# Pete's All Things Sun

## Solaris 11 Express

PETER BAER GALVIN

Peter Baer Galvin is the chief technologist for Corporate Technologies, a premier systems integrator and VAR (www.cptech.com). Before that, Peter was the systems manager for Brown University's Computer Science Department. He has written articles and columns for many publications and is co-author of the *Operating Systems Concepts* and *Applied Operating Systems Concepts* textbooks. As a consultant and trainer, Peter teaches tutorials and gives talks on security and system administration worldwide. Peter blogs at http://www.galvin.info and twitters as "PeterGalvin."
pbg@cptech.com

To some of us, Solaris 11 Express seems like yesterday's news. It bears much resemblance to the OpenSolaris distributions, which stopped when it was released. The first OpenSolaris release was OpenSolaris 2008.05, meaning that some sites have been using at least some of the new features for almost three years. However, many of you sysadmins have not played with either OpenSolaris or the just-released Solaris 11 Express. This column is for you.

This column is also for admins who are familiar with OpenSolaris but not yet with Solaris 11 Express. Here I will discuss what is new in Solaris 11 Express, differences from both Solaris 10 and OpenSolaris. The discussion is mostly technical but also includes details on the legal, support, and production status of both releases. If you are already familiar with Solaris 11 Express and its features, then feel free to move along—nothing to see here.

### Solaris 11 Express—Should We Care?

Certainly some sites became less interested in Solaris when Oracle bought Sun. Others became disenchanted with the long delay between Oracle's purchase and any word on the future of Solaris, and others when an internal Oracle memo about the future of Solaris made the rounds (see http://opensolaris.org/jive/thread.jspa?messageID=496203). Is Solaris still open? Or is it closed? And should we care about its state or even about its existence?

For many sites, Solaris is a key operating system and will likely continue to be key. They run production Solaris 10 (S10) or earlier releases and wait for production updates before upgrading their systems. For those sites, Solaris 11 Express (S11E) is good news. The Solaris 10 releases have had decreasing incremental improvements over time as Solaris engineers worked on the next generation of Solaris. Solaris 10 shipped on January 31, 2005, and several major changes, including ZFS, shipped in subsequent releases. However, it is unarguable that there are many nice new features in S11E that are not available in S10. Given the long gestation of OpenSolaris, it is encouraging that the first release (even though an "express" release) of Solaris 11 is available. S11E is available for SPARC and x86, affirming Oracle's plan to support both architectures going forward.

The "openness" of Solaris is likewise of importance to some sites, but not others. Many sites did not change their behavior when Solaris was open sourced, and so should not consider any change in openness to be important. However, some sites

do care about the open state of Solaris, whether for philosophical or practical reasons.

Philosophically, many feel that open source operating systems are fundamentally superior to closed ones. Sometimes this feeling of superiority comes from the benefit of non-employees contributing to the code, improving it faster than just the hired engineers could. Other times it comes from the idea of supporting open source efforts, or a basic belief that software's natural state is open and free.

Practically, some like open source because they can make their own distribution or base appliances or products on the source, whether a commercial or free effort. There are many instances of such efforts around OpenSolaris, and many feel that such an ecosystem is a contributor to the health of an operating system. Even more practically, at sites that simply run the commercial distribution based on an open source operating system, they can read source code for debugging, tuning, and general understanding.

Which brings us to the question of the current openness of Solaris. There has been no official word that OpenSolaris is dead or closed. In fact, even the leaked memo states that the CDDL license will not be removed from any code that was labeled with it, fundamentally leaving the code open. But the memo also states that source code will not be released until after the commercial version of Solaris that is based on it is released. Which leads us to the assumption that, once S11 ships, the source code for most of it will be made available. That would likely go far in terms of calming the fears of open source fans and Solaris fans alike.

We should also care about Solaris 11 and its future because, frankly, it has some very nice features—features that many sites would find useful and would enjoy using in production. I'll discuss those features right after addressing the legalities.

## Licensing and Support

An operating system can be feature-rich and still not be used, due to its costs and legal limitations. So, what is the status of S11E? The FAQ that was released by Oracle is a good place to start for all of the business-side details [1]. But, in summary, it seems that anywhere you are allowed to run S10, S11E is allowed. If S10 is licensed on a given server, then S11E can be used there as well. Separate support for S11E is available, just as with S10. Further, S11E is usable with a support contract for evaluation and development purposes under the Oracle Technical Network perpetual license. Fundamentally, S11E is a full, supported next-generation release of Solaris. Then why the "Express" designation? There are many changes from S10 to S11, and ISVs need to have a stable code base with which to port or validate their applications. "Express" is also a bit of a warning that S11 is young and may not be appropriate for production use. S11E is a full commercial release, but it should be used with caution until the first S11 release ships.

## Features

Several features of S11E were topics in this column as they came out in OpenSolaris. Therefore, rather than re-covering them, please see the appropriate columns.

- Crossbow, the network quality of service and virtualization feature set was discussed in the February 2010 *;login:*, [2].
- Another major new feature of S11E, and perhaps the biggest difference, is the new package management system. This, along with ZFS as the only allowed root

file system and the new installer and boot environment manager, were discussed in the December 2008 *;login:* [3].

- ZFS de-duplication was covered in the April 2010 *;login:* [4].

Some tried-and-true features of Solaris carry through to S11E, including binary compatibility with previous Solaris releases. Old package management still works in parallel with the new IPS packaging system, but there is no patching for those old packages (and certainly not for the new style), but there are certainly many new and different features.

