



ITSO -- Linux for zSeries update workshop

Introduction

Michael MacIsaac (mikemac@us.ibm.com), IBM US
(or)
Erich Amrehn (amrehn@de.ibm.com), IBM Germany



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation

LINUX is a registered trademark of Linus Torvalds

Penguin (Tux) compliments of Larry Ewing

Tivoli is a trademark of Tivoli Systems Inc.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.




This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.


Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.





- ★ Introduction
- Brief history and level set
- zSeries platform options
 - ▶ Break
- Distributions (SuSE, Red Hat, Turbolinux, others)
- Possible scenarios and solutions
 - ▶ Lunch
- Linux for zSeries in the enterprise
- Miscellaneous topics
 - ▶ Break
- Sizing and TCO explained
- Summary, future, Q&A
 - ▶ End of the day





Objectives for the Day

- **Learn a brief history of open source software and Linux**
- **Learn about the zSeries platform options**
- **Learn about the Linux distributions for zSeries**
- **Understand Linux installation steps**
- **Understand how to position Linux for zSeries**
- **Know the enterprise solutions available for Linux**
- **Understand sizing and TCO basics**



Who am I?

■ **Michael MacIsaac**

- ▶ ITSO project lead - 3 years (OS/390 UNIX, Linux)
- ▶ S/390 marketing technical support - 2 years (Wind/U, File and print)
- ▶ SP application programmer - 5 years (C, C++)
- ▶ GIS application programmer - 4 years (Fortran, REXX, C)
- ▶ Bachelors of Computer Science, Northeastern University, Boston
- ▶ E-mail: mikemac@us.ibm.com

■ **Erich Amrehn**

- ▶ ITSO project lead - 28 months (Linux, WCS on OS/390, USS, VM/ESA and some VSE books)
- ▶ S/390 VM/ESA Field support center (technical marketing support - 8 years)
- ▶ VM systems programmer - 5 years (REXX, EXEC, PIPELINES, ASSEMBLER)
- ▶ IBM Hardware technician 4 years (370-158, 43xx, Tapes)
- ▶ Bachelors of Telecommunication
- ▶ E-mail: amrehn@de.ibm.com





Who are you?

- Who is your employer?
 - ▶ IBM - sales or technical
 - ▶ Vendor - company that sells software
 - ▶ Customer - company that uses IT to be successful
- Do you use Linux for zSeries?
 - ▶ No plans
 - ▶ Plan to, but still don't
 - ▶ Use it, but not in production
 - ▶ Use it in production
- What is your primary desktop?
 - ▶ Windows
 - ▶ Linux
 - ▶ Other
- Is there any area I should emphasize, given the agenda and objectives?

@server



Background of the Workshop

- Based on redbooks written this year
 - ★ *Linux for IBM eserver zSeries and S/390: Distributions*
 - ▶ SG24-6264-00
 - ▶ Published Sept. 20, 2001
 - ▶ <http://www.redbooks.ibm.com/redpieces/abstracts/sg246264.html>
 - ★ *Linux for zSeries and S/390: ISP/ASP Solutions*
 - ▶ SG24-6299-00
 - ▶ Currently a redpiece
 - ▶ <http://www.redbooks.ibm.com/redpieces/abstracts/sg246299.html>
- Also based on a redpaper:
 - ★ *Building Linux Systems Under IBM VM*, Simon Williams
 - ▶ <http://publib-b.boulder.ibm.com/Redbooks.nsf/RedpaperAbstracts/redp0120.html>

@server



ITSO Residencies

- Usually 4 -- 6 week projects
 - ▶ Install, document new hardware and software
 - ▶ Suitable for IBMers, usually customers & business partners
 - ▶ Dates are usually negotiable
- Published in hard copy and on the Web
 - ▶ <http://www.redbooks.ibm.com/>
 - ▶ FatBrain publishes Redbooks on demand
- ITSO will pay travel and living expenses
 - ▶ In conjunction with IBM travel guidelines - see Web for details
- Unsolicited testimonial from a recent resident:
 - ▶ "Once again thanks for your input and help during my stay in Poughkeepsie, I thoroughly enjoyed the experience and the comments I received from people that read the draft copy so far is only good."



Advertisement





Notes:

@server



Agenda - section 2 of 9

- Introduction
- ★ Brief history and level set
- zSeries platform options
- Distributions (SuSE, Red Hat, Turbolinux, others)
 - Break
- Possible scenarios and solutions
 - Lunch
- Linux for zSeries in the enterprise
- Miscellaneous topics
 - Break
- Sizing and TCO explained
- Summary, future, Q&A

@server

History - Outline



- Open source software
- Linux in general
- Linux on S/390



@server

History: Open Source Software




- In the beginning, all source code was open
 - ▶ For example, IBM did not charge for software before 1970
- Ancient history
 - ▶ UNIX - 1969, Ken Thompson, Dennis Ritchie, AT&T bell labs
 - Sent source code to U. Cal at Berkeley in 1974
 - ▶ ArpaNet - 1969, DoD
 - First node-to-node message (UCLA and Stanford Research Inst.)
 - ▶ TCP/IP - 1974, Vint Cerf and Bob Kahn
 - sockets, e-mail, news groups, FTP, Gopher
 - ▶ PC - 1980s, Apple, IBM, the "clones"
 - ▶ GNU/FSF - 1984, Richard Stallman, MIT
 - Drafted the GNU General Public License (GPL) with the help of a lawyer
 - ▶ WWW - 1990s, Tim Berners-Lee, CERN
 - ▶ Linux - 1991, Linus Torvalds, University of Helsinki



@server

History - Culture clash



Lou Gerstner



Richard Stallman

@server

History: What is open source software?



■ Freely available software in source code format

- ▶ Terminology: freeware vs. open source software
- ▶ Free means "liberty", not "no cost"
- ▶ Term "open source software" is more conservative

■ Usually distributed via Internet or CD

■ Often developed in a collaborative community effort

- ▶ These communities are "meritocracies": hierarchies in which position is based on merit, not on title, nor organization

■ Not "Public Domain" - Frequent misconception

- ▶ Copyright law gives author sole rights
- ▶ Author determines appropriate uses of a work
- ▶ Author may grant permission via a license
- ▶ License grants permission but also may impose obligations

@server



History: Open source licenses

■ **Public Domain**

■ **GNU General Public License (GPL)**

- ▶ You can run, copy and modify the software
- ▶ You can redistribute and charge \$\$ for the software
- ▶ You cannot add restrictions to the software
- ▶ You must make the source code available
- ▶ If you include software which is GPL'd, your software must also carry the GPL (viral nature)

■ **GNU Library GPL (LGPL)**

■ **BSD-style licenses**

■ **Artistic License**

■ **Netscape, Mozilla Public Licenses**

■ **IBM Public License**



History: Linux in General

■ **The term Linux is commonly used for two things:**

- ▶ An operating system kernel written by Linus Torvalds
- ▶ One of many free operating systems called **distributions**
 - Include the Linux Kernel
 - Include much more open source software/freeware
 - Sometimes called GNU/Linux

■ **Linux is UNIX-like, but not a version of UNIX**

- ▶ Runs on many (>140) platforms
- ▶ Code is in a common base

■ **Originally written to run on Linus' 386 - 1991**

- ▶ First attempt at a port of Linux was to the Motorola 68000
 - Linus considered this a code-fork that created a Linux-like OS
- ▶ First serious effort for portable code was to DEC Alpha 1993-94
 - Linus did not want to create a 3rd code base
 - Kernel underwent a major rewrite - motivated by how to work in a community





History - Linux' 10 year birthday

From: torvalds@klaava.Helsinki.FI (Linus Benedict Torvalds)

Newsgroups: comp.os.minix

Subject: What would you like to see most in minix?

Summary: small poll for my new operating system

Date: 25 Aug 91 20:57:08 GMT

Organization: University of Helsinki

Hello everybody out there using minix -

I'm doing a (free) operating system (**just a hobby, won't be big** and professional like gnu) for 386(486) AT clones. This has been brewing since april, and is starting to get ready. I'd like any feedback on things people like/dislike in minix, as my OS resembles it somewhat (same physical layout of the file-system (due to practical reasons) among other things).

I've currently ported bash(1.08) and gcc(1.40), and things seem to work.

This implies that I'll get something practical within a few months, and I'd like to know what features most people would want. Any suggestions are welcome, but I won't promise I'll implement them :-)

Linus (torvalds@kruuna.helsinki.fi)

PS. Yes - it's free of any minix code, and it has a multi-threaded fs.

