

PETER BAER GALVIN

Pete's all things Sun: the death of Solaris . . .



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 coauthor 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://pbgalvin.wordpress.com> and twitters as "PeterGalvin." His All Things Sun Wiki is <http://wiki.sage.org/bin/view/Main/AllThingsSun>.

pbg@cptech.com

THIS MONTH'S ISSUE OF ;LOGIN: IS about security. Unfortunately for me, my previous column about security (April 2008 [1]) is still pertinent and mostly up to date. To catch up on what's new with Solaris security, I recommend you have a look at Glenn Brunette's blog [2], where Glenn recently posted a presentation about the state of Solaris 10 security.

Now that we have security out of the way, we can get to the meat (or potatoes if you are a vegetarian) of this PATS: the first production release of OpenSolaris: OpenSolaris 2008.11. The debate still rages about the future of commercial Solaris, including a recent article in *InfoWorld* questioning whether Solaris is on its deathbed [3]. In the meantime, Sun is moving forward with a fork in the road of Solaris, releasing OpenSolaris in November 2008 [4]. This release is really a merge of many of the benefits of Linux and Solaris and changes the battlefronts in the operating system wars. Certainly one view of this change is that merging Linux usability and manageability with the Solaris industry-leading feature set creates a best-of-breed operating system. Other operating systems are going to have a tough time competing with this new world.

Keep in mind that this first release of OpenSolaris seems to be aimed at the desktop, not the datacenter server. For example, there is no SPARC release, only x86. Further down the road, expect OpenSolaris to become more datacenter-ready. Now let's take a look at the feature set, and then get hands-on with the new OpenSolaris release.

OpenSolaris 2008.11 Features

The biggest difference between Solaris and OpenSolaris is the new package management facility. The new system is similar to those found in Linux distributions. It is full-featured, installing operating systems, operating system updates, and user packages (everything else). It also provides full management of packages. For example, `pkg refresh` updates the list of packages available from the Internet-accessible package repositories (authorities), and `pkg install Y` downloads package Y and all things that Y depends on, makes sure all versions match, and installs all those packages. The same `pkg install Y` later on upgrades package Y to the current version. Additional non-official-Sun repositories can be added to the system's list of "authorities" via `pkg set-authority`, as in `pkg set-authority -O httpd://pkg`.

sunfreeware.com:9000 sunfreeware. All package search and install commands executed on the system use all authorities set in this manner.

The desktop GUI and user-software stack have been upgraded as well, including Gnome 2.24, the eSpeak (GPLv3) voice synthesizer, Firefox 3, Thunderbird, the Songbird music player, the Transmission bittorrent client, xChat, a Backup GUI (using ZFS snapshots), the Meta Tracker object database, a new printer configuration tool, and a more refined look and feel.

Packages available in the default repository include top, sudo, emacs, and many more favorites (but not a complete set of open source tools, yet). Some large packages include a Web stack that contains MySQL, Apache, CVS, DTrace modules for Apache and Ruby, and more. A developer package includes Ruby, GRails, SunStudio, NetBeans, and Eclipse.

One “feature” that is blissfully missing is a patching facility. Rather, packages maintain their consistency. If there is a bug in a program that is in a package, a new version of that package is created. Updating that package leads to all packages it depends on being updated as well. Now the bug is fixed but no patching happened (and none of the nightmares of dealing with patches happened).

OpenSolaris itself is installed and managed via this package management system. Approximately every two weeks, when a new set of OpenSolaris packages is released, the single command `pkg image-update` downloads all the new OpenSolaris bits and installs them, essentially upgrading the operating system. There is also a GUI package manager that shows which packages are installed, which are available, and which have updates available. It is likely that there will be less frequent and more official releases of OpenSolaris, to make it easier for Sun support and ISVs to support their software on OpenSolaris.

ZFS is the root file system, and it includes all of the ZFS goodness that has been documented elsewhere, including snapshots, clones, replication, and almost trivial administration.

The boot environment now has a manager and integrates use of ZFS as the root file system to great advantage. For example, as part of a `pkg image-update`, a ZFS snapshot is first taken. Once the `pkg` command is done, the new version of the OS is enabled as the default boot environment. `beadm` allows the listing, creation, deletion, and activation of boot environments. For example, `beadm activate old-boot` would set `old-boot`, an existing boot environment, to be the default boot environment and a reboot would return the system to its pre-upgrade state. There can be many different boot environments from which to select. Unneeded boot environments can be deleted (and the ZFS file systems destroyed to reclaim space) via a command such as `beadm destroy old-boot`.

Virtualization options abound in OpenSolaris, including the XVM (Xen-based) hypervisor and Sun’s open source desktop VirtualBox. VirtualBox [5] is especially interesting and useful, as it runs on many platforms and can virtualize many x86-based environments (Windows, Linux, and Solaris among them).

