


Linux: An Introduction



© Anand Vaidya
vaidya.anand@gmail.com
Version: 2.4
Last Updated: 09-Jan-2009

Topics

- ◆ What is Linux?
- ◆ UNIX in 2 slides
- ◆ Why use Linux?
- ◆ Why Not Use Linux? Concerns.
- ◆ Who uses Linux & Where?

- ◆ Hands on : Demo
- ◆ Hands on : Commands

What is GNU/Linux?

- ◆ Just another OS like Windows XP or Mac OSX or Unix
- ◆ Provides an environment for users to run applications
- ◆ Distinctive features:
 - ◆ User friendly License terms
 - ◆ Committed volunteer and employed contributors
 - ◆ Wide adoption in corporate, education, government and military
 - ◆ Use for server, desktop, embedded
 - ◆ Usable in 80+ languages around the world

What is GNU/Linux?

Comprises of:

- ◆ Monolithic Kernel by Linus Torvalds & many contributors (since 1991)
- ◆ Various GNU, BSD & other software
- ◆ Started as a hobby/learning experiment since Linus found DOS inadequate and UNIX expensive for his i386 machine
- ◆ Modeled after free educational Minix OS
- ◆ Follows UNIX philosophy
- ◆ Source code freely available
- ◆ Different portions of code on a complete OS use different licenses:
- ◆ Eg: GPL , LGPL, BSD, MPL



Licensing & Rights

- ◆ Entire OS source code is free (Gratis) – free as in zero dollars
- ◆ Also Free - Full freedom to study, modify, redistribute. No payments or restrictions. (Libre') – Free as in Freedom
- ◆ Kernel – GPL licensed, many other packages follow GNU GPL, LGPL or BSD/MPL style license.
- ◆ GPL designed to protect end users & programmers rights.
- ◆ Copyrighted to the author
- ◆ NOT public domain
- ◆ Open Source (<http://www.opensource.org>)
- ◆ Know more about (L)GPL: <http://www.fsf.org>

Why Linux?

- ◆ Multitasking, SMP, NUMA, HA Clustering
- ◆ Protected virtual memory
- ◆ True multiuser capabilities
- ◆ Robust journalled filesystem choices: ext3, ext4, JFS, XFS, reiserFS, Btrfs
- ◆ Highly modular & scalable: From embedded systems to IBM mainframes (PDA/watch/mobile phone, SMP servers to 1024 CPU SGI boxes and IBM mainframes)
- ◆ Multiple hardware architecture: eg: Debian Linux is supported on 11 processor architectures (eg PPC32/64, IA32/64, ARM, SPARC, x86_64, SuperH etc)

Why Linux?

- ◆ Scales well SMP: 1000+CPU Itanium2 SGI Altix3000
- ◆ Clusters 4000+ CPU Itanium2 (PNNL), 10000 CPU clusters are in production
- ◆ Support for a wide array of hardware (SCSI, USB, scanners, printers, VGA, RAID, wireless, network hardware, FPGA, InfiniBand, 1/10GE, Myrinet)
- ◆ X11 windowing system for Graphics
- ◆ Various Desktop Environments: KDE, GNOME, Xfce, Windowmaker, Blackbox etc
- ◆ Powerful command line interface(shells)

Why Linux?

- ◆ Scripting:
 - ◆ Shell (bash, csh, zsh, ksh)
 - ◆ Perl, PHP, Python, Ruby, tcl/tk, lisp....
- ◆ Several compiled languages:
 - ◆ C, C++, pascal
 - ◆ Fortran (GNU f77, f90/f95, intel/ pathscale/ pgi f90)
 - ◆ Java
 - ◆ C# (Mono)
 - ◆ Smalltalk, ada, haskell
- ◆ Rich Development Environment (tools, IDE, code samples, community) kdevelop, anjuta, Eclipse, Netbeans, IntelliJ, Eric etc

Why Linux?

- ◆ Server Applications:
 - ◆ Web (Apache, lighttpd)
 - ◆ DNS (BIND, djbdns)
 - ◆ Mail (sendmail, postfix, qmail, pop3, imap)
 - ◆ LDAP (OpenLDAP)
 - ◆ Database (MySQL, Postgres, Oracle, DB2)
 - ◆ Fileserving (NFS, samba, appletalk)
 - ◆ Corporate messaging: Lotus Domino, Bynari Insight, Zimbra, openXchange
 - ◆ IM servers
 - ◆ Source code mgmt: CVS, SVN, DARCS, git
 - ◆ RADIUS
 - ◆ Too many to list...