It is NOT portable (uses 386 task switching etc.), and **it probably never will support anything other than AT-harddisks**, as that's all I have

@server



One Year Birthday



@server

One Year Birthday - May 17th 2001



December 1999

- Linux for S/390 patches to Open Source

January 2000


- Marist Distribution
- Joint customer project

February 2000

- Linux World NYC/Paris

May 2000

- Vista 2000
- SuSE and Turbolinux distributions announced
- **IBM Support for Linux for S/390 announced, May 17th**



August 2000

- Integrated Facility for Linux (IFL) announced
- Virtual Image Facility (VIF) for Linux announced
- IBM & ISV middleware announced
- IBM Global Finance announcements
- SuSE Install Fest

October 2000

- IBM commits 64-bit patches for zSeries
- Gigabit ethernet support ships
- Red Hat announces distribution
- Red Hat, SuSE and Turbolinux announce support for iSeries, pSeries, xSeries, and zSeries models

November 2000

- Turbolinux InstallFest
- SuSE Linux for S/390 available

December 2000

- IBM middleware for Linux for S/390 rolls out
- 64-bit patches for zSeries to Open Source
- Turbolinux available in Japan

January 2001

- Linux World NYC - Best hardware award for zSeries
- Turbolinux for S/390 available

March 2001

- mySAP.com announced for Linux for S/390

June 2001

- Turbolinux for S/390 v6.5 beta available
- Additional enhancements available on developerworks

September 2001

- Turbolinux for S/390 v6.5 GA



Quotes on Linux for zSeries (S/390)



"This is the first time that Linux will be run in a totally bulletproof environment. As 24*7 availability becomes more and more important on the Web, this will be absolutely crucial."

- Dale Vile, senior analyst - Bloor Research

"It is an extremely interesting move because the S/390 offers a range of possibilities that doesn't exist on other (computing) platforms,"

- Dan Kaberon, Parallel Sysplex Manager - Hewitt Associates LLC

"People had been moving more to the Unix camp; 2001 has been a year of resurgence for the mainframe"

- David Mastrobattista - Giga Information Group

"For a mature response to open-source software, consider IBM Corp.'s decision to offer Linux on its S/390 mainframes, the machines that run banks, factories, and the Pentagon."

- Hiwatha Bray - Boston Globe 5/18/2000

"It doesn't take a sleuth to see that as schools go on Linux, you could have 600 million students who have never used another operating system,"

- Frederick Berenstein - Linux Global Partners

This shows a degree of imagination from IBM that I haven't seen for some time. For once, they seem to be with the market instead of two to three years behind it."

- David Floyer, VP of Research, ITcentrix





Join the community!

■ LINUX-390 list server:

- ▶ Archives are on the Web at:
<http://www.marist.edu/htbin/wlvindex?linux-390>
- ▶ Subscribe and tailor by sending an e-mail to listserv@vm.marist.edu
 - In the first line put [sub linux-390 first_name last_name](#)
 - Follow and save the directions that follow
 - [set linux-390 mail](#) // 1 e-mail/message
 - [set linux-390 digest](#) // 1 e-mail/day
 - [set linux-390 nomail](#) // no e-mail
 - [get linux-390 log0109](#) // get Sept 2001 archives
 - [signoff linux-390](#) // to sign off the list
 - To append to the list send an e-mail to linux-390@vm.marist.edu

■ Try the code! (start on a PC)

- ▶ Linux is more accessible than any other OS in history



History: Web articles

- ▶ *S/390: The Linux Dream Machine* by Scott Courtney
 - <http://www.linuxplanet.com/linuxplanet/reports/1532/1/>
- ▶ *Has the Linux 'dream machine' arrived?* by Mary Jo Foley
 - <http://www.zdnet.com/zdnn/stories/news/0,4586,2491800,00.html>
- ▶ *Free Mainframe* by Laura Wonnacott
 - <http://www.internetweek.com/reviews01/rev073001.htm>
- ▶ *It's Official: IBM Announces Linux for the S/390*, by Scott Courtney
 - <http://www.linuxplanet.com/linuxplanet/reports/1848/1/>
- ▶ *OpenSources* by Linus Torvalds, Richards Stallman, Eric Raymond, Bruce Perens, many more
 - <http://www.oreilly.com/catalog/opensources/book/toc.html>





History: Web articles (cont'd)

- ▶ *How Big Blue Fell for Linux* by Andrew Leonard
 - http://www.salon.com/tech/fsp/2000/09/12/chapter_7_part_one/index.html
- ▶ *The Cathedral and the Bazaar* by Eric S. Raymond
 - <http://www.tuxedo.org/~esr/writings/cathedral-bazaar/>



Agenda - section 2 of 9

- Introduction
- ★ Brief history and level set
- zSeries platform options
 - ▶ Break
- Distributions (SuSE, Red Hat, Turbolinux, others)
- Possible scenarios and solutions
 - ▶ Lunch
- Linux for zSeries in the enterprise
- Miscellaneous topics
 - ▶ Break
- Sizing and TCO explained
- Summary, references, future, Q&A





Level set - What is Linux?

- **Linux is platform independence**
- **Linux is the industry's next UNIX**
- **Linux is the most obvious result of open source and open standards**



Level set - IBM's Linux strategy

- **Leverage Linux to create a pervasive application development and deployment environment that will drive applications growth**
 - ▶ Responsible participation with the Open Source community
 - ▶ Support a native Linux on all server platforms
- **Ease the deployment of Linux applications on IBM servers**
 - ▶ Develop IBM Linux based offerings
 - ▶ Platforms, services and packaged solutions
- **Expand IBM Linux Technology Center**
 - ▶ Partner and contribute IBM technology and skills to the open source community
 - ▶ Enhance IBM Linux and Open Source technical skills





Level set - IBM's Linux strategy

■ Irving Wladawsky-Berger on Linux and open standards

- ▶ Vice president of technology strategy IBM Server Group
- ▶ Technology is increasing 50-100% **per year**
- ▶ How does one manage this fantastic growth rate?
 - Project eLiza - self managing servers
 - Open standards
 - Open sharing of ideas
 - Professional communities such as physics or chemistry have always published information openly
 - Sometimes English is not an adequate language to describe how to share - for example
 - ♦ In Physics, mathematics is necessary
 - ♦ In software, source code is necessary
 - To share ideas a reference system is needed - for example IPv6
 - It is reasonable for higher level software to be proprietary
- ▶ Linux
 - Is clearly the fastest growing operating system
 - "In IBM, we are committed to embracing Linux across everything that we do"
 - The number one market segment where our customers are using Linux is workload consolidation

■ See <http://www.ibm.com/news/us/2001/08/15.html>



Level set - IBM's commitment to Linux

Hardware	Software	Services	Alliances	Open Source
				
Servers Desktops	All Major Applications	Support And Training Consulting	Distributor Partnerships	Code Contributions Technical Resources
zSeries, pSeries, iSeries, xSeries, ThinkPads, NetVista	WebSphere, Domino, DB2, Tivoli, VisualAge Java, MQ Series, ViaVoice	Learning Services, Redbooks, Supportline, Consulting Services	Major 4 Distributors Caldera, Red Hat, SuSE, Turbolinux	Manpower, Knowledge, Technology





Level set - S/390 design

■ Linux for S/390 design

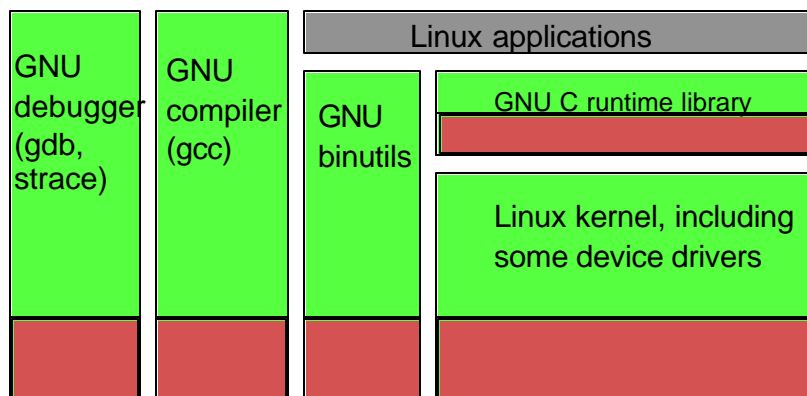

- ▶ Linux is Linux, S/390 is S/390
- ▶ Runs Native, in an LPAR or on VM/ESA
- ▶ Pure Linux, in ASCII environment
- ▶ Exploits IBM S/390 hardware, including IEEE floating point

■ Not a unique version of Linux or other operating system

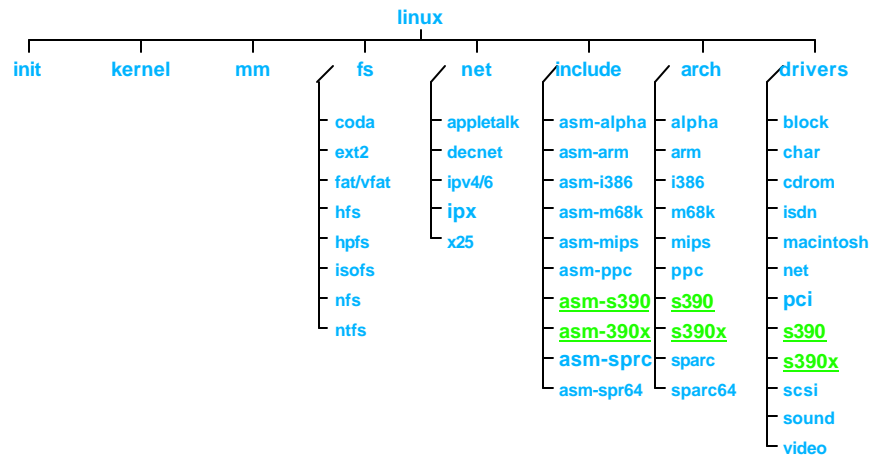
■ Not a replacement for other IBM S/390 operating systems



