



IBM

 creative
commons

 BY:



Linux Distributions

*What they are, how they work,
which one to choose*

Avi Alkalay <avix@br.ibm.com>

+55-11-2132-2327

*Linux, Open Standards Consultant
IBM Corporation*

ibm.com/open || ibm.com/linux

IBM



Before You Start...

This presentation, specially the CentOS part, is not fully endorsed by IBM.

It is not, in any terms, a way to convince clients the distribution they should choose, leaving this decision only to them.

It may be used to explain to customers what is and how a Linux distribution is built, and leave them with parameters to decide about what is important to each of them, and to fulfill their specific needs.

ibm.com/open || ibm.com/linux

IBM®



What is a Linux Distribution



Ecosystem

(A set of packaged software
(as RPM files) with the
distro's digital signature)

(An installation program,
which is by himself one of
the packaged software)

(People that build value
networks with SW and HW
vendors, community, etc)



What are RPM Packages



- ◆ A core OS technology created to ease software management
- ◆ Similar to a ZIP file, plus meta-information like sw version, description, installation date, etc
- ◆ Responsible for sw installation, deinstallation and upgrades
- ◆ Makes consistency checks for different components inter-relations and dependencies
- ◆ Provides automatic execution of custom scripts before and after a sw is (de)installed
- ◆ Provides file corruption check mechanism





IBM

“The Best Linux Distribution”

ibm.com/open || ibm.com/linux

IBM



***Fact 1: All Enterprise* Linux
distributions are excellent***

***Fact 2: All Enterprise Linux
distributions fulfill real world needs***

*Talking about Red Hat Enterprise Linux and SLES

ibm.com/open || ibm.com/linux

IBM



Fact 3: To choose a Linux distribution is more like to choose a tomato sauce

(it is more related to personal taste and proximity to you, than to the quality of the product)

ibm.com/open || ibm.com/linux

IBM



Important Things to Think About

1. Which distro provider I have closer commercial relations with ?
2. Which distro I have more experience with ?
3. Who has best price model for the value provided ?
4. Which distro is certified by the ISV sw and IHV hw of some specific project ?
5. Unless you know what you are doing, be responsible and use an Enterprise distribution



Fact 4: To use at most 2 distributions is not a bad thing

- ◆ There are not big technical differences between them. A RHEL sysadmin can handle SLES too, and vice-versa
- ◆ Diversity provides choice, and make providers fight for you, in quality and in price
- ◆ The magic number is 2. Not 1. Not 3.



Instead of asking:

Whats the best distro ?

Start thinking about:

What is the best distribution for my company and my specific project ?

ibm.com/open || ibm.com/linux

IBM



What Customers look for on Enterprise Distributions ?

They look for 'support', which really means:

1. A partner available now and in the long term, to transfer operational risks
2. Fast access to quality updates
3. Large set of certified hw and ISV sw
4. Availability of pre-tested complex solutions ecosystem

ibm.com/open || ibm.com/linux

IBM®





IBM



Popular Distributions

*Aspects of Red Hat, SuSE,
Fedora, OpenSuse, Debian, etc*

ibm.com/open || ibm.com/linux

IBM



|  SLES |  RHEL |
|---|--|
| ↑ Include Java, Flash and other popular no-cost closed-source software | ↔ Strict policy to include only Open Source and patent-free software |
| ↔ Stick with what already works | ↔ Include last stable innovations from the Open Source world |
| ↓ Uncommon interpretation of standards like the FHS, LSB, jpackage | ↑ Common sense strict conformance to Linux standards, including jpackage |
| ↓ Naming conventions (packages, folders, filenames, etc) have some "SuSE" signature | ↑ Naming conventions are generic and Red Hat-independent |
| ↓ Different source packages (kernel, libc, etc) for different architectures | ↑ Same source packages across all platforms, with intelligent build system |
| ↓ Not all source packages available to the public | ↑ All source packages available to the public |
| ↑ Unified, componentized and more complete configuration tools via YaST, for TUI and GUI | ↓ Small separate configuration tools for the generic situations, mostly for GUI |
| ↔ Building an ecosystem of ISVs, IHVs and users | ↑ Strong ecosystem established |
| ↔ Company: High level products like mail server, firewall, etc, plus Novell closed-source products, including Java MW | ↔ Company: PostgreSQL, Java MW, and infrastructural products like GFS, Cluster System, Directory Server, all OSS |