Why Linux?

- ◆ Scientific Apps: (Bio)
 - ◆ BLAST, bioperl, biopython, biojava...
 - ◆ Accelrys DiscoveryStudio, Materials Studio
 - ◆ IBM DiscoveryLink
 - ◆ CNS, HMMer, Phrep, Phylip, VMD
 - ◆ Over 200 opensource BioInformatics apps listed at www.bioinformatics.org

Why Linux?

- ◆ Scientific Apps: (Math)
 - ◆ Matlab, Mathematica, GNU Octave, R
 - ◆ NAMD, AMBER, Gaussian
 - ◆ Fluid dynamics: Fluent, ANSYS CFX
 - ◆ Material Science: Materials Studio, CNX
 - ◆ Animation, image processing: PRMan, 3dsmax backburner, maya, blender, gimp, filmgimp
- ◆ Weather: MM5
- ◆ More...SourceForge lists 100K+ projects (2008)

Why Linux?

- ◆ Virtualization and Grid computing:
 - ◆ Linux runs on IBM pSeries, Mainframes (zSeries hypervisor)
 - ◆ Xen VMM from Cambridge Univ
 - ◆ VMWare – WS, GSX, ESX Servers
 - ◆ Parallels VM
 - ◆ Kernel Virtualization – KVM
 - ◆ Other virtualization solutions: OpenVZ, User-Mode Linux, vserver etc
 - ◆ Intel VT and AMD-V hardware virt assists
- ◆ Linux is extremely popular in the Grid computing market

Networking

- ◆ Born and brought up on the Internet
- ◆ Very strong networking support: IPv4, Ipv6
- ◆ Range of networking apps, drivers
 - ◆ VPN
 - ◆ Firewall (packet-filtering, mangling)
 - ◆ Routing, bridging, QoS, VLAN tagging
 - ◆ Wireless
 - ◆ ADSL, SDSL, ATM, SONET, Ethernet, iSCSI
- ◆ High Performance Interconnects: 10Gb Ethernet, Quadrics, Myrinet, Infiniband etc
- ◆ Large set of network hardware drivers and protocols

Secure By Design

- ◆ Designed groundup for multiuser, multitasking
- ◆ Strict separation between user and kernel
- ◆ Chroot, jails, Sane defaults

- ◆ Enhancements:
 - ◆ NSA SE linux, RSBAC, file ACLs
 - ◆ GrSecurity, Openwall, LIDS patches
 - ◆ Non exec stack, canary markers (eg Adamantix Debian). No Execute (x86_64)

- ◆ Advanced package management:
 - ◆ just install necessary software, no more.
 - ◆ run only necessary services
- ◆ Extensive logging: almost every activity is logged
- ◆ Easily securable, auditable

CPU support

- ◆ Supports many CPU architectures:
 - ◆ x86 Intel IA32 and clones (AMD, Transmeta, Cyrix, VIA etc)
 - ◆ IA64 Intel 64bit Itanium2
 - ◆ AMD64 (Opteron and Athlon64) and Intel EM64T
 - ◆ IBM PowerPC 32/64 bit
 - ◆ PPC970 (Macs), IBM POWER64
 - ◆ Sun Sparc64, Alpha64, MIPS
 - ◆ ARM, Hitachi, SuperH, Motorola Coldfire, Intel XS-cale and many more embedded processors

Standards

- ◆ Linux and applications adhere closely to various standards and Linux has some of its own:
 - ◆ POSIX, ANSI C, various IEEE, ANSI, W3C
 - ◆ IETF RFCs (eg mail, web, routing, dns)
 - ◆ FHS, LSB, freestandards.org
 - ◆ CORBA, SQL...
 - ◆ Emblix (embedded linux -japan), CELF, DataCentre Linux
 - ◆ LI18N – Internationalisation

Distributions

- ◆ Distributions:
- ◆ =kernel+applications+branding
 - ◆ Select, compile and package applications
 - ◆ Test (sometimes thoroughly)
 - ◆ Provide Installer (text/GUI)
 - ◆ Technical support: phone, email, web
- ◆ Differences in distros: packaging, support, price, supported architectures, target customers

Distributions

- ◆ Popular distributions:
 - ◆ Debian (related: Knoppix, (k)ubuntu, Mepis, Xandros)
 - ◆ RedHat Enterprise (related: Fedora, CentOS)
 - ◆ SUSE Enterprise (related: OpenSUSE)
 - ◆ Mandriva
- ◆ Slackware
- ◆ Gentoo
- ◆ Embedded Linuxes
- ◆ Other specialised distros:
 - ◆ single floppy, NPACI ROCKS, BioBrew, DSL (damn small linux)

Distros-@Home?