S/390 modifications



Internals: Kernel code tree



@server

Level set: Linux for S/390 native drivers

■ DASD

- ▶ ECKD - 3380, 3390 - can now be a module, dasd_mod.o
- ▶ Minidisk

■ Network

- ▶ OSA token ring, ethernet - lcs.o (OCO)
- ▶ OSA-Express Gigabit Ethernet - qeth.o, qdio.o (OCO)
- ▶ Channel To Channel, Escon
- ▶ IUCV - VM and VIF only
- ▶ CLAW (Cisco Mainframe Channel Connection)
 - ▶ <ftp://ftp.utsglobal.com/pub/c7000/>
- ▶ Hipersockets

■ Others

- ▶ Hardware console, 3215
- ▶ XPRAM - Expanded memory - can be a module
- ▶ Tape
- ▶ Cryptographic coprocessor

@server



Object Code Only drivers

■ Some people accept the need for OCO drivers

- ▶ lcs.o
- ▶ qdio.o (protocol driver) and qeth.o (device driver)

■ Some people don't like OCO drivers


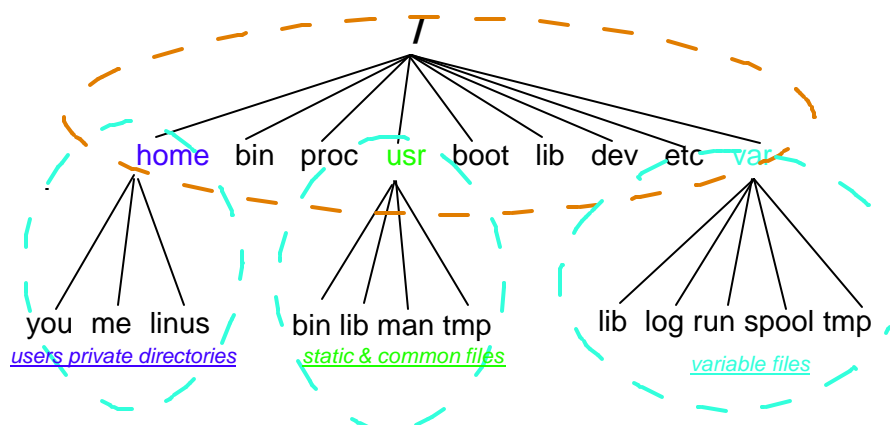
- ▶ Red Hat surprised the community not shipping OCO modules

■ An issue comes up with compatibility - different kernel levels and architectures

- ▶ Drivers advertised on developerWorks site:
 - linux-2.4.7 64-bit 2001-08-20
 - linux-2.4.5 31-bit 2001-08-03
 - linux-2.4.5 64-bit 2001-06-29
 - linux-2.4.5 2001-06-18
 - linux-2.2.19 2001-04-26
 - linux-2.4.3 2001-04-04
 - linux-2.2.18 2001-03-22
 - linux-2.4.0 and linux2.4.0-test11 2000-12-21 & 2001-01-16



Level set - File system tree





Level set - What S/390 does NOT give you

- Monitor, graphics card
- Speakers, sound card
- Mouse, keyboard
- Local printer
- CD, floppy disk
- PCMCIA, Advanced Power Management (APM)
- Many interrupts: Ctrl-Alt-Del, power-fail, power-restore



Notes:





Agenda - section 3 of 9

- **Introduction**
- **Brief history and level set**
- ★ **zSeries platform options**
 - ▶ Break
- **Distributions (SuSE, Red Hat, Turbolinux, others)**
- **Possible scenarios and solutions**
 - ▶ Lunch
- **Linux for zSeries in the enterprise**
- **Miscellaneous topics**
 - ▶ Break
- **Sizing and TCO explained**
- **Summary, references, future, Q&A**




Platform options - Outline


- **Native, LPAR, VIF, z/VM**
- **Possible Hardware**
 - ▶ MP3000 (7060), G5 (9672), G6 (9672), zSeries (2064)
 - ▶ Deprecated: G2, G3, G4, P/390
- **Memory**
 - ▶ Main Memory, Expanded Memory or Virtual Memory (z/VM and VIF)
- **Network options**
 - ▶ OSA, 2216, CTC, VCTC, IUCV
- **DASD**
 - ▶ Mdisk vs. Full Volume DASD



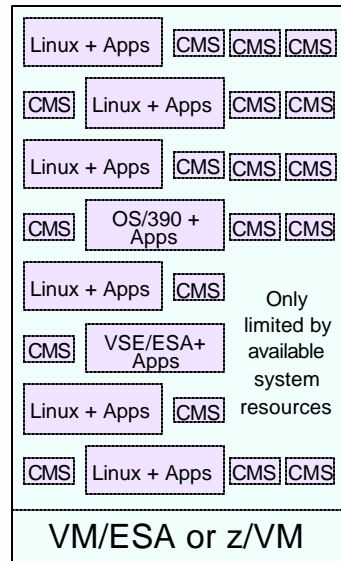
Platform options to run Linux for zSeries



1. Native Processor



2. Logical Partitions



3. VM Guests



Different PU types and memory



■ zSeries

- ▶ CPs, ICFs, SAPs and IFLs
- ▶ Add the number of available PUs for each type
- ▶ CPs and IFLs can be dedicated and shared
- ▶ Max memory 64GB - can be split up into XSTOR and main memory (31-bit)

■ G5, G6

- ▶ CPs, ICFs, SAPs and IFLs
- ▶ CPs and IFLs can be dedicated and shared
- ▶ Max memory 2GB main and 8GB XSTOR

■ MP3000

- ▶ CPs, SAPs, no ICFs, IFL - just added
- ▶ Max of 2 PUs
- ▶ CPs can be dedicated and shared
- ▶ Max memory 4GB

Acronym key:
 PU - Processing Unit
 CP - Central Processor
 ICF - Internal Coupling Facility
 SAP - System Assist Processor
 IFL - Integrated Facility for Linux

■ G2, G3, G4 - deprecated due to lack of IEEE FP instructions






LPARs with IFL "engines"

- All mainframes must have at least one CP
- IFLs and CPs cannot be mixed in an LPAR
- You MAY run:
 - ▶ Linux native
 - ▶ Linux under z/VM 4 or higher
 - ▶ z/VM native with CMS-only applications.
 - ▶ z/VM under z/VM with CMS-only applications.
- The following will **NOT** IPL on a IFL or as a guest of a VM system running on an IFL
 - ▶ VM/ESA (VM 3 or lower)
 - ▶ Any form of z/OS, OS/390 or MVS
 - ▶ DOS/VS
 - ▶ VSE
 - ▶ TPF



OSA Network options (LPAR, z/VM)

- OSA Open Systems Adapter
 - ▶ OSA2-ENTR (Ethernet/Token Ring)
 - ▶ OSA2-FE (Fast Ethernet 10/100)
 - ▶ OSAX-FE (Fast Ethernet Express)
 - ▶ OSAX-GB (Gigabit Express)




SC Duplex MM



SC Duplex SM



OSA card drivers and modes (LPAR, z/VM)



@server

OSA options on G5, G6 (LPAR and z/VM)



Attribute	OSA2-ENTR	OSA2-FE	OSAX-FE (OSE)	OSAX-FE (OSD)	OSAX-GB (OSD)
Connection	Ethernet or Token Ring	Ethernet	Ethernet	Ethernet	Ethernet
Mb/sec throughput	10 (ethernet) or 16 (T/R)	100	100	100 (200 if full duplex)	1000
Cable	copper	copper	copper	copper	fiber-optic
HCD/IOCDS CHP type	OSA	OSA	OSE	OSD	OSD
Ports/CHP	2	1	1	1	1
Path to multiple IP addresses	OSA/SF	OSA/SF	automatic	automatic	automatic
Max IP addresses	16	16	16	240	240
Access IP addresses on other LPARs?	Yes	No	Yes	Yes	Yes
Linux device driver	lcs.o	lcs.o	lcs.o	qdio.o	qdio.o

@server

OSA options on zSeries (LPAR and z/VM)



card (mode)	OSAX-FE (OSE)	OSAX-FE (OSD)	OSAX-GB (OSD)
Connection	Ethernet	Ethernet	Ethernet
Mb/sec throughput	100	100 (200 if full duplex)	1000
Cable	copper	copper	fiber-optic
HCD/IOCDS CHP type	OSE	OSD	OSD
Ports CHP	1	1	1
Max IP addresses	32	480?	480?
Linux device driver	lcs.o	qdio.o	qdio.o



OSA Options in etc/modules.conf



■ 1 chp, 1 port/chp

► options LCS noauto=1 devno_portno_pairs=0x2000,0

■ 1 chp, 2 ports/chp

► options LCS noauto=1 devno_portno_pairs=0x2000,0,0x2000,1

■ 2 chp, 1 ports/chp

► options LCS noauto=1 devno_portno_pairs=0x2000,0,0x3000,0

■ 1 chp

► options qeth qeth_options=noauto,0x4000,0x4001,0x4002,eth0

■ 2 chp

► options qeth qeth_options=noauto,0x4000,0x4001,0x4002,eth0,0x5000,0x5001, 0x5002,eth1





