

Sine Nomine Associates

Debian/390: So What's the Big Deal?

David Boyes
Sine Nomine Associates
WAVV Colorado Springs 2005

Presentation Download

This presentation will be available for download from:

www.sinenomine.net

What is Debian?

- Debian is a aggressively open-source Linux distribution for a large number of architectures
 - Debian 3.1 “Sarge” and 3.0 “Woody” both support: alpha, arm, hppa, i386, ia64, m68k, mips, mipsel, powerpc, sparc, s390
- Debian is a very large distribution:
 - 3.0: 8465 packages (s390, 5 May 2005)
 - 3.1: 15115 packages (s390, 5 May 2005)
 - (totals include contrib and non-free packages)

What Else is Debian?

- Unlike Red Hat or SuSE, the Debian Foundation is a non-profit organization
 - Debian Community large and active
- Debian Social Contract: http://www.debian.org/social_contract
- Very, very big on policy
 - Ensures a lot of consistency
 - Also makes change slow
- Names releases after “Toy Story” characters.

Why Debian?

- More packages than any other Linux distribution
 - Admittedly, many are useless for S/390 users
 - Most are probably useless for most users...
- Very clear demarcation between Free Software and Non-Free (as determined by license)
 - No dependency on proprietary tools
- Maintenance and security fixes are much, much easier than SuSE or RedHat (if you have a network connection)

Why Debian on S/390?

- No pay-to-play
 - You don't need an expensive support contract to get basic security fixes
 - You can buy support if you like – no lock-in to single vendor

- Much easier to create lightweight Linux instance
 - “Minimal” SuSE needs 2 3390-3s these days
 - Debian runs quite well in 250 cylinders

Why do we like it?

- Staying up to date is trivial:

`'apt-get update && apt-get dist-upgrade'`
- Software distribution and packaging model is better suited to appliance (R/O) deployment
- Very friendly for CLI-only servers:

administer everything via ssh and screen installations.

Why do we like it?

- Easy to do much more modular systems than with SuSE
 - SuSE loves to do everything-and-the-kitchen-sink versions of apps, which pull in huge numbers of barely-related libraries
 - Debian typically splits complex apps into multiple packages

Example: for the Bacula application, you have bacula-director-common coupled with one of bacula-director-mysql, -pgsql, or -sqlite, depending on which back end you prefer

- Configuration is separate from code in most packages.
- This lets us do little appliance machines much more easily

Why do we like it?

- Small dedicated-function appliance machines are what we think Linux on the mainframe should be:
 - SSLSERV is the best known
 - SMTPPLUS for modern mail handling under VM
 - Bacula with DFSMSHsm or VM tape-mount back ends for zero-additional-cost file-level backup to the same tape silos
 - Traditional Samba file/print services
 - Experimental work, like the iptables-based firewalling IP stack enhancement project

Little Appliance Machines

- We like the appliance approach for a number of reasons
 - Separation of function between virtual machines is the VM way of doing things: one service per virtual machine.
 - We can package them as black boxes: you don't have to care how it works. Run the installer, and start using the service.
 - We can sell and support them cheaply, individually or bundled; separation of function makes it easy to refactor them for different customer needs
 - No one has to pay SuSE's support fee in order for us to support them for our customers; while SuSE's charges may be reasonable for general-purpose Linux servers, they're high for single-purpose black-box virtual machines

Some Annoying Debian “Features”

- Lots of politics in the development community
 - Everyone’s a volunteer, so the Debian Project Leader’s power to enforce change is limited
 - A fair number of prima donnas and hissy fits
- This leads to slow development times: Sarge is very late (but coming soon! Really!)
- No complete zSeries (64-bit) port yet.
 - Kernel work completed; matter of doing a full build and regression test
 - We have an OSDL machine reserved for it; after Sarge release we’ll take a look at effort involved.

Recent Rumors

- Are just that: rumors
- 390 remains solidly on the supported platform list for Debian
- 390 is not one of the “problem” platforms
 - Sarge for 390 has been ready for months
 - Significant work on debian-installer came from 390 folks
- We have all the source. We can rebuild it... 8-)
- SNA and others remain committed to a low-cost, up to date Debian platform on 390 and 390x

Debian-installer and Sarge

- Please try Debian-installer and Sarge (RC3 is the latest)
 - We're using Sarge in production lots of places now; it's as stable as Woody is
- Debian-installer much easier to use than Woody installer was
- Get it from <http://www.debian.org/devel/debian-installer/>

Differences in Debian

- Configuration files differ in some places
 - /etc/network/interfaces instead of /etc/sysconfig/network
 - Runlevel 2 is default, not 3

- Less commercial application support
 - We're working with vendors to change this
 - Debian gaining traction in x86 world

Differences in Debian

- Configuration philosophy a bit different: designed for easy manipulation by programs, but all text files so easy for humans too
- Application defaults under `/etc/default`, sort of like `/etc/sysconfig` in SuSE/RH
 - `/etc/sysconfig` is supported, but is not part of the formal Debian policy
 - Détente in progress between the `/etc/default` and `/etc/sysconfig` camps on which is “preferred”
- Policy ensures man pages for everything (none of this “info” nonsense!)