- ◆ My Recommendations:
- ◆ Home usage: (user-friendly, mostly GUI, Multimedia rich)
 - ◆ Kubuntu (<http://www.kubuntu.org>) default KDE desktop or Ubuntu with GNOME
 - ◆ Other choices: OpenSUSE, Fedora
- ◆ Why ubuntu?
 - ◆ Easy to use
 - ◆ Very large & good community
 - ◆ free support
 - ◆ Optional paid support
 - ◆ Patches for atleast 18months
 - ◆ 1000's of apps. One click install. No compiling, no messing with dependencies

Distros- @Work?

- ◆ Commercial apps (Oracle etc): use recommendation by vendors (usually RedHat EL or SUSE EL)
- ◆ Debian/Ubuntu (maybe FreeBSD) for all other fully OSS apps: DNS, Mail, Firewall, NFS/samba, web, ftp servers, mysql, postgres etc...
- ◆ Power user desktops: Let them choose! (Gentoo, Debian/Ubuntu, SUSE, Redhat etc)
- ◆ Use RHEL or SLES if you need certifications (eg: EAL4)
- ◆ Low IT skilled workers: Thin clients (LTSP)

Distros-Which One?


- ◆ Thin Clients (eg: Libraries, Kiosks etc)
- ◆ Emergency Rescue, Security work:
 - ◆ Choose from one of the several single cd specialised distro. Watch <http://freshmeat.net> for announcements or searching.
- ◆ Don't forget the free BSD UNIXes viz. FreeBSD, OpenBSD and netBSD. Each has significant advantages.

UNIX in 2 Slides

- ◆ Been around for 30+ years, Elegant, simple design
- ◆ Dennis Ritchie, Ken Thompson, AT&T
- ◆ Extensively studied, commented
- ◆ and documented

Many commercial implementations:

- ◆ SUN Solaris
- ◆ IBM AIX
- ◆ HP UX
- ◆ SGI IRIX
- ◆ Digital UNIX (Tru64)
- ◆ Many other flavours... ever heard of unicos?




Thompson and Ritchie, the creators of UNIX, working on a PDP-11 machine.

UNIX in 2 Slides

UNIX Philosophy

- ◆ (almost) Everything is a file
- ◆ Build small programs: One well defined and do it very well
- ◆ All programs
 - ◆ accept Input from stdin,
 - ◆ write Output to stdout and
 - ◆ write Error to stderr
- ◆ Pipes connect program I/O like plumbing, data flows like water
- ◆ Strict separation between privileged (kernel) and user modes.



Thompson and Ritchie, the creators of UNIX, working on a PDP-11 machine.

Concerns with Linux

Beware:

- ◆ Many concerns expressed: some genuine, some bogus
- ◆ Don't trust whitepapers and journalists , analyst firms too much
- ◆ Many 'studies' on Linux funded by competitors and use meaningless data to provide wrong conclusions
- ◆ Search the internet for such less than honest behaviour by companies frightened by Linux's rapid ascendancy

Concerns with Linux

- ◆ Hard to install/use - Situation improving
- ◆ Too many distros, incompatible packages-being addressed by standards-LSB, FHS
- ◆ RPM hell – solved by yum, apt, urpmi etc.
- ◆ Non-availability of popular desktop applications (eg: MS Office, photoshop, Quicken, Games)
- ◆ Hardware - driver availability eg: USB-ADSL modems, winmodems, WinPrinters, VGA adapters etc. (esp. for no-brand/cheap hw)
- ◆ Viability of linux vendors & business Model (– Not a major concern nowadays)

Concerns

- ◆ Support :
 - ◆ Home / small user: *Who will help me?*
 - ◆ Enterprise: *Who will take the blame when there is a problem?*

Concerns

- ◆ Immature - compared to UNIXes
- ◆ SMP Scalability not great - compared to UNIX
- ◆ Fear of Code Forking (hasn't happened in 10+ years, probably won't)
- ◆ Legal issues:
 - ◆ GPL not tested in US courts, lawyers have questions regarding strength, though situation improving... (eg several German cases, IBM vs SCO)
 - ◆ Unknown/hidden Patents, claims by patent trolls
 - ◆ Copyrights violations – so far clean
- ◆ Microsoft's vehement, no holds barred effort to destroy Linux & GPL.