LPAR setup

■ IOCDs

- ▶ Limit the I/Os Linux has access to
 - By using the PARTITION parameter on the CHPID statement:
 - CHPID PATH=(54),TYPE=CNC,PARTITION=(LIN01,REC)
 - CNTLUNIT CUNUMBER=0008,PATH=(54),UNIT=OSA
 - IODEVICE ADDRESS=(5400,02),UNITADD=00,CUNUMBER=(0008),UNIT=OSA
 - By using the PARTITION parameter on the IODEVICE statement:
 - CHPID PATH=(54),TYPE=CNC,PARTITION=((LIN01,PRMVS1,PRMVS2),SHARED)
 - CNTLUNIT CUNUMBER=0008,PATH=(54),UNIT=OSA
 - IODEVICE ADDRESS=(5400,02),UNITADD=00,CUNUMBER=(0008),UNIT=OSA,PARTITION=(LIN01)
 - IODEVICE ADDRESS=(5402,14),UNITADD=??,CUNUMBER=(0008),UNIT=OSA,PARTITION=(LIN01)

@server



LPAR and or Native setup (cont'd)

■ Swap space

- ▶ Recommend just a swap file to start
- ▶ The **fdasd** command will allow you to partition DASD

■ Booting the installation kernel - 3 options

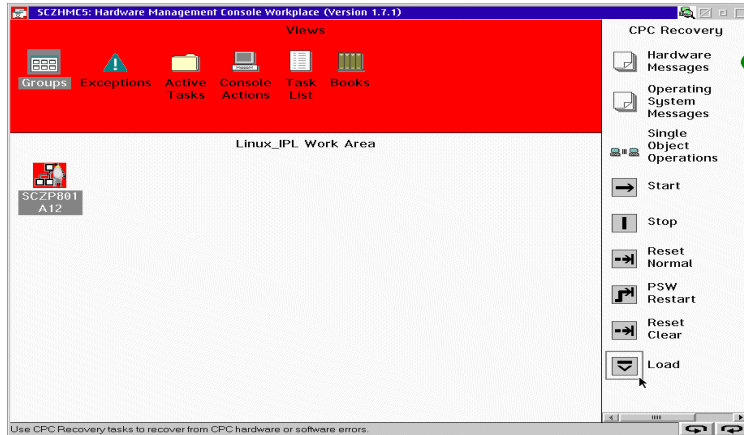
1. Via tape
2. From HMC CD-ROM or FTP server
 - Need correct EC-level
 - G5 or G6 service element
 - CMOS Driver 26W (MCL DR 26W)
 - Microcode fix MCL F99918.025
 - Microcode fix MCL F99933.009
 - Microcode fix MCL F99934.002
 - Multiprise 3000 v Microcode fix EC F34643 MCL048
 - Microcode fix MCL F34663.087
3. Using ICKDSF bootstrap
 - OS/390 can be used to prepare a Linux IPL DASD
 - See <http://pucc.princeton.edu/~rvdheij/linuxipl.html>

@server

LPAR setup - prepare the tape


■ Booting the installation kernel from tape

- ▶ Preparing the Tape
 - FTP the tape boot files with a record length of 1024 bytes (kernel, RAMdisk, parameter file)
 - Get a tape drive (attach)
 - Generate the tape from z/VM, OS/390 or VSE
- ▶ Select the LPAR on the HMC and click **Load**