The installer is new, simpler, and better, due to its support of ZFS as the only root file system. Also gone is the old jumpstart, replaced with an "automated installer" that allows customized hands-off installation of multiple Solaris systems.

Unfortunately, there is no true upgrade path from S10 to S11E. This marks a large change from previous Solaris releases, but is an indicator of just how different S11E is from S10. It is possible that there will be an upgrade path included in a future S11 release, but there is no path available today. There is some help for S10 systems, though, in the form of a new "Solaris 10 Container" feature within S11E. An S10 system can be archived and installed within an S10 container inside S11. Oracle says that all applications that ran on S10 will run within an S10 container on S11E. But note that if the S10 system has containers, it cannot be run inside S11E within an S10 container. In other words, there is no concept of containers-within-containers, so all containers in the S10 system must be removed before attempting to encapsulate that system within an S10 container. And even though there is no patching of S11E, there is still patching with Solaris 10 containers (as well as S9 and S8 containers, of course).

S11E itself will be upgradable to S11 when it ships using the new package management system. Much like Linux, S11 can determine the list of all packages that need to be updated, including kernel packages, and download and install them. A new boot environment is created as necessary, and the system can then be booted into the new or previous environments. There is now a "fast boot" option that skips the hardware diagnostics phase of booting. There is also a new GUI "Update Manager" tool which lets the admin download and install packages that are not yet on the system or update those that are. For example, a simple `% pfexec pkg install gcc-3` brings in the Gnu C compilation environment via the command line. Or within the update manager, typing the search term "emacs" lists both installed and available-to-install packages that match the term. Selecting the ones desired and clicking "install/update" installs or updates the packages as appropriate.

ZFS and the underlying file system structures have new features beyond what is available in S10. ZFS gets deduplication and encryption, both major features. The new ZFS diff feature will show what changed between two snapshots. CIFS is now a fully integrated kernel feature, rather than a set of user-land programs. Also, all of the various SCSI protocol implementations that were within S10 have been merged into a single COMSTAR (Common Multiprotocol SCSI Target) facility. At the user level, the "time slider" GUI tool is quite an improvement. It can automate the creation of ZFS snapshots and can also help users visualize the snapshots by showing which files were available at any given time (between the oldest existing snapshot and the current version of the file system). Erwann Chenede has a complete blog posting, including a video demo, exploring time slider [5].

Interesting new commands include zonestat to display per-zone information, flowstat to show Crossbow networking information, and dlstat to show network link statistics. While on the topic of networking, many sysadmins will be pleased to hear that DTrace has been enhanced to be able to observe much deeper into the networking stack, including to and from Solaris containers. Each container can now have one or more dedicated virtual network ports, rather than a dedicated or shared hardware port as in S10. Also, there is a software layer 3/layer 4 load balancer included in S11E. Many other networking improvements are included, such as improved InfiniBand, link protection, and bridging and tunneling technologies.

Security changes are also numerous, including cryptographic framework improvements, trusted extension enhancements, and in-kernel pfexec implementation.

Oracle has published some "what's new" documents that go into detail on all of these changes and many more [6, 7, 8].

So, is Solaris 11E ready for use? Development and Q/A environments can use it as is, giving sysadmins experience with the new features and getting developers ready for the new package management system. In my uses, it has been rock-solid. The performance, security, and feature enhancements are all impressive, useful, and welcome. And Oracle is placing its trust in S11E where its hardware is—Solaris 11E is an option on the Exadata and Exalogic appliances. Apparently, the improved InfiniBand stack made S11E more appropriate than S10 for those systems, among other reasons.

## The Future

Oracle has stated that Solaris 11 (the official production release) will be available in 2011. Even if you choose not to try S11E, its features will be available within a reasonable amount of time in the S11 release. But S11E, whether you choose to use it for production, testing, or just exploration, should be a very practical next step into the use of Solaris in your environments.

## Tidbits

Life is good for production and performance-oriented sysadmins, because there is a new release of the Chime performance monitoring tool. Chime has been a bit of a proof-of-concept and science experiment until now, but the new features make it quite useful and more powerful. The DTraceToolkit scripts have been incorporated, making Chime a good first stop for lighting up various aspects of the system to understand their performance characteristics. Chime is available for download from http://hub.opensolaris.org/bin/view/Project+dtrace-chime/.

#### References and Resources

For full details on the history of Solaris releases see http://en.wikipedia.org/wiki/Solaris_(operating_system). Other details of S11E are available from http://www.oracle.com/us/products/servers-storage/solaris/index.html, http://forums.oracle.com/forums/category.jspa?categoryID=303, http://www.oracle.com/technetwork/server-storage/solaris11/overview/index.html, and http://www.oracle.com/technetwork/server-storage/solaris11/documentation/index.html.

[1] http://www.oracle.com/technetwork/server-storage/solaris11/overview/faqs-oraclesolaris11express-185609.pdf.

[2] http://www.usenix.org/publications/login/2010-02/pdfs/galvin.pdf.

[3] http://www.usenix.org/publications/login/2008-12/pdfs/galvin.pdf.

[4] http://www.usenix.org/publications/login/2010-04/pdfs/galvin.pdf.

[5] http://blogs.sun.com/erwann/entry/zfs_on_the_desktop_zfs.

[6] http://www.oracle.com/technetwork/server-storage/solaris11/documentation/solaris-express-whatsnew-201011-175308.pdf.

[7] http://www.oracle.com/technetwork/server-storage/solaris11/documentation/s11sysadminwp101109final2-186770.pdf.

[8] http://www.oracle.com/technetwork/articles/servers-storage-admin/solaris11-dev-186782.pdf.