Usage: Who & Where

- ◆ Substantial section of internet powered by Linux
 - ◆ Amazon
 - ◆ Google
 - ◆ Ebay
 - ◆ akamai

Usage: Who & Where

- ◆ Research / Education:
 - ◆ Clusters: NASA, LLNL, LBNL, CERN etc
 - ◆ LANL 2048 CPU cluster 9.2TF
 - ◆ OSU, MIT, UCB, UCSD...
- ◆ Grid: US Teragrid runs on Linux, Butterfly.net (Gaming on grid)
- ◆ IBM BlueGene (64000 processor, petaflop capability) to run Linux
- ◆ Indiana State (USA)– 22000 students on Linux desktops - Aug2006

Who & Where

Telco/HSP:

- ◆ Telstra (Aussie telecom)
- ◆ Many ISPs for web/mail/ftp (slicehost, usonyx, rackspace etc)

Airlines:

- ◆ Korean Air - flight bookings

Financial:

- ◆ Wall St: Schwab, Fidelity
- ◆ Central Bank of India 1100 branches. LIC of India -4000+ branches
- ◆ First National Bank (SA) – 12000 desktops

Petroleum:

- ◆ Shell (& many others), Veritas

Who & Where

Petroleum:

- ◆ Shell (& many others), Veritas

Automotive:

- ◆ Crash simulations (Ford, Merc, Toyota)

Movie animation industry:

- ◆ Pixar, Weta Digital (Nemo, Titanic, Shrek ...)

Manufacturing:

- ◆ Burlington's coat factory (POS)
- ◆ Boeing, Nokia, Sony, Matsushita, NEC, Fujitsu, Motorola
- ◆ Small and Medium Enterprises: Too numerous to list

Users: Who & Where?

- ◆ Desktop:
 - ◆ US DoD uses Linux extensively(see Mitre report)
 - ◆ Extremadura, Spain 80K desktops in schools
 - ◆ Munich, Germany – entire city-14000 desktops
 - ◆ Vienna, Austria -entire city
 - ◆ Enthusiasts – atleast 20 million?

Users: Who & Where?

- ◆ IT cos using linux on desktops:
 - ◆ Oracle (RAC, ~ 15000 desktops)
 - ◆ IBM has about 14000 linux desktop users (2003 – Desktop Linux Conference) Planning to go 50-60K)
 - ◆ Novell

Users: Who & Where (SG)?

- ◆ NIE, RP, SP, NYP , NUS , NTU, SMU, SIM
- ◆ Research: A-STAR (ICES, BII, IMCB, DSI, GIS, I2R, NGP)
- ◆ National Library Board
- ◆ Defence: Mindef deployed OpenOffice on Windows (no word on linux yet)

FSF & GNU

- ◆ FreeSoftware Foundation founded by Richard Stallman
- ◆ Sister orgs in EU, LatAm and India
- ◆ Started gcc, emacs and GNU project
- ◆ Wrote GPL, LGPL licenses
- ◆ GPL = GNU General Public License
- ◆ GNU = GNU's Not Unix (recursive acronym)
- ◆ LGPL = Library (or Lesser) GPL
- ◆ AGPL = Affero GPL



<http://www.fsf.org>

FSF & GNU

- ◆ Most of the commercial licenses are designed to take away or restrict user freedom
- ◆ GPL is designed to protect users' freedom (to examine, modify, share the code.)
- ◆ Linux depends on GNU software to a large extent.
- ◆ But GNU project has no working kernel. Depends on Linus for his kernel.
- ◆ GPL allows **commercial for-profit use.**
- ◆ GPL restricts **monopoly rents.**

Please refer to Licensing Slides for detailed discussion

Linux @ Home?

- ◆ Linux @ Home is Real, Feasible (and FUN too)
- ◆ Homes need very different applications compared to work:
 - ◆ ADSL/Phone dialup vs LAN
 - ◆ Rich multimedia & Entertainment apps vs limited intranet, client/server Office apps
 - ◆ Connectivity to multitude of (cheap) peripherals – colour printers, scanners, mobile phones, PDA, tablets, joysticks etc
 - ◆ No stringent requirements
 - ◆ Ease of use (untrained people)
 - ◆ Inadequate backups, security processes

Linux @ Home?