The security model of OpenSolaris is stronger than standard Solaris as well. By default, root is a role, not a user. Roles are part of the Solaris privilege umbrella of services and are a key part of the Role Based Access Control (RBAC) facility. Because of this change no logins are allowed to the root account. Rather, the user specified at install time has root as a role. Users with that role can `su` to root or can preface commands with `pfexec` to execute them with root role privileges. Root can be added as a role to other accounts via `usermod -R root username`. Such a change would allow that user to be able to assume the root role (by providing the root password). If you really need root to be an account, that change can be made with the command `rolemod -K type=normal root`.

One final change of note is the preeminence of the GNU tools and the bash shell. The bash shell is the default shell on the system, and the GNU version of tools such as tar come before the Solaris-standard versions of those tools in the default search path. This change freshens the user experience (and improves compatibility with Linux), but it can break some backward compatibility between Solaris and OpenSolaris. This is especially true where scripts are involved. So care should be taken (and testing done) when taking Solaris components and installing them and using them in OpenSolaris.

This list of variations from Solaris is extensive. However, believe it or not, the two share many features. These include secure-by-default settings, DTrace, zones/containers, SMF, and ZFS. Also, the System V package system is included (for backward compatibility). This allows old-school commercial Solaris packages to be installed on the system alongside new-style packages.

How to Get It

OpenSolaris is distributed as an ISO image of a LiveCD. A download of 700 MB, a blank CD and a CD-ROM writer later, and booting a system via that CD brings the new world of OpenSolaris. The LiveCD has an installer that puts OpenSolaris into a designated disk slice. Of course, another great way to try OpenSolaris without any long-term commitment is to install it via the ISO image into a virtual machine in your favorite virtual machine tool. For my testing I ran OpenSolaris natively on a generic white-box PC, as well as within VMWare Fusion (on Mac OS X).

The official OpenSolaris site is <http://www.opensolaris.com/get>; it has an ISO image of the 052008 release (and probably the 2008.11 release by the time you read this). Oddly enough, other ISO images are available from <http://genunix.org/>. Check that site for the latest available ISO images.

Hands On

For the purpose of these tests I installed OpenSolaris `osol-0811-96.iso` from genunix.org in VMWare Fusion 2.0 under Mac OS X 10.5.5. In VMWare I chose Solaris 10 64-bit as the guest operating system. The install went smoothly, and even VMWare Tools installs properly (but is missing some functionality).

Rather than focus on the GUI during my exploration, which is very Linux-like, I spent my time in the administration, which is different from both Linux and Solaris in some important regards.

Let's have a look at the state of the system post-install. Note that I shortened the prompts for ease of reading.

```
$ more /etc/release
```

```
OpenSolaris 2008.11 snv_96 X86
Copyright 2008 Sun Microsystems, Inc. All Rights Reserved.
Use is subject to license terms.
Assembled 29 August 2008
```

```
$ df -kh
```

Filesystem	Size	Used	Avail	Use%	Mounted on
rpool/ROOT/opensolaris	11G	2.6G	8.1G	25%	/
swap	2.0G	304K	2.0G	1%	/etc/svc/volatile
/usr/lib/libc/libc_hwcapi.so.1	11G	2.6G	8.1G	25%	/lib/libc.so.1
swap	2.0G	12K	2.0G	1%	/tmp
swap	2.0G	52K	2.0G	1%	/var/run
rpool/export	8.1G	19K	8.1G	1%	/export
rpool/export/home	8.1G	56M	8.1G	1%	/export/home
rpool	8.1G	58K	8.1G	1%	/rpool
/hgfs	16G	4.0M	16G	1%	/hgfs

```
$ zpool status -v
```

```
pool: rpool
state: ONLINE
scrub: none requested
```

config:

NAME	STATE	READ	WRITE	CKSUM
rpool	ONLINE	0	0	0
c4t0d0s0	ONLINE	0	0	0

errors: No known data errors

\$ zfs list

NAME	USED	AVAIL	REFER	MOUNTPOINT
rpool	4.66G	8.01G	57.5K	/rpool
rpool@install	17K	-	57.5K	-
rpool/ROOT	2.61G	8.01G	18K	legacy
rpool/ROOT@install	0	-	18K	-
rpool/ROOT/opensolaris	2.61G	8.01G	2.55G	legacy
rpool/ROOT/opensolaris@install	66.2M	-	2.44G	-
rpool/dump	1023M	8.01G	1023M	-
rpool/export	55.4M	8.01G	19K	/export
rpool/export@install	15K	-	19K	-
rpool/export/home	55.4M	8.01G	55.3M	/export/home
rpool/export/home@install	18K	-	21K	-
rpool/swap	1023M	9.01G	16K	-

Mirroring the disk is a one-command effort (assuming I have a similarly sized disk named c4t1d0).

```
# zpool attach rpool c4t0d0s0 c4t1d0s0
```

```
# zpool status -v
```

```
pool: rpool
state: ONLINE
status: One or more devices is currently being resilvered. The pool will continue to function,
possibly in a degraded state.
action: Wait for the resilver to complete.
scrub: resilver in progress for 0h0m, 4.46% done, 0h5m to go
```

config:

NAME	STATE	READ	WRITE	CKSUM
rpool	ONLINE	0	0	0
mirror	ONLINE	0	0	0
c4t0d0s0	ONLINE	0	0	0
c4t1d0s0	ONLINE	0	0	0

errors: No known data errors

Certainly that is easier than the old disksuite efforts that were needed to mirror root! Similarly easy steps could replace a failed member of the mirror set with a replacement, for example.

