:
   (a) Red Hat Enterprise Linux
   (b) SuSE Enterprise Linux
2. Free (community-supported) GNU/Linux distributions:
   (a) CentOS
   (b) Scientific Linux
   (c) Debian
   (d) Ubuntu
   (e) Fedora Core
3. Non-Linux operating Systems:
   (a) Solaris
   (b) Free BSD

The committee quickly eliminated Solaris and Free BSD from consideration. There is no Matlab support for Solaris on our x86 cluster architecture, and no announcement of support is forthcoming from Mathworks. Given the heavy use of Matlab in the LSC for scientific analysis, this eliminates Solaris from consideration, despite its advantages of native ZFS file system support, and better support for NFS than Linux. The committee also feels that, although BSD has an application binary support for Linux, there is too little experience in the LSC to seriously consider BSD for the LDG. There was also concern that the community and user support for BSD is not as extensive as GNU/Linux. This factor, coupled with the lack of back-porting of security patched, would means that the LSC would have to take on a larger support role for a BSD-based choice. The committee agreed that this was not viable.

The committee consider the use of virtualization to allow a cluster administrator to install, for example, Solaris on the compute nodes, but then present a GNU/Linux distribution to user via a virtual machine. The consensus was that virtualization technology is not sufficiently mature at present to allow this.

We also eliminated the commercial GNU/Linux distributions due to the their prohibitively large cost. It is the opinion of the committee that the financial burden imposed by these operating systems is too great for them to warrant serious consideration. The experience of the Penn State tier 2 center suggests that the functionality provided by these commercial distributions can be obtained from the derived community-supported versions.

Given the remaining free GNU/Linux distributions, we eliminated Scientific Linux in favor of CentOS. Both of these distributions are derived from the commercial RHEL distribution, however CentOS appears to be more tightly coupled into the Red Hat project and has a larger community support base than Scientific Linux. We also eliminated Fedora Core based on our previous experience with this distribution. The release schedule of Fedora Core is far too rapid to provide the long term support required for the next generation operating system, and there have been major concerns raise by administrators about the quality of Fedora releases.

Having eliminated the above operating system, the remaining choices (CentOS, Debian and Ubuntu) were evaluated in detail by the committee, and our findings are reported below.

3 Factors Considered

3.1 Security

CentOS

With regards to security, CentOS 5 should more than suit our needs, as the LSC is already used to dealing with Red Hat distributions and we already are familiar with good practices to secure Red Hat boxes. This distribution also has been very recently released, and as it should inherit security updates and bug fixes for RHEL 5, CentOS 5 should be well supported well past our required life.

Ubuntu/Debian

Ubuntu releases have a standard life time of at least 18 months. Every 18-24 months they will release a long term support (LTS) version which has security support for 3 years for desktop and 5 years for servers. The last LTS release was 6.06 (June 2006) and the next LTS will most probably be 8.04 (April 2008).

(side note: Debian guarantees security support for 12 month after the next release is out, e.g. if the next release is out in 18 months time, the overall lifetime is 30 months)

3.2 Stability/OS support

CentOS

CentOS 5 inherits the stability of RHEL 5. This should provide us a very stable base upon which we can build our wanted/needed customizations. The effects of using a ”bleeding edge” distribution should have resulted in a more pleasant experience for our users from the standpoint of what packages were available, how new the versions of those packages were, and how quickly new projects/packages were available. Unfortunately, I think the dynamism that should have resulted in these pluses for our users actually hindered things sufficiently that any benefits were outweighed by the instability of the distribution as a whole. CentOS 5 will not always have the latest and greatest sought by some users, but the base should be stable enough that we can add custom updates as needed.

Ubuntu/Debian

Ubuntu tries to utilize the latest stable versions available of each package for each release, especially the Linux kernel. The distribution within itself is coherent and partitioned into four parts: ”main” and ”restricted” are actively supported by Canonical staff while ”universe” and ”multiverse” are community supported repositories. The overall quality of the last releases was very good and seamless upgrades from one relese to the next is supported where possible - also from LTS to LTS.

Since Debian an Ubuntu share a wide range of packages it is usually very easy to take a Debian package and simply rebuild it against an Ubuntu system and vice versa).