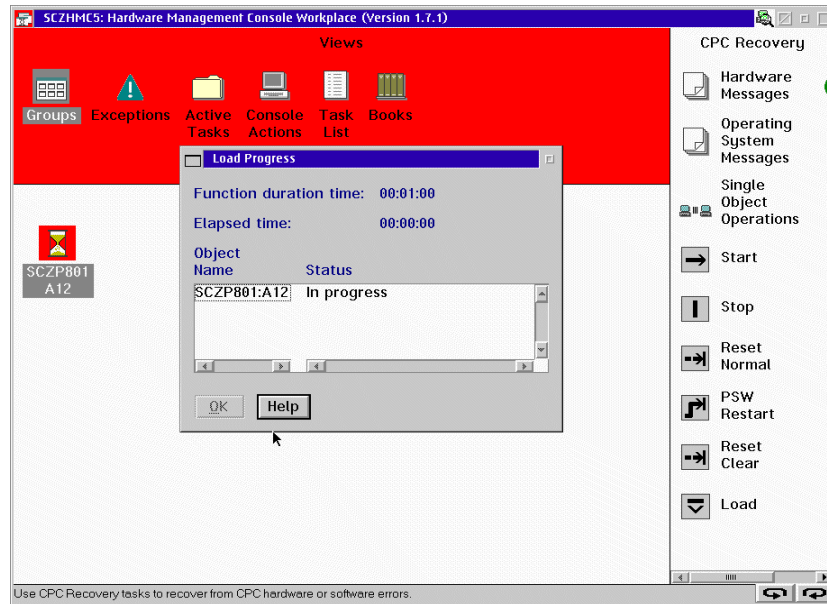
LPAR setup - specify tape address

■ Booting the installation kernel from tape

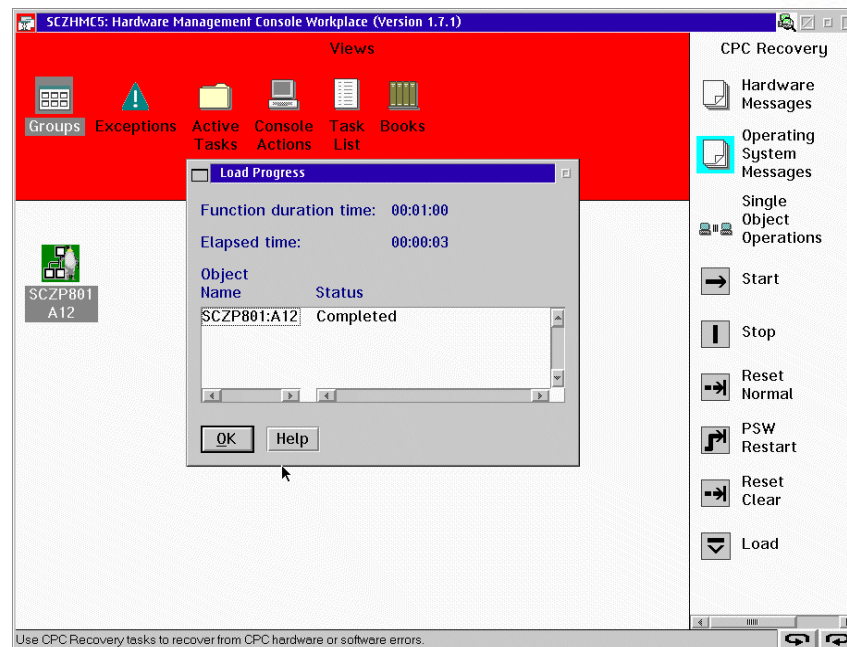


LPAR setup - IPL begins

■ Booting the installation kernel from tape

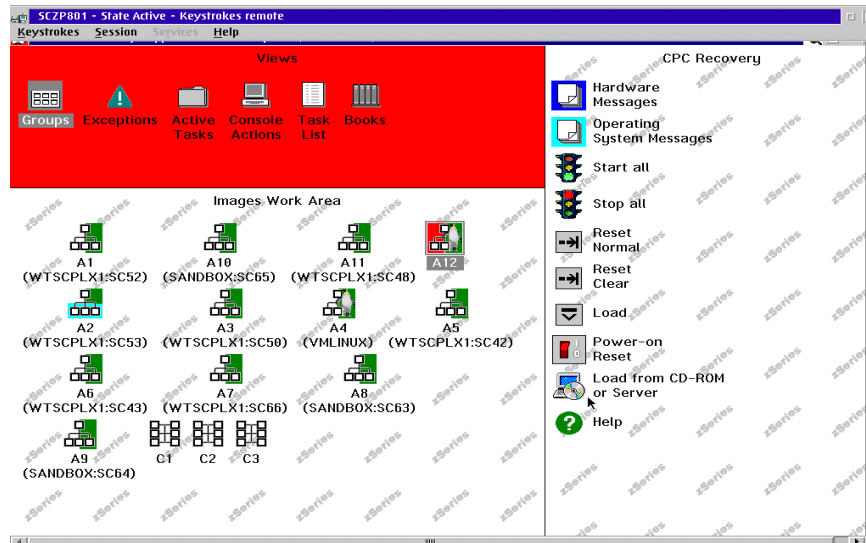


LPAR setup - IPL successful

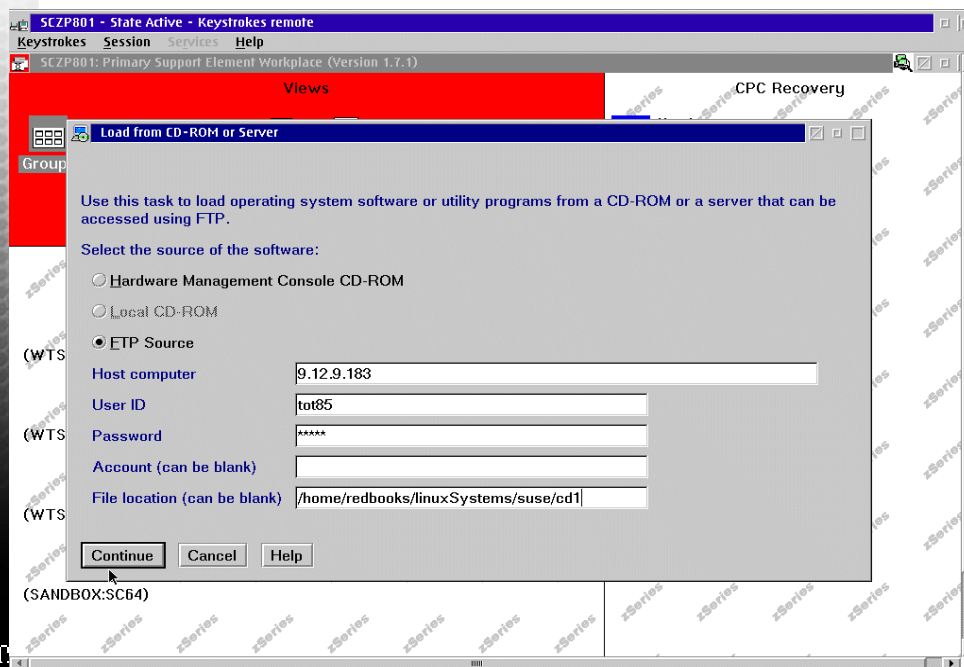


LPAR setup - boot from CD or FTP server

- ▶ Make sure the distribution CD or FTP server has access to kernel-file, RAMdisk-file, parameter file
- ▶ Select the Linux LPAR from "Images Work Area"
- ▶ Double-click "Load from CD ROM or Server"



LPAR setup - boot from FTP server

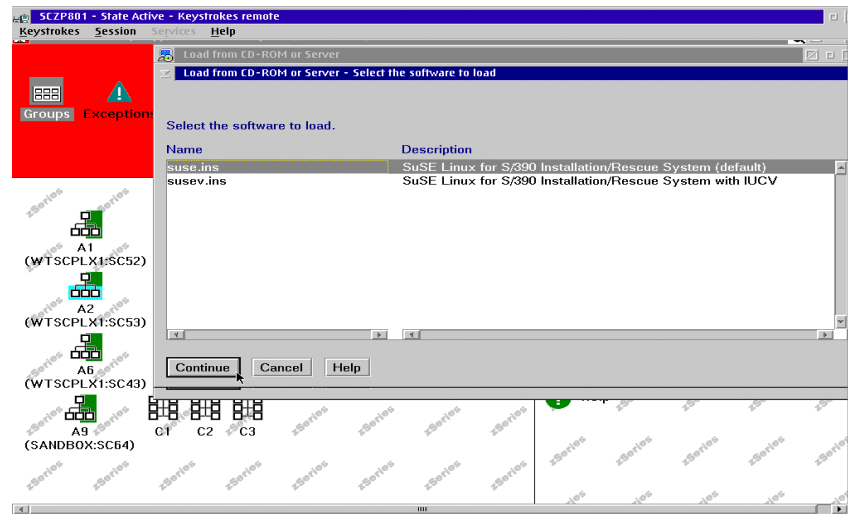


LPAR setup - Load from FTP server (cont'd)



■ Booting installation kernel from HMC CD-ROM or FTP

- The root directory of the load source is scanned for files whose names end in .ins.

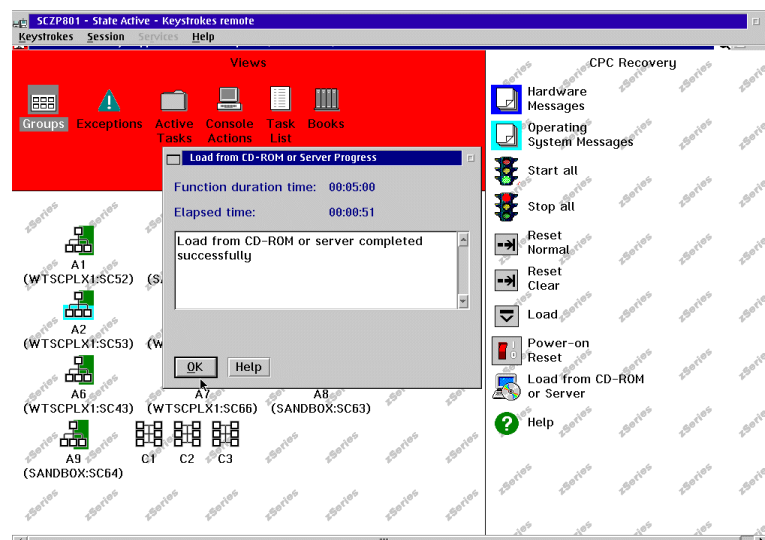


@server

LPAR setup - Load from FTP server (cont'd)



- Confirmation box says "Retrieving code from source"
- Click on "Reset Clear" - confirmation box says "Loading data into system"



@server



LPAR setup

- ▶ When you get this far, double-click "Operating System Messages"
- ▶ You should see the Linux kernel booting
 - We describe for SuSE in "SuSE - boot installation system"
 - We describe for Turbolinux in "Turbolinux - boot installation system"

@server



Notes:

@server



Agenda - section 4 of 9

- **Introduction**
- **Brief history and level set**
- **zSeries platform options**
 - ▶ Break
- ★ **Distributions (SuSE, Red Hat, Turbolinux, others)**
- **Possible scenarios and solutions**
 - ▶ Lunch
- **Linux for zSeries in the enterprise**
- **Miscellaneous topics**
 - ▶ Break
- **Sizing and TCO explained**
- **Summary, future, Q&A**

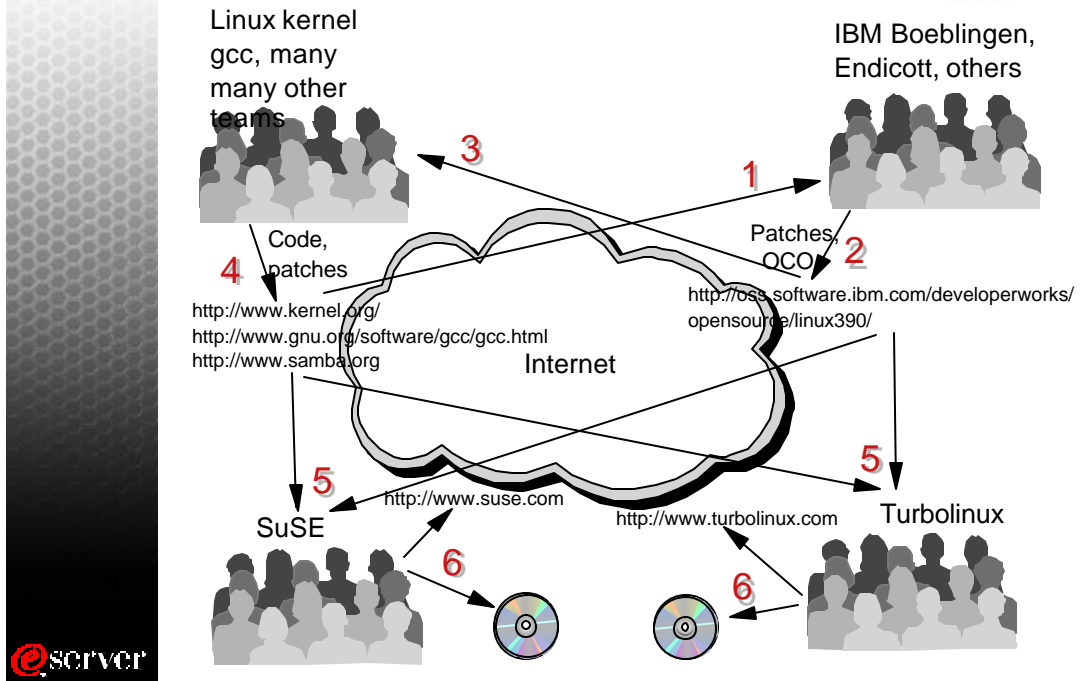


Agenda for the distribution part

- **How are S/390 distribution created?**
- **Overview of distributions**
 - ▶ SuSE
 - ▶ Red Hat
 - ▶ Turbolinux
 - ▶ Marist "file system"
 - ▶ Caiman from LinuxKorea
 - ▶ Think Blue from Millenux
 - ▶ Others
- **Details of SuSE install**
- **Details of Turbolinux install**
- **Discussion of Red Hat install**



How are S/390 distributions created?



Distributions - SuSE

■ SuSE "GA"

- ▶ First major distribution, November 10, 2000
- ▶ Based on SuSE 7.0 for the PC
- ▶ 3 .iso CD images

■ Documentation on CD 1

- ▶ l390gp3.pdf - *Preparing for Installing SuSE LINUX for S/390*, April 5, 2001 LNUX-1001-01, 76pages
- ▶ l390ga3.pdf - *Installing SuSE LINUX for S/390*, April 5, 2001 LNUX-1002-01 - 146 pages
- ▶ manual.pdf - *SuSE 7.0 Installation, Networking, Know How* - 271pages

■ GA 2 - when???

