Using Debian to provide services to thousands

Faidon Liambotis
paravoid@debian.org

DebConf10
August 3rd, 2010

Who am I?

- Debian Developer,
  paravoid@debian.org
- Servers & Services Engineer,
  Network Operations Center,
  Greek Research and Technology Network

grnet
Networking Research and Education

What I'll talk about

Services & Infrastructure

Requirements

Setup

Future

Challenges

Contributing back

grnet
Networking Research and Education

Greek Research and Technology Network

- aka GRNET
- The Greek “NREN”
- Public company providing Internet & computing services to the academic, research and educational community
- Mostly big customers (universities) but occassionally providing services directly to users (e.g. ADSL)
- Part of the pan-European GÉANT network
- Operates GR-IX, the Greek Internet Exchange

grnet
Networking Research and Education

Network Operations Center

- Day-to-day operation of network and services
- Also design/capacity planning
- Newly-founded (2009)
- Predated by an outsourced operations model (“VNOC”)
- 5 netadmins, 5 sysadmins, 1 developer,
  1 operations manager

grnet
Networking Research and Education

Services in general (1/2)

- Two main groups of services
- Internet connectivity services
  - DWDM, Ethernet, IP etc.
  - Mostly being done by proprietary equipment (Alcatel, Cisco, Juniper)
  - Not much you can do about that
- System/computing services
  - “Server stuff”
  - Will focus on these (I'm a sysadmin!)

grnet
Networking Research and Education

System/computing services (2/2)

- Network-supporting (DNS, monitoring, authentication)
- Identity (LDAP, RADIUS, SAML federation)
- PKI
- Storage *(Dropbox-like)*
- VVoIP & video-conferencing
- Video on demand/IPTV
- …and other minor stuff (mail, web, NNTP, …)
- VPS, Colocation, Dedicated servers

grnet
Networking Research and Education

Infrastructure

- Two privately-owned data centers (around 50 racks in total)
- Colocating racks in *hundreds* of facilities across Greece
- Over 200 physical servers
- …and several hundred more virtual servers
- Two SANs (NetApp & EMC) with over 500TB of storage

grnet
Networking Research and Education

Past

- < 2009
- Design & operations outsourced to multiple mostly-disjoint teams
- A heterogeneous multi-vendor environment
- CentOS, Scientific Linux, Debian, SUSE, Solaris (8, 9, 10), FreeBSD, HP-UX, …
- Software (and updates) typically installed from source
- Monolithic servers, each running tens of unrelated services

grnet
Networking Research and Education

Transition to an in-house NOC

- From dozens of admins to a handful
- *Hundreds* of machines
- …with more coming (a need to scale)
- Decision to redesign from scratch
- Three basic principles:
  - Virtualization
  - Automation
  - FL/OSS & Debian

grnet
Networking Research and Education

Virtualization

- Taking advantage of x86 virtualization
- Running a separate VM per core service/role
- Also handing out VMs to customers (VPS)
- Initial pilot on newly-acquired equipment:
  - HP BladeSystem c7000
  - 12 blades with 16/32GB RAM
  - 4Gbps network on each server
  - FibreChannel to an IBM/NetApp SAN
  - 168 × 1TB SATA + 56 × 300GB FC disks
- In production for over a year

grnet
Networking Research and Education

Automation

- Too many machines, too few people
- Virtualization $\rightarrow$ more systems
- Not a luxury anymore but a real necessity
- Centralized provisioning
  - backup
  - monitoring
  - admininistration tools
  - etc.
- Centralized auditing

grnet
Networking Research and Education

Why FL/OSS & Debian (1/3)

- Team *and* management both believe in Free Software
- A convenience
  - Usually of better quality
  - Easier (& cheaper) to work with
  - Fixing problems ourselves
- …and belief in the Free Software values
  - Long-time supporters of FL/OSS in Greece
- Price is also a factor…
- Being in the research & technology section helps

grnet
Networking Research and Education

Why FL/OSS & Debian (2/3)

- Conscious choice to switch **everything** to Debian
- We were all familiar with it
- “The Universal Operating System”
- Huge collection of software
- Well-integrated system
  - Things are where you expect them to be (`/opt/csw` anyone?)
  - Services are well integrated with each other

grnet
Networking Research and Education

Why FL/OSS & Debian (3/3)

- `stable` means rockin' stable!
- Convenient release cycle
  - Yes, it is!
  - Not too frequent, not too infrequent
- Security support
  - For everything in the archive, not just the core system
  - For sufficient time after the current+1 release
  - On-time, usually without breakages

grnet
Networking Research and Education

Virtualization solutions (1/2)

- First quick-n-dirty attempt:
- `lenny`'s Xen hypervisor/kernel + libvirt + custom tools

grnet
Networking Research and Education

Virtualization solutions (1/2)

- First quick-n-dirty attempt:
- `lenny`'s Xen hypervisor/kernel + libvirt + custom tools
- FAIL:
  - Xen is buggy
  - Xen is hard to operate
  - Live migrations never worked properly
  - Debian is deprecating Xen
    (along with the rest of the world)
  - libvirt is too complex

grnet
Networking Research and Education

Virtualization solutions (2/2)

- Second attempt: Google's Ganeti + KVM

grnet
Networking Research and Education

Virtualization solutions (2/2)

- Second attempt: Google's Ganeti + KVM
- A huge success so far!
- Benefits:
  - Cluster management!
  - Fast-moving/adapting communities (both KVM & Ganeti)
  - Working live migrations, KSM, booting over HTTP, …

grnet
Networking Research and Education

Virtualization solutions (2/2)