3.3 Hardware Support

OS Support (quality and duration at least until AdvLIGO shutdown)

CentOS

With regards to Hardware Support, the incremental releases should suffice in addressing any needs for future driver updates; and UWM already needs to build custom kernels with modified drivers anyway, so needing to rebuild our own is something we’re already used to.

The CentOS system has often rebuilds of their install media that will result in hardware support for the newest hardware. During the OS-Committee run we tested this by installing the latest rebuild of the 4 year old CentOS 3 branch on the latest and greatest Socket F Opteron hardware. Although there were some issues with other types of SATA controllers, this had more to do with the 2.4 kernel rather than RH’s investment in keeping an old base out of date. We did not experience that with the 2.5 year old CentOS 4 branch which is based on a 2.6 kernel.

Ubuntu/Debian

All hardware covered by the vanilla kernel is also supported by Ubuntu, plus many modules are provided by extra packages to allow newer hardware or hardware requiring binary blobs to work.

Ubuntu is currently released for three architectures, x86, AMD64 and UltraSparc T1.

(side note: If more architectures are needed, Debian would be more appropriate)

3.4 Virtualization Support

Virtualization support (e.g., LSC laptop virtual machines)

What are our options for creating an LDG virtual machine that would run analysis jobs under Condor?

CentOS

RHEL (and thus CentOS) is a target of any major virtualization project (eg. Xen and KVM), so kernel/OS level virtualization should already be well supported/maintained, and RHEL is a major target of the Condor group, so any Condor virtualization we may wish to pursue should also be well covered.

Ubuntu/Debian

QEMU, kvm, OpenVZ and Xen are supported by own packages. Vmware works without problems.

3.5 Middleware Support

1. Condor: Support for Condor universes: vanilla, standard, VM, MPI.
2. Globus and grid middleware support. pyGlobus.
3. LDG Client/Server Bundle: client/server, segfind client/server.

CentOS

RHEL, and hence CentOS is already a primary target of each of the middleware products we currently use, and so our current middleware will be fully supported.

Ubuntu/Debian

Although not directly supported by the maintainers, our middleware seems to work fine on Debian. Initial tests on Ubuntu were also successful. However, it would be nice to have more generic support for Debian/Ubuntu, especially the Condor standard universe which is not presently supported.

3.6 Application Support

Analysis package support: Matlab, Root, DMT, LDAS, C/C++ code.

CentOS

Matlab installs easily and should be well supported (again thanks to RHEL 5); other packages should already be supported or be quite easily, given our current use of FC4. A concern in this regard is how well the gnu compilers and build utilities will hold up over time; an example that springs to mind is having to custom update automake/autoconf before the LSC moved beyond RH9.

Ubuntu/Debian

Matlab, Mathematica install fine and are easy to use on Ubuntu, ldas was installed and briefly tested back in 2001 or 2002. It worked a bit, but it was a pain to set-up.

3.7 Filesystem support

Cluster filesystems, native QFS, ZFS.

CentOS

A full array of journaled file systems are available; cluster file systems SHOULD be well supported, but the committee did not have time to evaluate any. ZFS is only available via FUSE for any Linux distribution due to license limitations. QFS support is available for RHEL, which would allow native mounting of the LIGO Lab archive file systems.

Ubuntu/Debian

All major filesystems (either clustered/local) are supported (either through the kernel or through FUSE). ZFS is only available via FUSE for any Linux distribution due to license limitations.

3.8 Support for Scientific Users and Developers

Desktop and laptop support for software developers.

CentOS

There appear to be no major disadvantages to using CentOS for users, however as users tend to stick with newer (bleeding edge) desktops such as Fedora and Ubuntu, CentOS will be a stable base that doesn't change. This means any development done on such systems may run into issues. We currently run into this as people attempt to run their python 2.5 scripts on our python 2.4 system. This can be mitigated with ldcg binaries and custom RPMS, as well as a VMware Virtual Machine to provide an LDG reference platform for users.

Ubuntu/Debian

Ubuntu was mainly a desktop centric distribution, but has added extra support for servers and clusters in 2006. The packages are (almost) the same, so developers will find the very same environment on their desktops/laptops as they will find on their servers.

3.9 Ease of Transition

Ease of transition from the current Reference Platform (e.g., cluster wide installation tools, how much LSC software needs to be modified).