Now let's look at our boot environment.

```
$ beadm list -a
```

BE/Dataset/Snapshot	Active	Mountpoint	Space	Policy	Created
-----	-----	-----	-----	-----	-----
opensolaris	NR	/	2.61G	static	2008-09-05 15:04
rpool/ROOT/opensolaris@install	—	—	66.17M	static	2008-09-05 15:31

There is only one boot environment. By default OpenSolaris calls it, well, "opensolaris."

Shall we upgrade this installation to a newer OpenSolaris release?

```
# pkg refresh
Fetching catalog 'opensolaris.org'...
# pkg image-update
Fetching catalog 'opensolaris.org'...
```

```

Checking that SUNWipkg (in '/') is up to date...
WARNING: pkg(5) appears to be out of date, and should be updated before running image-update.

Please update pkg(5) using 'pfexec pkg install SUNWipkg' and then retry the image-update.
bash-3.2# pkg install SUNWipkg
Fetching catalog 'opensolaris.org'...

```

```

Creating Plan -

```

DOWNLOAD	PKGS	FILES	XFER (MB)
SUNWipkg	1/1	110/110	0.38/0.38
Completed	1/1	110/110	0.38/0.38

PHASE	ACTIONS
Removal Phase	12/12
Install Phase	7/7
Update Phase	151/151

PHASE	ITEMS
Reading Existing Index	8/8
Indexing Packages	1/1

```

bash-3.2# time pkg image-update
Fetching catalog 'opensolaris.org'...

```

```

Checking that SUNWipkg (in '/') is up to date...

```

```

Creating Plan -

```

PHASE	ITEMS
Indexing Packages	1/545
...	
Indexing Packages	545/545

DOWNLOAD	PKGS	FILES	XFER (MB)
SUNWpython-cherrypy	0/549	0/8025	0.00/219.59
...			
SUNWipkg-brand	549/549	7708/7708	218.50/218.50
Completed	549/549	7708/7708	218.50/218.50

PHASE	ACTIONS
Removal Phase	1/3281
...	
Removal Phase	3281/3281
Install Phase	3290/4450
...	
Install Phase	4450/4450
Update Phase	86/13668
...	
Update Phase	13668/13668
Reading Existing Index	1/9
...	
Reading Existing Index	9/9
Indexing Packages	1/549
...	
Indexing Packages	549/549

The upgrade, start to finish, took 21 minutes. A reboot brought the system up running the new version of OpenSolaris. It is simply amazing how much easier it is to perform administration tasks such as disk mirroring and upgrading than with standard Solaris. Certainly I prefer this platform to the old one.

The Future

As to the future, that remains to be seen. Without a doubt, OpenSolaris makes a better desktop operating system than Solaris. Also without a doubt, Solaris makes a better datacenter operating system at this point. The lack of SPARC support and application vendor support for OpenSolaris will keep it out of the datacenter for most uses.

What about Sun's official view on the future of Solaris versus OpenSolaris? Jim McHugh, Vice President of Solaris marketing, responded to my question with this quote:

Solaris is Sun's flagship operating system. OpenSolaris is the open source community and release cycle that is based on 20 years of Solaris development and innovation and where the next generation of Solaris is being built. Solaris and OpenSolaris releases, since they're actually from the same source foundation, are very similar but have somewhat different target audiences. The Solaris release hallmark is its long life cycle and mission-critical enterprise support. OpenSolaris releases come out every six months and thus have the latest and greatest features, making it ideal for developers, Web 2.0 companies, enthusiasts, and startups.

For users facing challenging business and technical requirements, such as lowering costs, simplifying system administration, and maintaining high service levels, the Solaris platform is the ideal choice. It is supported on over 1000 x86 and SPARC systems and its innovative features like ZFS, Dtrace and Containers deliver breakthrough virtualization capabilities, data management, advanced security, and world record performance.

Over time this may change, but it seems that we will have two major Solaris releases, much as Linux distributions frequently have a commercial release and a free release. Such is the new, confusing, exciting world of Solaris.

Next Time

Sun has a new NAS product line that should be released by the next column deadline. An exploration of the features, functionality, and differentiators of that set of products should make for fun writing (and, hopefully, fun reading).

RESOURCES

- [1] <http://www.usenix.org/publications/login/2008-04/pdfs/galvin.pdf>.
- [2] <http://blogs.sun.com/gbrunett/category/Solaris+10+Security>.
- [3] http://www.infoworld.com/article/08/09/24/39NF-linux-killing-solaris_1.html.
- [4] A talk about the state of OpenSolaris is found at <http://www.opensolaris.org/os/project/indiana/files/OpenSolarisTownHallv7.pdf>.
- [5] VirtualBox is available from <http://www.virtualbox.org>.