↑ favorable ↓ unfavorable ↔ goot to know

ibm.com/open || ibm.com/linux



- ◆ SuSE and Red Hat testbed for next enterprise versions
- ◆ Same flavor of their mother “E” distributions
- ◆ Available only on some platforms (ia32, x86_64)
- ◆ Not committed to build an ecosystem of ISVs and IHVs
- ◆ Not committed to stability
- ◆ Open development and testing with community help





IBM

CentOS Linux

A very special Linux distribution

ibm.com/open || ibm.com/linux

IBM





***Fact: CentOS is byte-identical
to Red Hat Enterprise Linux***

ibm.com/open || ibm.com/linux

IBM®



CentOS Aspects

- ◆ Made possible because Red Hat has a policy that all they write must be Open Source, including build process, integration, debugging methods, etc. So...
- ◆ ...to be Open Source means to be reproducible
- ◆ Available on all platforms (ia32, x86_64, ia64, s390x, ppc)
- ◆ Binary RHEL-ready hw drivers are technically CentOS-ready too
- ◆ ISV sw for RHEL behaves exactly the same on CentOS
- ◆ Coherent version numbers
CentOS 3.2 equals to RHEL 3 update 2





But how CentOS does that ?


***Answer is in the core of how the RPM
technology works, and on how RHEL
source-packages are deployed***

ibm.com/open || ibm.com/linux

IBM



RPM: The Raw Source Code



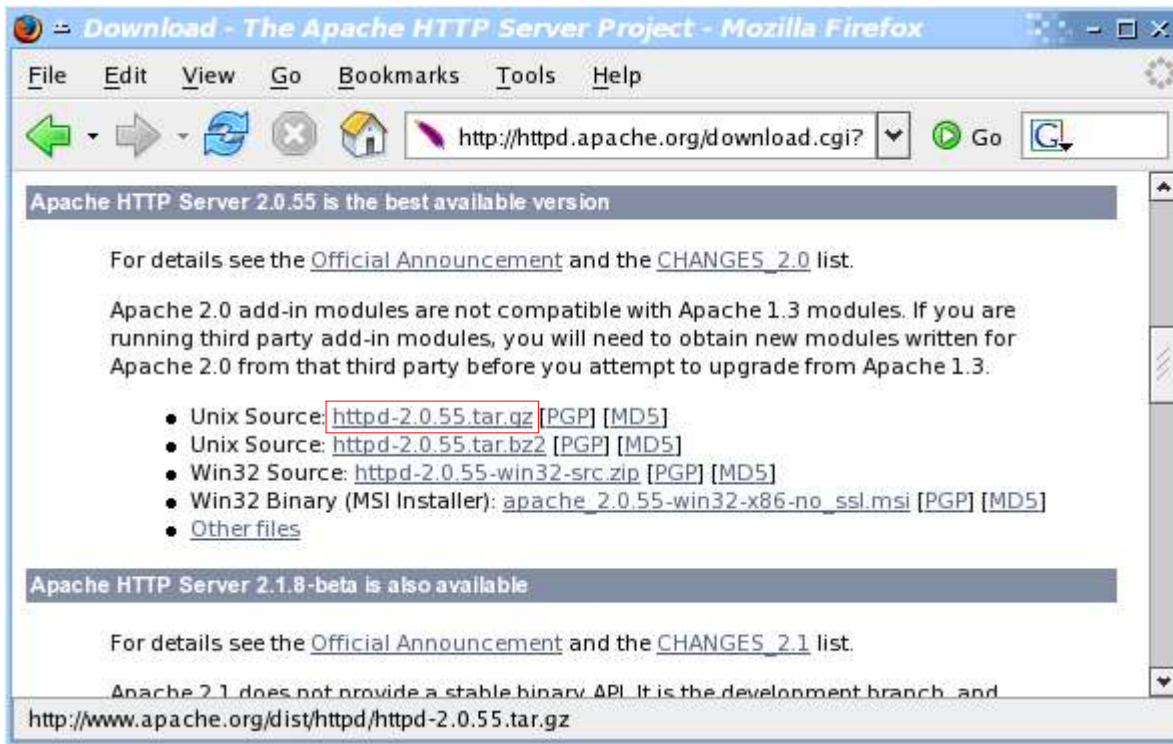
- ◆ The base original thing to start with
- ◆ Collection of sources in C, C++, documentation, etc
- ◆ Published with a version number on the sw home page
- ◆ Use to be difficult to compile, and needs customizations