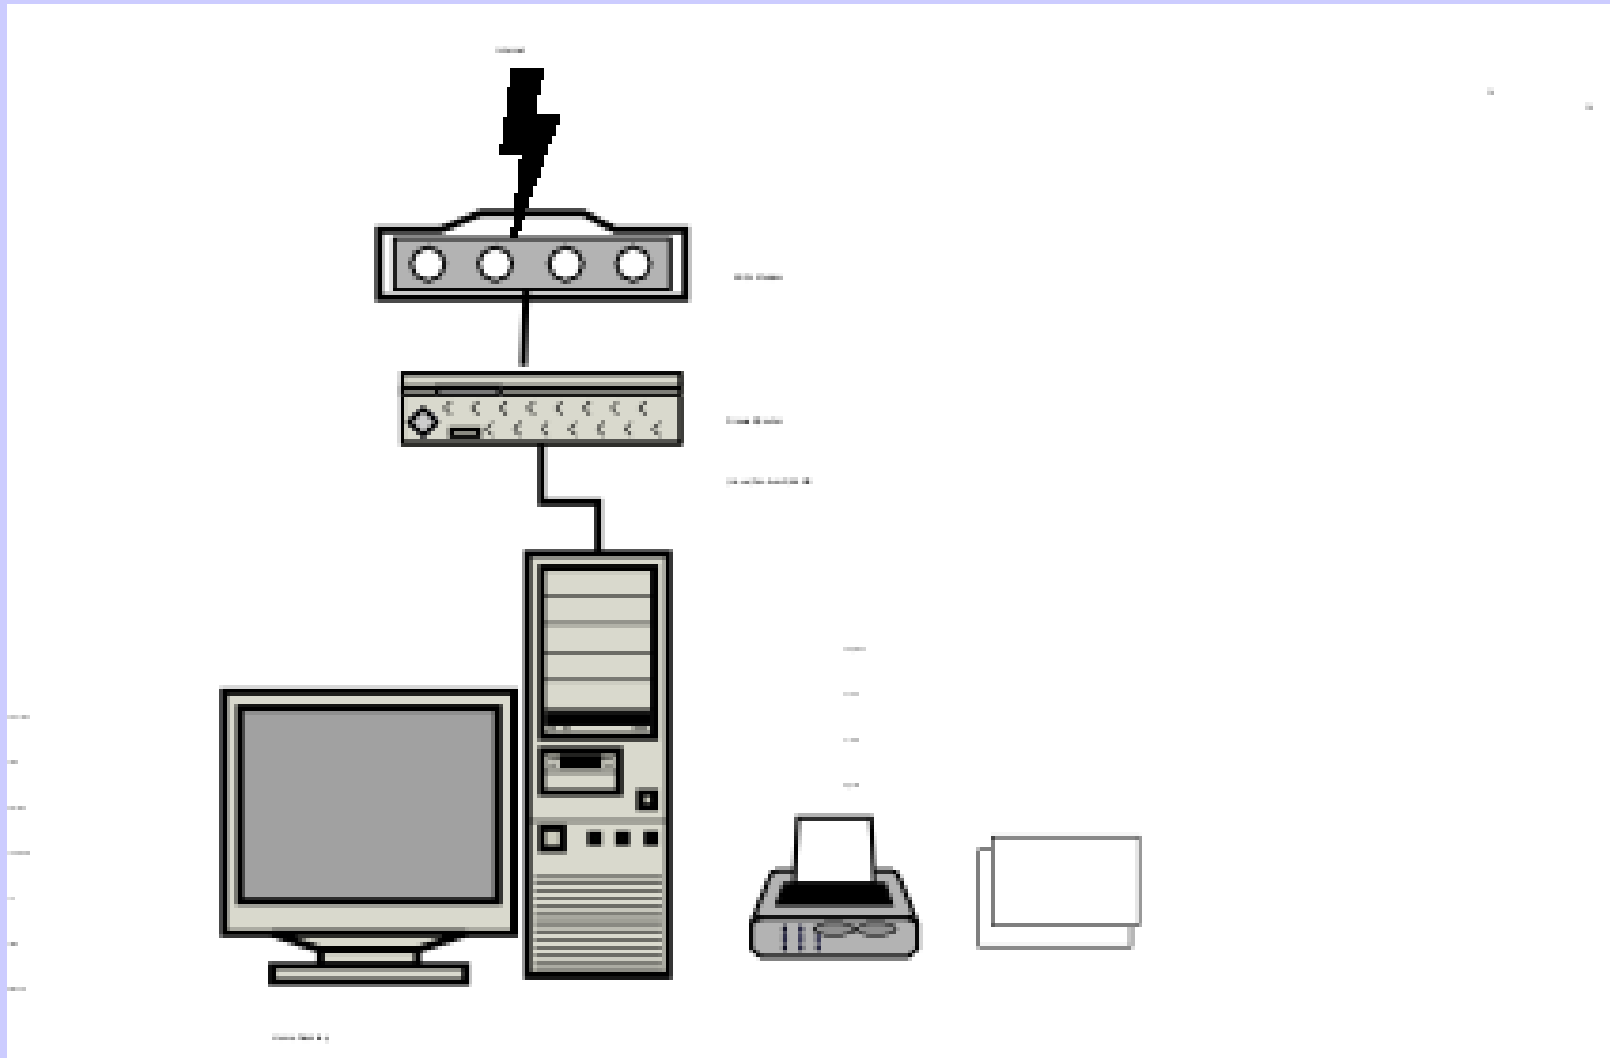
- ◆ Security is equally important
 - ◆ Direct connection to Internet at home vs professional firewall, IDS & network admins.
 - ◆ e-banking, CPF, e-Library, credit card purchases, e-trading accounts
- ◆ Shared usage @ home vs 1 user/desktop at office
- ◆ Kids, grandma, non-IT literate members vs IT-literate workers at the office (PS: not all office workers are smart or IT literate.)