■ Download

- ▶ From: <ftp://ftp.suse.com/pub/suse/s390/suse-us-s390>
- ▶ For IBMers (faster): <ftp://ftp3.linux.ibm.com/pub/SuSE/s390/en/>



Distributions - SuSE

- **First mainstream distributor**
- **Most complete in terms of packages**
 - ▶ 1068 binary RPMs
 - ▶ 689 source RPMs
- **LVM support**
- **YaST is a value add (some might disagree)**
- **/etc/rc.config is unique**
- **Run levels differ from LSB**
- **Installation details will follow**



Distributions - Turbolinux

- **"GA" v6.0 - February 2001**
 - ▶ 1 .iso CD image
 - ▶ Download from: <ftp://ftp.turbolinux.com/pub/product/s390>
 - ▶ For IBMers (faster): <ftp://ftp3.linux.ibm.com/pub/TurboLinux/s390/en/>
 - ▶ 2.2.16 kernel
 - ▶ 518 binary RPMs
- **Beta v6.5 - June 2001**
 - ▶ 1 .iso CD image
 - ▶ <ftp://ftp.turbolinux.com/pub/beta/s390/>
- **GA v6.5 - September 2001**
 - ▶ <http://www.turbolinux.com/products/s390/download.html>
 - ▶ LVM support, tape support





Distributions - Turbolinux

■ **Documentation on CD or Web**

- ▶ IBM_Install_Dec0400.pdf - *Preparing for Installing Turbolinux for S/390*
- ▶ IBM_Prep_Install_Dec0400.pdf - *Installing Turbolinux for S/390*
- ▶ TLS6zSeriesUserGuide.pdf - *Turbolinux for zSeries and S/390: User Guide*

■ **Installation details will follow**



Distributions - Red Hat

- **Public beta from May, 2001 based on 7.1**
- **No ISO images - must be recursively downloaded**
- **Download from either:**
 - ▶ <ftp://ftp.redhat.de/pub/s390>
 - ▶ <ftp://ftp.redhat.com/pub/redhat/linux/rawhide/s390>
- **Documentation**
 - ▶ Many "borrowed" manuals
 - ▶ None Red Hat specific
- **2.2.19 kernel (recent discussion of 2.4.9)**
- **Need tool such as wget to recursively download**
- **Basic setup tool simply named setup**



Distributions - Red Hat (cont'd)



■ Status (from linux-390 list server, Sept 13th)

...
ftp://ftp.redhat.de/pub/s390 and also
ftp://ftp.redhat.com/pub/redhat/linux/rawhide/s390/

are both updated to the newest versions and everything looks much more stable than any previous release. Also several big endian bugs have gone away etc.

If everything is going ok, I'll switch tomorrow from a 2.4.7-Torvalds based kernel to a standard Red Hat Linux kernel including all necessary s390 patches.

So definitely testing time, as the kernel can be updated separately. About updates from previous versions: *This is completely untested by us. It might work for you, it might give you a non-working system.*

For installation, you need the complete "Red Hat" directory and also the startup kernel/initrd from "images".
good luck,

— Florian La Roche

■ 96MB of storage (memory) is recommended during installation

- ▶ This does NOT include any swap space

@server

Distributions - Marist file system



■ Marist

- ▶ "Updated" version May 1, 2001- kernel, parm file, RAMdisk, large file system
- ▶ <ftp://linux390.marist.edu/pub/update>
- ▶ Differs - a compressed archive of a running file system
 - Pros: can get a system up easily
 - Cons: finding RPMs and upgrading can be difficult
- ▶ The format makes it more of a "Roll Your Own (RYO)" system
- ▶ 2.2.16 kernel
- ▶ Includes the **chkconfig** tool

@server



Distributions - Caiman from LinuxKorea

■ **Caiman v1.0**

- ▶ Public version May, 2001 from LinuxKorea
- ▶ Download: <ftp://linux390.linuxkorea.co.kr/caiman>
- ▶ 2.2.16 kernel
- ▶ Documentation: <ftp://linux390.linuxkorea.co.kr/caiman/doc/>
- ▶ Korean Language support
- ▶ LVM support
- ▶ Based on Debian distribution
 - ─ Only distribution of the 6 to be based on Debian
- ▶ Includes **dpkg** command
- ▶ Includes **apt-get** command
 - ─ Receive and install packages from the specified FTP server
- ▶ Includes **dselect** tool
 - ─ Graphical tool for managing packages
- ▶ Includes **alien** command
 - ─ Converts RPMs to "deb"s (dpkg format)



Distributions - Think Blue from Millenux

- **Based on Red Hat rawhide**
- **Only distribution with 64-bit kernel**
- **Founded by former employees of Thinking Objects**
- **ThinkBlue**
 - ▶ 31-bit
 - ▶ Download from <ftp://linux.zseries.org/pub/RedHat/RPMS/s390/>
- **Think Blue 64**
 - ▶ 64-bit v7.1 from Millenux Inc.
 - ▶ Download from <ftp://linux.zseries.org/pub/ThinkBlue64-7.1/>
 - ▶ 2.4.3 kernel
 - ▶ LVM support
 - ▶ ext3 support
 - ▶ Many RPMs available at <ftp://linux.zseries.org/pub/RedHat/RPMS/s390/>





Distributions - other

■ **Debian**

- ▶ There has been talk and discussion for a while
- ▶ Some are opposed to the OCO drivers

■ **LaPlace**

- ▶ Available to IBMers
- ▶ Latest system from IBM developers
- ▶ Similar to Marist - "a file system"

@server



Installation checklist

■ **Installation checklist**

- ▶ Prepare your environment LPAR or z/VM
 - LPAR prepare the IOCDS
 - z/VM prepare the directory entries for the Linux guests
- ▶ Plan the DASD environment
 - LPAR prepare the real DASDs
 - z/VM Minidisk or DASD
- ▶ Plan the Networking environment
- ▶ Complete the worksheet
- ▶ Booting the installation Linux system
 - Linux kernel
 - RAMdisk
 - Parameter file
- ▶ Complete the installation using the distributor's tool

@server



SuSE basics

■ 3 CDs

- ▶ Order from SuSE (150\$)
- ▶ Or download from their FTP site
 - <ftp://ftp.suse.com/pub/suse/S390/suse-us-s390/>
- ▶ Verify the iso images after download using MD5SUM from SuSE
- ▶ Using the ISO images simply “loopback” mount them as follows:

```
# mkdir /suse
# cd /suse
# mkdir cd1 cd2 cd3
# mount -o loop,ro /your location/suse-us-s390-cd1.iso /suse/cd1
# mount -o loop,ro /your location/suse-us-s390-cd2.iso /suse/cd2
# mount -o loop,ro /your location/suse-us-s390-cd3.iso /suse/cd3
```



SuSE - what is on the CDs

■ 3 CDs 1068 RPMs and 689 SRPMs in 46 series

- ▶ CD 1 - base 727 RPMs in 29 series
 - Linux Base System, Programs that don't need X , Development (C, C++, Lisp, etc.), Emacs, Games and more
 - all about graphics, Japanese packages, Desktop Environment, KDE applications, Network-Support (TCP/IP, UUCP, Mail, News)
 - Commercial Software, Perl modules, Security related software, Sound related stuff, Spell checking utilities and databases
 - Tcl/Tk/TclX, Tcl-Language and Tk-Toolkit for X, TeX/LaTeX and applications, Base X Window System - XFree86, 3D stuff for X11
 - X Applications, Development under X11, Development under X11, Window Manager and Desktop
- ▶ CD 2 - 308 RPMs and 439 SRPMs in 14 series
 - Development (C, C++, Lisp, etc.), Documentation, Games and more, GNOME - GNU Network Object Model Environment
 - All about graphics, Korean packages, Components for a SGML system, XView (OpenLook, Applications), Source packages
- ▶ CD 3 - 33 RPMs and 250 SRPMs in 3 series
 - TeX/LaTeX and applications, Base X Window System - XFree86, Source packages





SuSE Installation basics

■ Installation options and DASD / Directory requirement

Installation Type	Package Count	Space Required	Directory	Size
SuSE Almost Everything	834	2.72G	/bin	5.7 MB
			/boot	1.7 MB
			/dev	28 KB
SuSE Development System	319	1.18G	/etc	3.6 MB
			/lib	12 MB
SuSE DMZ Base System	96	309.1M	/opt	31 MB
SuSE Minimum System	75	160.4M	/root	20 KB
			/sbin	8.2 MB
SuSE Network Oriented System	337	918.1M	/tmp	4 KB
			/usr	654 MB
			/var	19 MB
SuSE Default System with Server Applications and Development	244	784.3M	Total	736 MB



SuSE installation worksheet

Data	Value
VM Guest user ID	
VM guest password	
Disk with FTP program	
Devices	
First DASD (root file system)	
Second DASD (swap space)	
Tape unit (not needed when installing from VM reader)	
Network data	
Host name	
IP address	
Net mask	
Broadcast address	
Gateway address	
IP address of the DNS	
DNS search domain	
NFS, FTP or SMB server data	
IP address or hostname	
Path to CD-ROM data	
Linux data	
Root password	



SuSE installation - FTPing from VM



```
ftp 9.12.0.88
...
ftp> cd /suse/cdl
ftp> locsite fix 80
ftp> cd suse/images
ftp> get parmfile suse70.parmfile
150 Opening BINARY mode data connection for 'parmfile' (38 bytes).
ftp> bin
ftp> get vmrdr.ikr suse70.kernel
150 Opening BINARY mode data connection for 'vmrdr.ikr' (1511884 bytes).
ftp> get initrd suse70.ramdisk
150 Opening BINARY mode data connection for 'initrd' (9867061 bytes).
ftp> quit
TYPE SUSE70 PARMFILE
ramdisk_size=32768 root=/dev/ram0 ro
```