ibm.com/open || ibm.com/linux

IBM®



RPM: The Raw Source Code



Apache Web Server source code being distributed on his web site: generic and not flavored to any distribution


- ◆ GZipped TAR is an archive format similar to ZIP, but better
- ◆ Includes all Apache source in C language, documentation etc

ibm.com/open || ibm.com/linux

IBM



RPM: The Distribution-specific Patches




- ◆ Modifications to the original source code
- ◆ Distribution bug fixes
- ◆ Distribution specific addons, additional sources, plugins
- ◆ Backports of advanced, next-generation features

ibm.com/open || ibm.com/linux

IBM



RPM: The RPM Specification (spec) File




- ◆ Package summary, description, version, release, etc
- ◆ The pre-install, post-install, etc scripts
- ◆ How to integrate patches and sources, and compile all
- ◆ Where to install the compiled files on the filesystem

ibm.com/open || ibm.com/linux

IBM



RPM: The Source-RPM File



- ◆ All three components grouped in one distributable file
- ◆ Source-RPMs are automatic cooking machines
- ◆ They contain, in open formats, all the flavor a distribution wants to give to certain software

ibm.com/open || ibm.com/linux

IBM



RPM: Source-RPM Files




RHEL 4AS updated Apache (httpd), Kernel and other source-RPMs: freely available on Red Hat's FTP server

- ◆ Use to be freely distributable under GPL license




RPM: The Build Process



- ◆ A standard single `rpmbuild` command generates the installable RPM file
- ◆ No logic, customization, brain or magic is needed in this step




RPM: The Generated RPM File



- ◆ The RPM file is generated with distro's digital signature
- ◆ The usable RPM file use to be not freely-distributable on commercial distributions



RPM: The Generated RPM File



The screenshot shows a web browser window displaying the index of RPM files on the OpenSuSE FTP site. The browser's address bar shows the URL <http://ftp.opensuse.org/pub/opensuse/distribution/SL-10.0-OSS/inst-source/suse/i686>. The page title is "Index of /pub/opensuse/distribution/SL-10.0-OSS/inst-soi". The main content is a table with three columns: "Name", "Last modified", and "Size". The table lists several RPM files, including "db-4.3.27-7.i686.rpm", "db-devel-4.3.27-7.i686.rpm", "db41-4.1.25-78.i686.rpm", "db41-devel-4.1.25-78.i686.rpm", "qlibc-2.3.5-40.i686.rpm", and "qlibc-devel-2.3.5-40.i686.rpm". The status bar at the bottom of the browser window shows "Done".