CentOS

CentOS 5 should take little work to update to. Rebuilding OS installs should be relatively quick, and rebuilding our own software products went pretty smoothly with minimal effort. As FC4 is not very different from CentOS 5, we will not have to rewrite all of our kickstart configurations and install scripts from scratch. This will be a trivial migration on the system side. LSC end-user software (LAL, LALApps and Matlab) has been tested on CentOS 5 and found to be fully functional.

Ubuntu/Debian

Ubuntu can either use a mechanism called ”pre-seeding” which mimics a interactive set-up with the standard text based installer, but a much more powerful way to install it is using the (Debian-)package FAI (Fully Automatic Installation). At AEI/Potsdam this system has been used very successfully over the past years and initial tests at AEI/Hannover showed that it is a very versatile package allowing easy installation of heterogeneous environments with a common base system and only a few extra bits depending on the hardware or type of target machine. It also provides means to install images, thus installing Fedora, CentOS, Suse or even Gentoo is possible/seems to be possible.

There is even a kickstart compatibility package which promises to reuse existing scripts, but that has not yet been tested.

3.10 Ease of Maintenance

Ease of LSC software maintenance and distribution (e.g., packaging)

CentOS

CentOS uses the same RPM/yum packaging as our current reference platform, so little extra effort would be involved transitioning.

Ubuntu/Debian

Currently not too easy to install LSC software, since Debian/Ubuntu is not an officially supported platform. pacman based installations work most of the time, but rpm-based software needs to be repackaged. For most programs passing the rpms through alien works, but might need some tweaking afterwards. Would be much better if the LSC could simply offer a repository which users can add to their system and use. Also this would be much faster and user friendlier than pacman installs and upgrades would be much easier to handle.

3.11 Cost

Not just cheaper is better, but consider the new PSU model of paying for a higher level of support on servers but having a free option for cluster nodes.

CentOS

CentOS is free, the cost of retraining should be minimal, and the time to convert should be minimal; overall this should be no more expensive than updating Fedora releases. Since CentOS is tightly

coupled to RHEL, it will not prohibit LDG sites from installing RHEL on key servers, if they so desire.

Ubuntu/Debian

Debian/Ubuntu is available free of charge but support can be bought. Ubuntu/Canonical offer paid support with current yearly prices are $750 for 9x5 or $2750 for 24x7 support for servers or $1200/$4000 for cluster support (http://www.ubuntu.com/support/paid).

The Ubuntu community is known to be the most active community right now to help with problems, either by IRC, launchpad.net, it's large wiki section or mailing lists.

3.12 Uniformity

Meets as many site specific issues as possible to ensure as wide an adoption as possible. Is it possible to achieve uniformity through virtualization?

CentOS

Uniformity is a potential weakness of CentOS; users may not find all they want and need; but in our opinion is that with some reasonable amount of effort, we could support whatever our users required.

Ubuntu/Debian

Debian/Ubuntu can be run on laptops, servers, compute nodes without much hassle and present the users a uniform look and feel everywhere. Even if users run many different versions of Ubuntu (or even Debian in parallel) the basic look and feel is still the same, regardless if people were still using 5.10 or 7.04.

4 LSC Computing Group Impact

It is very likely that CDS will standardize on CentOS, Debian or Ubuntu, depending on what the LDG adopts (barring any technical difficulties, but we don't expect to encounter any).

No LDAS issues were reported to the committee.

5 Conclusion

After consideration of the factors described in detail in the next sections, the conclusion of the committee is that there are two operating systems suitable for the next LDG OS: CentOS 5 and Ubuntu Linux. CentOS 5 is a free community-suppoted GNU/Linux distribution derived from the commercially supported Red Hat Enterprise Linux (RHEL) distribution. Ubuntu is largely a sponsored GNU/Linux distribution managed by Canonical Inc. with an additional community-supported part.

It is the opinion of the committee that migration from the current LDG operating systems (Fedora Core 4 and Debian) to either CentOS 5 or Debian/Ubuntu is practical and requests the LSC Computing Committee to either select one of these operating systems, or to charge this committee with specific instructions to allow a choice to be made.

A viable solution might also be to choose to support two distributions to minimize the effect of transitions from the current situation. In the long run this would most certainly mean a slightly increased work-load for packagaing LSC related software, but might also increase the portability of software solutions.