@server

SuSE installation - FTPing from VM (cont'd)



```
FILEL SUSE70 *
SUSE70  RAMDISK  A1 F      80      123339      2409  9/06/01  6:31:18
SUSE70  KERNEL   A1 F      80      18899       338  9/06/01  6:29:46
SUSE70  PARMFILE A1 F      80           1         1  9/06/01  6:26:17

TYPE SUSE70 EXEC
/* REXX LOAD EXEC FOR SUSE */
SAY 'LOADING FILES FOR SUSE 7.0 INTO READER...'
'CP CLOSE RDR'
'CP PURGE RDR CLASS L'
'CP SPOOL PUN * RDR CLASS L'
'PUNCH SUSE70 KERNEL A (NOH'
'PUNCH SUSE70 PARMFILE A (NOH'
'PUNCH SUSE70 RAMDISK A (NOH'
'CP SPOOL PUN * RDR CLASS A'
'CP SPOOL RDR KEEP CLASS L'
'CP IPL 00C CLEAR'
```

@server

SuSE installation - boot installation system



suse70

Linux version 2.2.16 (root@ikr_rdr.suse.de) (gcc version 2.95.2 19991024 (release)) #1 SMP Tue May 1 11:44:38 GMT 2001

Command line is: ramdisk_size=32768 root=/dev/ram0 ro

We are running under VM

This machine has an IEEE fpv

Initial ramdisk at: 0x02000000 (10132480 bytes)

Detected device 2928 on subchannel 0000 - PIM = 80, PAM = 80, POM = FF

Detected device 2929 on subchannel 0001 - PIM = 80, PAM = 80, POM = FF

Detected device 0191 on subchannel 0002 - PIM = F0, PAM = F0, POM = FF

.....

== Welcome to SuSE Linux S/390 7.0

First, select the type of your network device:

- 0) no network
 - 1) OSA Token Ring
 - 2) OSA Ethernet
 - 3) OSA-Express Gigabit Ethernet
 - 4) Channel To Channel
 - 5) Escon
 - 6) IUCV
 - 7) CLAW (Cisco Mainframe Channel Connection)
- Enter your choice (1-7): **2**

@server

SuSE installation - network setup (cont'd)



To set up the network, you have to read and confirm the license information of the network device driver provided by IBM.

Do you want to see the license (Yes/No) ? **yes**

...

Do you agree with this license (Yes/No) ? **yes**

Ok, now we can set up the network configuration.

Please enter the device address of the network device

Ask your system administrator for the correct settings.

If there is only ONE OSA network device attached to your machine, you may type "auto" for automatic detection; use this option carefully.

Network device address (e.g. FC20): **2928**

Please enter the relative port number on device address 2928

Relative port: **0**

@server

SuSE installation - network setup (cont'd)



Please enter your full host name (e.g. s390.suse.com): **vmlinux5.itso.ibm.com**
Please enter your IP address: **9.12.6.76**
Please enter the net mask: **255.255.255.0**
Please enter the broadcast address (9.12.6.255):
Please enter the gateway address: **9.12.6.75**
Please enter the IP address of the DNS server (leave blank for none): **9.12.14.7**
Please enter the DNS search domain (e.g. suse.com): **itso.ibm.com**
Please enter the MTU (Maximum Transfer Unit, leave blank for default: 1492):
Configuration for eth0 will be:
Full host name : vmlinux5.itso.ibm.com
IP address : 9.12.6.76
Net mask : 255.255.255.0
Broadcast address: 9.12.6.255
Gateway address : 9.12.6.75
DNS IP address : 9.12.14.7
DNS search domain: itso.ibm.com
MTU size : 1492
Is this correct (Yes/No) ? **yes**
...
Please enter the temporary installation password: *********

@server

SuSE Installation - start networking and inetd



Temporary installation password set.
restarting syslogd:
Jun 13 11:24:32 vmlinux5 syslogd 1.3-3: restart.
ifconfig eth0 9.12.6.76 netmask 255.255.255.0 broadcast 9.12.6.255 mtu 1492
/sbin/ifconfig eth0 :
eth0 Link encap:Ethernet HWaddr 00:06:29:6C:CB:CE
inet addr:9.12.6.76 Bcast:9.12.6.255 Mask:255.255.255.0
UP BROADCAST RUNNING MULTICAST MTU:1492 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:100

Trying to ping my IP address:
...
Trying to ping the IP address of the Gateway:
...
Trying to ping the IP address of the DNS Server:
...
Network Setup finished, running inetd...
#cp disc

@server

SuSE Installation - telnet and YaST to finish



Welcome to SuSE Linux 7.0 (s390) - Kernel 2.2.16 (ttyp0).

vmlinux5 login: **root**

Password:

>>> >>> >>> >>> >>> >>> SuSE Linux S/390 7.0 <<< <<< <<< <<< <<< <<<


1. If you want to check which devices the dasd driver can see,
run 'insmod dasd probeonly' (may take long) and check
the output of 'cat /proc/dasd/devices'.
Remove the dasd driver with 'rmmod dasd' afterwards.

2. Choose the device numbers you want to use for SuSE Linux S/390
!!! BE CAREFUL WHEN SELECTING DASDs - !!!
!!! YOU MAY DESTROY DATA ON SHARED DEVICES !!!

3. Enter 'insmod dasd dasd=<list of devices>'
Remember to separate devices by commas (<dev_no>,<dev_no>),
syntax for ranges is <from_dev_no>-<to_dev_no>
like
'insmod dasd dasd=FD00-FD0F,FD40,FD42,FD80-FD86'

Note: When updating, you have to load the dasd driver with the
same DASD order as in the installed system - see documentation
for further information.

4. Start installation or update with 'YaST'.



SuSE installation - formatting DASD devices



SuSE Instsys vmlinux5:/root # **insmod dasd dasd=201-206**

Using /lib/modules/2.2.16/block/dasd.o

SuSE Instsys vmlinux5:/root # **cat /proc/dasd/devices**

0201(ECKD) at (94:0) is dasda:active at blocksize: 4096, 468000 blocks, 1828 MB
0202(ECKD) at (94:4) is dasdb:active at blocksize: 4096, 468000 blocks, 1828 MB
0203(ECKD) at (94:8) is dasdc:active at blocksize: 4096, 18000 blocks, 70 MB
0204(ECKD) at (94:12) is dasdd:active at blocksize: 4096, 36000 blocks, 140 MB
0205(ECKD) at (94:16) is dasde:active at blocksize: 4096, 36000 blocks, 140 MB
0206(ECKD) at (94:20) is dasdf:active at blocksize: 4096, 36000 blocks, 140 MB

We format three volumes to be used as follows:

201 /dev/dasda	boot volume
202 /dev/dasdb	home directories
203 /dev/dasdc	swap space

We format the first volume and make an ext2 file system

SuSE Instsys vmlinux5:/root # **dasdfmt -f /dev/dasda -b 4096**

I am going to format the device /dev/dasda in the following way:

Device number of device : 0x201

...

Type "yes" to continue, no will leave the disk untouched: **yes**

[SuSE Instsys] vmlinux5:/root # **mke2fs -b 4096 /dev/dasda1**

We repeat for the other two volumes, then run YaST

SuSE Instsys vmlinux5:/root # **dasdfmt -f /dev/dasdb -b 4096**

SuSE Instsys vmlinux5:/root # **dasdfmt -f /dev/dasdc -b 4096**

SuSE Instsys vmlinux5:/root # **yast**



SuSE Installation - YaST

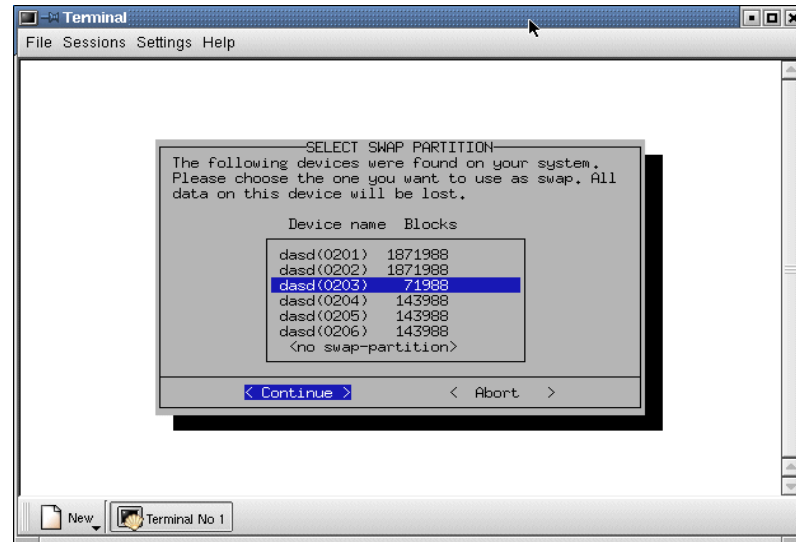


On initial three panels we select:

"English",

"Installation from an FTP site"

"Install Linux from scratch"