| Name | Last modified | Size |
|---|-------------------|------|
| Parent Directory | | - |
| MD5SUMS | 15-Sep-2005 02:58 | 360 |
| db-4.3.27-7.i686.rpm | 09-Sep-2005 21:55 | 1.5M |
| db-devel-4.3.27-7.i686.rpm | 09-Sep-2005 21:55 | 4.1M |
| db41-4.1.25-78.i686.rpm | 09-Sep-2005 21:55 | 646K |
| db41-devel-4.1.25-78.i686.rpm | 09-Sep-2005 21:55 | 1.2M |
| qlibc-2.3.5-40.i686.rpm | 09-Sep-2005 20:29 | 3.0M |
| qlibc-devel-2.3.5-40.i686.rpm | 09-Sep-2005 20:29 | 3.4M |

Apache/2.0.49 (Linux/SuSE) Server at ftp.opensuse.org Port 80

Some binary RPMs freely available from a free distribution: OpenSuSE


- ◆ SuSE is a commercial distribution which doesn't provide binary RPMs, different from OpenSuSE (the SuSE beta) which provides RPMs

ibm.com/open || ibm.com/linux

IBM



RPM: Designed to Rebuild on Any System



CentOS and our Support Aspects


Customers look for these 'support' aspects:

1. A partner available now and in the long term to transfer operational risks
↓ CentOS: Not provided by CentOS
2. Fast access to quality updates
↔ CentOS: A few hours delay compared to RHEL
3. Large set of certified hw and ISV sw
↔ CentOS: Technically inherited from RHEL, but not certified
4. Access to pre-tested complex solutions ecosystem
↔ CentOS: Technically inherited from RHEL, but not certified