- Second attempt: Google's Ganeti + KVM
- A huge success so far!
- Benefits:
  - Cluster management!
  - Fast-moving/adapting communities (both KVM & Ganeti)
  - Working live migrations, KSM, booting over HTTP, …
- But:
  - Ganeti didn't have shared storage/SAN support (we fixed it and are contributing it back)
  - KVM has bugs of its own (we worked around them)

grnet
Networking Research and Education

Automation tools (1/2)

- Bootstrapping:
  - `d-i` & preseeding
  - `debirf` for rescue

grnet
Networking Research and Education

Automation tools (1/2)

- Bootstrapping:
  - `d-i` & preseeding
  - `debirf` for rescue
- Monitoring:
  - Alerts via Nagios
  - Metrics via `munin`
  - …but evaluating `collectd`
  - Centralized `rsyslog`

grnet
Networking Research and Education

Automation tools (1/2)

- Bootstrapping:
  - `d-i` & preseeding
  - `debirf` for rescue
- Monitoring:
  - Alerts via Nagios
  - Metrics via `munin`
  - …but evaluating `collectd`
  - Centralized `rsyslog`
- Backups:
  - Bacula
  - On-site & off-site file-level backups

grnet
Networking Research and Education

Automation tools (2/2)

- Provisioning/management:
  - Puppet
  - Class-based provisioning
  - Integrated with the OS, not bypassing it
  - Integrated with all of the above

grnet
Networking Research and Education

Automation tools (2/2)

- Provisioning/management:
  - Puppet
  - Class-based provisioning
  - Integrated with the OS, not bypassing it
  - Integrated with all of the above
- Data mining & hardware inventory
  - `facter` (from Puppet)
  - + custom facts (like `apt` updates)
  - + a (Django-based) internal web application

grnet
Networking Research and Education

Using Debian (1/2): base system

- Almost exclusively `lenny`
- backports.org
- internal backports
- …and a few custom packages
- Nothing too out of the ordinary!
- Few unfortunate exceptions (multimedia etc.)

grnet
Networking Research and Education

Using Debian (2/2): software used

Exim Ganeti Debian Drupal proftpd IRRToolset rsyslog KVM Subversion Quagga Cacti Collectd vsftpd Sendmail Asterisk MySQL Rails Linux HAProxy ejabberd GnuGK pmacct dbeacon PowerDNS sflowtools Bacula OpenLDAP Nagios Open-iSCSI FreeRADIUS Openfire lldpd OpenSIPS Sympa OpenVZ inn dvblast RadSecproxy munin Ruby Xen Puppet nginx Perl JBoss OpenSSL Tomcat Pure-FTPd OpenSSH rancid mrtg Apache Mercurial Django Shibboleth vlc Dovecot OpenVPN Mailman BIND redmine Postfix git SpamAssassin mongrel ffmpeg Python PostgreSQL Samba …

grnet
Networking Research and Education

What's coming

- Moving to “cloud” services (IaaS)
- Project to provide VMs to the community
- New datacenter (2010), pilot with 128 servers
- To be filled (2011) with around **1000** servers
- Evaluation of software on-going
- Aiming to run Debian on all of them!

grnet
Networking Research and Education

Challenges (1/5): Debian's competition

- Main “competitors” (for our use cases) are:
  - Microsoft (Windows, Hyper-V)
  - Red Hat (RHEL, RHEV)
  - Novell (SLES)
  - Canonical (Ubuntu)
  - VMware (ESXi)
  - Citrix (XenServer)

grnet
Networking Research and Education

Challenges (2/5): hardware vendors

- Debian not very well supported by the “enterprise” world
- Server vendors usually ship (nasty) RPMs
- Storage vendors (NetApp, EMC) integrate *really* well with (at least) VMware
- No hardware specifications mention Debian (and we procure via public tenders…)
- There are a few exceptions (mostly HP, thanks!)

grnet
Networking Research and Education

Challenges (3/5): software support

- Debian is generally known but not very well accepted
- No single vendor who can support you
- *Some* independent contractors who offer support
- …but you might just as well support it yourself
- Usually hard to convince management about that
- On the other hand: no contract ⇒ reduced cost

grnet
Networking Research and Education

Challenges (4/5): upstream troubles

- Debian rarely the upstream
- Xen is an non-mainline patch
  - really hard to support properly
- KVM follows the RHEL/RHEV release cycle
  - fast-changing
  - $\Rightarrow$ frequently broken
  - no real stable branch

grnet
Networking Research and Education

Challenges (5/5): doing uncommon things

- The difference between a `grave` and an `important` bug
- Hard for maintainers to test big or exotic setups
- Not a lot of people do *that*:

```
ls -la /dev/sd* | wc -l
 3664
```

- Or live migrations with shared storage
- Or iSNS

grnet
Networking Research and Education

Contributing back (1/2)

- Being good Free Software citizens
  - Reporting bugs (properly)
  - Providing patches/doing upstream work
  - Maintain packages in Debian :)
- Donating/sponsoring resources
  - Currently colocating 15Us of Debian (DSA) equipment!

grnet
Networking Research and Education

Contributing back (2/2)

- Releasing our code under FL/OSS licenses
  - There's `code.grnet.gr`
  - …plus other contributions
- Encouraging other organizations to use Debian
  - Academic community (both inside and outside of Greece)
  - Goverment organizations

grnet
Networking Research and Education

Questions?

grnet
Networking Research and Education