

> Adventures with illumos

Peter Tribble

Theoretical Astrophysicist

Sysadmin

(DBA)

Technology Tinkerer

> Introduction

- Long-time systems administrator
- Many years pointing out bugs in Solaris
- Invited onto beta programs
- Then the OpenSolaris project
- Voted onto OpenSolaris Governing Board
- Along came Oracle...
- illumos emerged from the ashes

> key strengths

- ZFS – reliable and easy to manage
- Dtrace – extreme observability
- Zones – lightweight virtualization
- Standards – pretty strict
- Compatibility – decades of heritage
- “Solarishness”

> Distributions

- Solaris 11 (OpenSolaris based)
- OpenIndiana – OpenSolaris
- OmniOS – server focus
- SmartOS – Joyent's cloud
- Delphix/Nexenta/+ – storage focus
- Tribblix – one of the small fry
- Quite a few others

> Solaris 11

- IPS packaging
- SPARC and x86
 - No 32-bit x86
 - No older SPARC (eg Vxxx or SunBlades)
- Unique/key features
 - Kernel Zones
 - Encrypted ZFS
 - VM2

> OpenIndiana

- Direct continuation of OpenSolaris
 - Warts and all
- IPS packaging
- X86 only (32 and 64 bit)
- General purpose
- JDS desktop
- Generally rather stale

> OmniOS

- X86 only
- IPS packaging
- Server focus
- Supported commercial offering
- Stable components can be out of date

> XStreamOS

- Modern variant of OpenIndiana
- X86 only
- IPS packaging
- Modern lightweight desktop options
- Extra applications
 - LibreOffice

> SmartOS

- Hypervisor, not general purpose
- 64-bit x86 only
- Basis of Joyent cloud
- No inbuilt packaging, pkgsrc for applications
- Added extra features
 - KVM guests
 - Lots of zone features
 - Overlay networks (vxlan)

> OpenSXCE

- Really does look like SXCE
- SVR4 packaging
 - With pkgutil
- SPARC and x86
- Erratic development and releases

> Dyson

- Debian with illumos kernel
 - Aka Debian/kOpenSolaris
- 64-bit x86 only
- <http://osdyson.org>

> DiLOS

- Illumos with dpkg/apt
- Xen dom0
- Including dpkg zones
- <http://www.dilos.org>

> Tribblix

- SPARC and x86
- Multipurpose
- SVR4 packaging
- Lightweight desktops
- Retro feel
- Modern components
- Emphasis on easy admin and zones

> Storage Vendors

- Nexenta
 - And OEM such as Dell's zNAS
- Delphix
- Tegile
- Coraid (???)
- Others...

> Lost without trace

- Belenix
 - Introduced the technology behind the live CD
 - Eventually folded
- Illumian
 - Website, iso, then nothing
- StormOS
 - Superseded by osdyson
- Korona
 - KDE atop OpenSolaris

> Decisions - SPARC

- SPARC is pretty limited
- Modern SPARC (T4 or later)
 - Oracle Solaris
- Coolthreads (T1/T2/T3)
 - Solaris 10/11, Tribblix, OpenSXCE
- Older SPARC (US-II/III/IV)
 - Solaris 10, Tribblix, OpenSXCE

> Decisions - x86

- 32-bit rules out Solaris 11, SmartOS
- Want KVM?
 - SmartOS, OmniOS, OpenIndiana
- Want modern desktop?
 - XstreamOS
- Want sparse-root zones?
 - Tribblix

> Decisions - packaging

- IPS
 - Solaris 11, OmniOS, XStreamOS, OI
- Debian
 - Dyson, DiIOS
- SVR4
 - Tribblix, OpenSXCE
- RPM
 - (Belenix)
- PKGSRC
 - SmartOS

> Decisions - environment

- Cloud/hypervisor
 - SmartOS
- Commercial apps
 - Solaris
- General server
 - OmniOS
- Desktop
 - XStreamOS, Tribblix

> Decisions - applications

Providing a significant variety of canned applications is not a characteristic of any Solaris or illumos variant.