IBM



Avi Alkalay <avix@br.ibm.com>

11-2132-2327

*Linux, Open Standards Consultant
IBM Corporation*

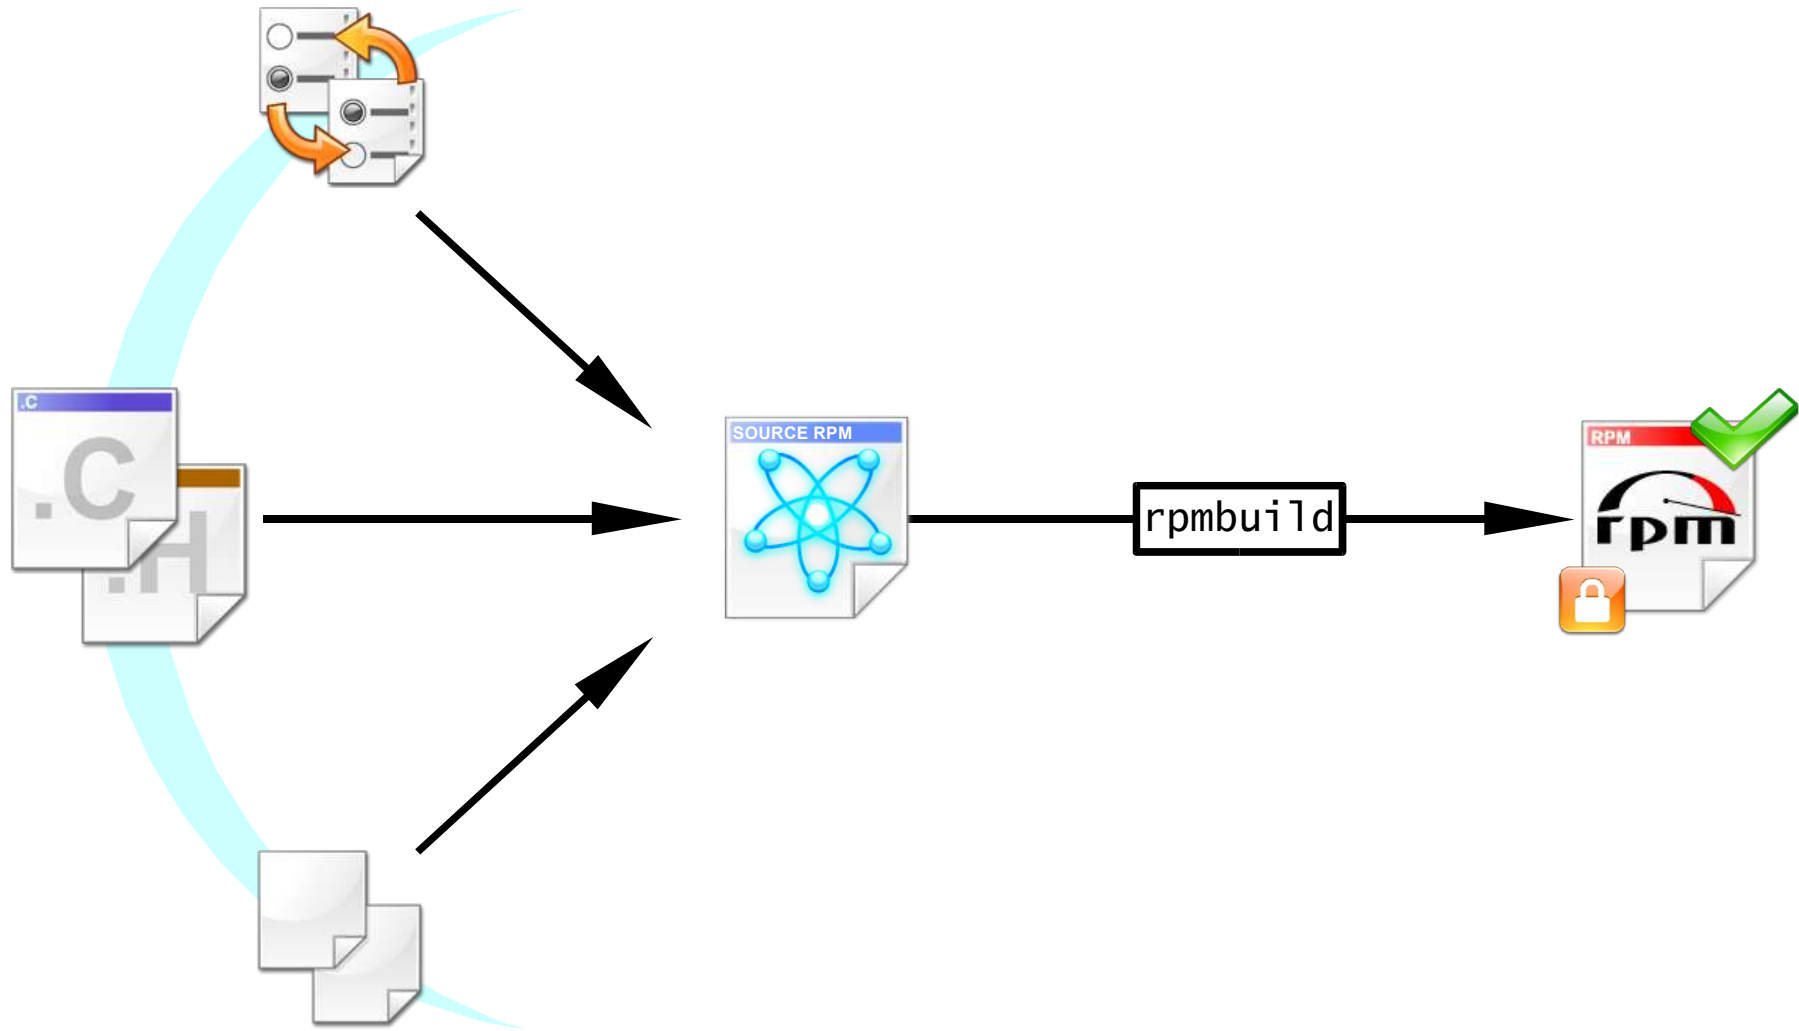
Thank You !

ibm.com/open || ibm.com/linux

IBM



RPM: The Whole Process



ibm.com/open || ibm.com/linux

IBM