Linux @ Home? Why?

- ◆ Free
- ◆ Learn programming, media file manipulation
- ◆ Secure – no trojans, viruses, worms, backdoors, activeX, adware, hijacking
- ◆ Easy to lockdown (kids)
- ◆ Remote login (eg login while travelling, from office)
- ◆ Great Choice of apps
- ◆ Multiuser
- ◆ No need for frequent hardware/software upgrades
- ◆ Recycle old hw to new uses

Linux @ Home?

- ◆ A typical home 'network'



Links

News:

- ◆ <http://lwn.net>
- ◆ <http://www.linuxjournal.com>

HW compatibility:

- ◆ <http://www.linuxcompatible.org>

Discussion:

- ◆ <http://slashdot.org>

Singapore/ASEAN:

- ◆ <http://www.lugs.org.sg>
- ◆ <http://linux.meetup.com/5/>

Links

Global:

- ◆ <http://www.linux.org>
- ◆ <http://www.fsf.org>
- ◆ <http://www.opensource.org>
- ◆ <http://www.eff.org>
- ◆ <http://www.kde.org>
- ◆ <http://www.gnome.org>
- ◆
- ◆

Links

Software Listing:

- ◆ <http://freshmeat.net>
- ◆ <http://sf.net>

Linux Vendors/Projects:

- ◆ <http://www.debian.org>
- ◆ <http://www.kubuntu.org> & <http://www.ubuntu.com>
- ◆ <http://fedoraproject.org>

Commercial:

- ◆ <http://www.canonical.com>
- ◆ <http://www.redhat.com>
- ◆ <http://www.suse.com> Or <http://www.novell.com>
- ◆ <http://www.mandriva.com>

Thanks!

Copying or Redistributing

© Copyright Anand Vaidya
email: vaidya.anand@gmail.com

- This presentation can be redistributed as follows:
 - No commercial re-distribution: eg, as part of a for-profit CDROM. Seek my permission first.
 - Must attribute the document creator.
 - Share alike: If you use this document and enhance it or modify, share the modifications or the modified document, which means, I apply:
 - <http://creativecommons.org/licenses/by-nc-sa/3.0/> license



Thanks!

Get these slides at:

[http://www.asianlinux.org/documentation/presentations/linux-intro-*](http://www.asianlinux.org/documentation/presentations/linux-intro-*.)

Thanks for your time. If you have any feedback, corrections or questions please contact me: Anand Vaidya, vaidya.anand@gmail.com

This document was created with OpenOffice 2.x on Debian Linux