(My personal preference has always been to manage any applications independently of the hosting OS.)

(And it's not clear that other OS platforms actually do any better in this regard.)

> Why build your own?

- I mean, weren't there enough distros?
- Because it's hard
- Understand the inner workings
- Satisfy the target audience
- Make a flexible platform for development of new ideas
- Didn't like other distros!

> Tribblix values

- Opinionated
- Modern components
- Retro styling
- Lightweight window managers
- SVR4 packaging
- Lightweight and fast
- Simplicity and “just works”

> Build - challenges

- Understand live boot
 - Most distros work the same way
- Create build and install from scratch
 - Here, Tribblix is completely different
 - Upgrade is yet another challenge
- Buggy environments (VBox, LDOMs)
 - So, who do I blame when it breaks?
- Fundamental design flaws
 - `driver_classes`, `driver_aliases`, `name_to_major`, `devlink.tab`

> Desktop variety

- Xfce default
- Enlightenment E17
- Afterstep, Windowmaker
- Icewm, openbox, fluxbox
- CDE
- Pekwm, fvwm
- Awm

> Zones

- Traditionally baked into packaging
 - Whole root, sparse root, ipkg
- That's the wrong way round
 - Modifications require changing packaging
 - And potentially all extant packages....
- In Tribblix, it's reversed
 - Zones understand packaging
 - Simple packaging helps here!

> Zone flavours

- Sparse root
 - Mount /usr from the global zone
- Whole root
 - Copy of the global zone
- Partial root
 - Just copy some of the software (maybe add more)
- Alternate root
 - Like sparse root, but with a different set of software
- Alien root (under development)
 - Run another distro in a zone
- And potentially more in the works...

> Tribblix Zones

- Do not fiddle with packages or packaging
 - Bad: hollow packages, inherit-pkg-dir, attributes in manifests
- Installing files is a pathname filter
 - Not present, private copy, mounted from, shared
 - Can use arbitrary sources of data
- Needs a registry of pathnames
 - SVR4 contents file
 - Trivial manipulation via shell commands
 - (Can do this with IPS as well, if you wanted)
- Allows easy creation of new zone flavours

> Building for Tribblix

- Minimal modifications
 - Only so it will build/run
- Simple configure/make/install/package
 - Prefix=/usr
 - Simple build script to reduce typing
- Try to minimise dependencies
 - Sometimes get locked downrev
 - Dependency hell beckons

> Managing packages

- Don't!
 - It's the wrong level of abstraction
- Tribblix has overlays (think clusters)
 - Useful unit of functionality
 - Example: xfce
 - Many fewer, easier to have a mental map
- Ignore package dependencies
 - Can be wrong or inappropriate
 - Overlays are tested as a unit

> ZAP!

- Zip Archive packaging
 - Packages are zipped SVR4 filesystem format
- Also the name of the command
 - Manage software (preferably overlays)
 - Never execute the old SVR4 commands
- Also manages zones
 - Because a zone is just a container for software

> Tribblix futures

- Zones and app deployment
 - Sparse-root, whole, partial, alien
 - Thinking like Docker and/or Rocket
- Simplify administration - “just works”
 - Make internals even more invisible
- Modern application stacks
 - Many on top of go
 - Integrated with zones and zfs (see first point above)

> illumos Directions

- ZFS (OpenZFS)
- XPG7 standards
- Missing pieces from open code
- Cleaning cruft
 - But can we preserve heritage?
- LX brand (native Linux emulation)
 - Hence Docker support

> Potholes

- Not enough time/people
- Fragmentation
 - All the work done at the distro level
- SPARC port struggling
- No cgo yet

> Further reading

<http://www.illumos.org/>

<http://www.tribblix.org/>

<http://www.petertribble.co.uk/>

<http://operatingsystems.io/>

Peter.tribble@gmail.com