Differences in Debian

- Default application choice sometimes different
 - Usually more aggressive about adopting new tech

Example:

- Bind 9 as nameserver
 - Exim 4 as MTA
- Less automated X and Desktop configuration
 - We've never had a problem with it, though

Differences in Debian

- Default to .deb package format rather than RPM
 - APT does better dependency tracking than RPM
 - APT combines both package management and delivery
 - More sophisticated package architecture policy
 - dpkg/apt are your package-maintenance tools, not rpm
 - RPM conversion tools allow coexistence

- Much less forcibly-GUI-oriented
 - In general, GUI tools bolt on top of command-line interfaces and are not required
 - Makes it easier to programmatically control service behaviour and inject new service hooks
 - Means you can install/run with much less machine

Porting Applications to Debian

- Toolchains are essentially the same as all other Linux distributions
- Build processes are essentially the same
- Only major differences are the ones noted earlier – locations for config files and runlevels
- Unofficially, binaries compiled for other distributions generally work without modifications – but test, test, test!

“Officially” Porting Applications to Debian

- If you have the source, it's fairly easy:
 - Good tutorial on creating control files, etc. at <http://www-106.ibm.com/developerworks/linux/library/l-debpkg.html>
 - The trick is building a package that's well-behaved by Debian's stringent criteria
 - The Debian Policy MATTERS. RTFM.
- If you don't have all the source code, then things get a little tougher...

Case Study: SSLSERV

- The core SSLSERV interface code exists as IBM-packaged code on 4TCP40 493 (for z/VM 4.4), or VMSYS:4TCP40.BINARY.
- It's based on the AIX GSKit for its crypto routines
 - GSKit is binary-only
 - Very old GSKit, linked against an ancient libstdc++
 - Doesn't use HW Crypto

Getting all the pieces together

- Install the alien package to deal with RPMs...
 - alien -t to create a tarball
 - Everything is in vmssld.tgz
- Build fails: you need libstdc++-libc6.1-2.so.3, which Debian has never heard of
 - rpmfind.net to the rescue: found it in the compat libraries in Fedora
 - Ran alien -t to unpack, removed everything but the libraries I needed, ran alien -d to repack into .deb, and installed it.

Why not a proper libcompat?

- Libcompat is GPL, so I shouldn't have rebuilt it from source, right?
- To build the environment to build libcompat requires building an ancient GCC (2.95.3) with 433 distinct patches applied to the GCC source tree.

Building SSLSERV

- With the correct libraries and kernel headers, the build went fine and I had a working vmssl binary
- Needed to configure the system: rip out existing TCP/IP definitions
 - rc-update.d service remove
 - Get rid of /etc/network/interfaces

Fixing up startup

- SSLSERV startup scripts clearly written by VM/Rexx programmers, not Linux/sh people
 - = versus == was the giveaway
 - Some of it didn't do what they thought it did
- Scripts written for early versions of SuSE and RH
 - Danger Will Robinson!
These scripts should NOT be used as models!
- Backgrounding was hideous: ended up writing a wrapper to daemonize vmssl.
- Added some tests to determine if running on Debian and do things the Debian way if so
 - start-stop-daemon, not startproc
 - Logfile location

Why no SSLSERV .deb ?

- What's the point? Once SSLSERV is installed, you have no networking and you don't have a general-purpose Linux guest anymore
- This is just a drop-in virtual machine: the whole thing functions as a black box. Without networking, apt-get is useless and the packaging system becomes irrelevant
- Maintenance shouldn't be your problem for an appliance machine: it should be ours

SSLSERV Wrapup

- The fact that it's Debian under the hood is irrelevant
 - Except that it means it can be supported cheaply
 - And we can maintain the build environment easily
- To you, the user, it's a black box: restore the DDR image, configure your TCP/IP stack, and away you go.
- We feel other services (like SMTPPLUS) should be packaged the same way.

Debian Distribution Availability

- Download from www.debian.org
 - On your own to do the integration
- Commercial CD packages and support available for most architectures from many sources
 - Buy your support from your favorite source – no lock-in extortion
 - AFAIK, SNA only source for commercial Debian/390 support

SNA's Debian Offering

- Native CMS-based install from tape to quick start install server
- Installation server provides CD images as pre-configured Debian network install source
 - Debian is happiest installing over the network, but...
 - WAN bandwidth may be limited
 - Your S/390 or zSeries may not be able to see the outside world
- Client systems install quickly within the machine
 - 8-10 minutes from 1st boot to ready-to-use.
- 24x7 OS and Application Support

SNA's Debian Offering

- Why do it this way?
 - We install a LOT of virtual machines
 - WAN bandwidth isn't infinite
 - Installs over guest LAN/hipersocket are REALLY fast
- Current install server based on Debian 3.0r4
 - Client systems can be ANY release
 - Also supports Sarge (3.1) images if you have disk space

Forthcoming Debian Offering

- 3.1 “Sarge”-based
 - Once “Sarge” becomes Debian “stable”
 - We use “Sarge” a lot, but are not going to try to support a prerelease distro, except in the context of a “black box” appliance
- Going to use Debian pool filesystem tree
 - Will require LVM on installation server
- Will be based on new Debian-installer
 - Preseeding makes configuration much simpler
 - Just have to update the parmfile cards

Debian Wrapup

- Debian is a means to an end, not an end in itself
 - A modular, flexible, extensible, easily supported, low-cost general-purpose development system
 - A base platform for embedded single-function appliances
 - A path to a design where each appliance is building block in overall modular service architecture

Contact Info

David Boyes
Sine Nomine Associates
dboyes@sinenomine.net
+1 703 723 6673