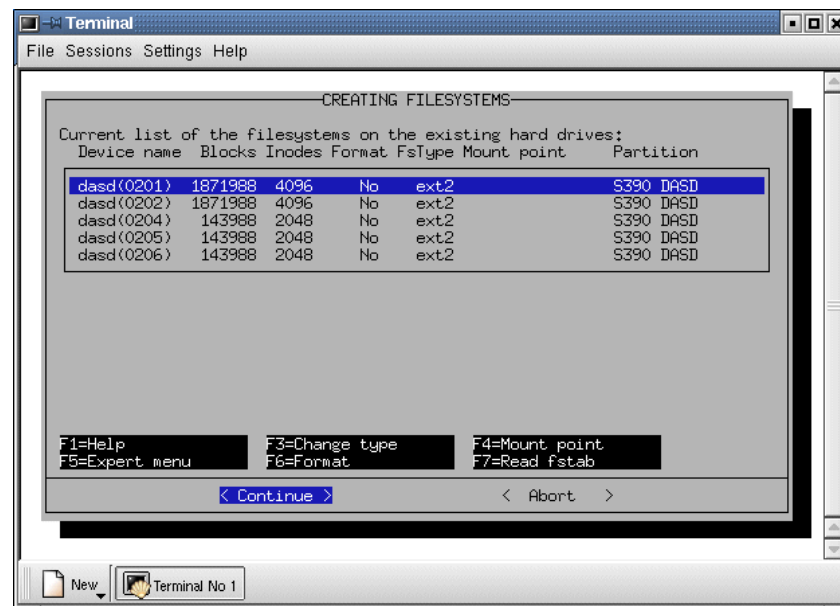
@server

SuSE Installation - YaST Creating Filesystems



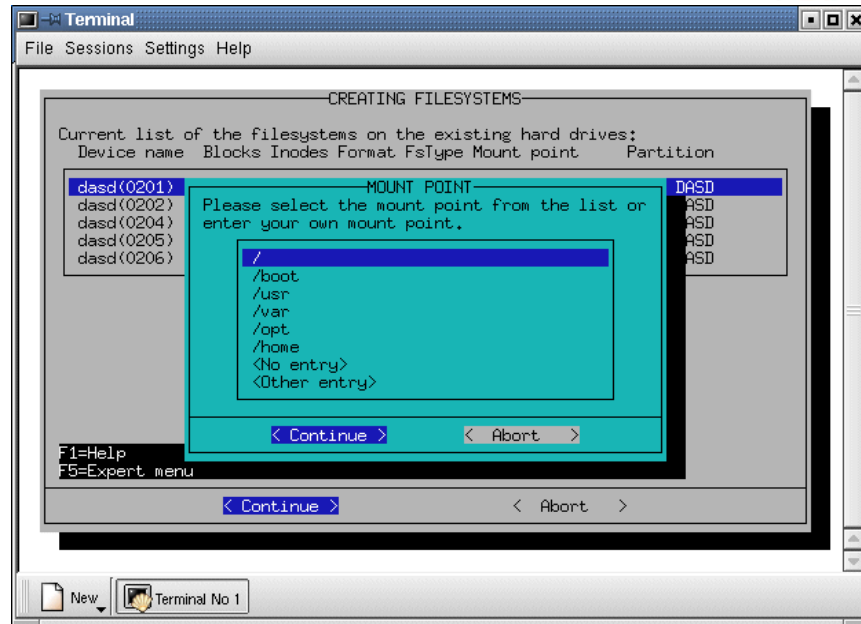
On Partition Harddrives panel, we select:

"Do not partition"




@server

SuSE Installation - YaST mount point panel



@server

SuSE installation - YaST - FTP settings panel

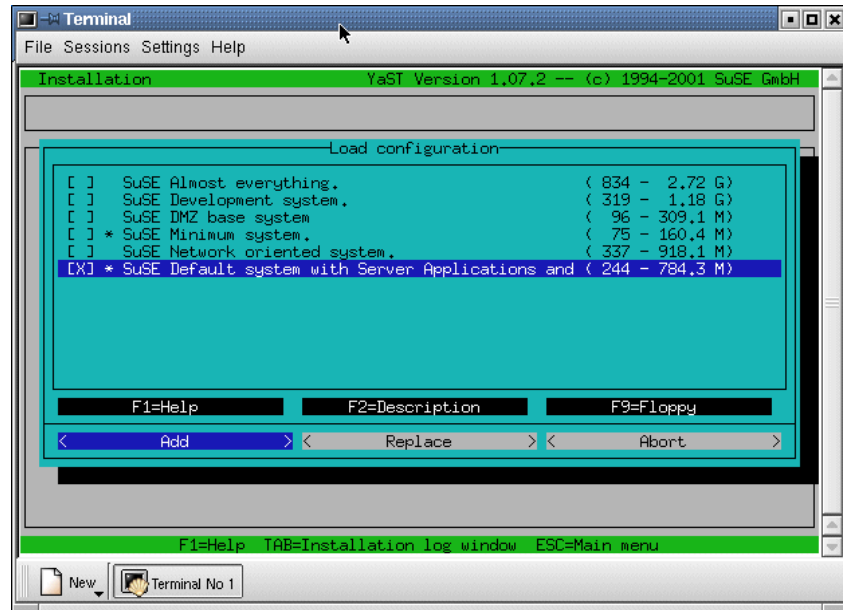


@server


SuSE installation - YaST Load configuration

On Installation menu we select:

Load Configuration



SuSE installation - YaST installation menu



SuSE installation - YaST installation log



```
Terminal
File Sessions Settings Help

Installation YaST Version 1.07.2 -- (c) 1994-2001 SuSE GmbH

INSTALLATION COMPLETE. <TAB> brings you to the Installation log window.

Logfile:
sdb
150 Opening BINARY mode data connection for "/suse/cd2/suse/doc3/sdb_en.rpm"
(2125889 bytes).
sdb_en
150 Opening BINARY mode data connection for "/suse/cd2/suse/doc3/susehlf.rpm"
(556000 bytes).
susehlf
150 Opening BINARY mode data connection for "/suse/cd2/suse/doc3/xf86html.rpm"
(4254595 bytes).
xf86html
150 Opening BINARY mode data connection for "/suse/cd2/suse/doc3/xman.rpm"
(1378122 bytes).
xman
Closed connection to [9.12.9.183]

Totally installed: 244

Base system: SuSE-Linux-US-S390 7.0.0-0

INSTALLATION COMPLETE.

F1=Help TAB=Installation log window ESC=Main menu
```

@server

SuSE Installation - packages installed



```
Terminal
File Sessions Settings Help

Installation YaST Version 1.07.2 -- (c) 1994-2001 SuSE GmbH

Logfile:
150 Opening BINARY mode data connection for "/suse/cd2/suse/doc3/sdb_en.rpm"
(2125889 bytes).
sdb_en
150 Opening BINARY mode data connection for "/suse/cd2/suse/doc3/susehlf.rpm"
(556000 bytes).
susehlf
150 Opening BINARY mode data connection for "/suse/cd2/suse/doc3/xf86html.rpm"
(4254595 bytes).
xf86html
150 Opening BINARY mode data connection for "/suse/cd2/suse/doc3/xman.rpm"
(1378122 bytes).
xman
Closed connection
Totally installed: 2

Base system: SuSE-Li

INSTALLATION COMPLETE.
INSTALLATION COMPLETE.

F1=Help TAB=Installation log window ESC=Main menu
```

@server



SuSE Installation - YaST panels

On the **SELECT KERNEL** panel we choose:

Default kernel for S/390 (with support for tape IPL)

On the **TIME ZONE CONFIGURATION** panel we choose

EST5EDT

On the **ADJUSTMENT OF HARDWARE CLOCK** panel we choose

Local time

On the **ENTER YOUR HOSTNAME** panel we take the defaults

For the question **Do you want to use TCP/IP in loopback mode only?** we choose

Real network

For the question **Do you want to use the machine as DHCP client?** we choose

No



SuSE Installation - YaST network address

ENTER THE NETWORK ADDRESSES

Please enter the data required for the configuration of your network. These are the IP address you want to give the machine currently being installed (e.g. 192.168.17.42) and the netmask of your network. The latter is 255.255.255.0 for most of the (smaller) networks, but you may wish to set it to a different value. If you need a gateway to access the server, please enter the IP address of the gateway host.

Type of network: [eth0]

IP address of your machine: :9.12.6.76:

Netmask (usually 255.255.255.0): :255.255.255.0:

Default gateway address (if required): :9.12.6.75:

IP address of the Point-to-Point partner:

Maximal Transfer Unit - MTU :1492:

(in doubt leave field empty)

< Continue > < Abort >





SuSE installation - YaST panels

On the **START INETD?** panel we choose

Yes

On the **START THE PORTMAPPER?** panel we choose

Yes

On the **START NFS-SERVER?** panel we choose

Yes

On the **ADJUST NEWS FROM-ADDRESS** panel we take the default

For the question **Do you want to access a nameserver?** we choose

Yes

On the **NAMESERVER CONFIGURATION** panel we take the defaults

On the **SENDMAIL CONFIGURATION** panel we choose

Host with permanent network connection (SMTP)



SuSE Installation - YaST - SuSEconfig runs

```
Terminal
File Sessions Settings Help

OUTPUT of SuSEconfig
Started the SuSE-Configuration Tool.
Running in full featured mode.
Reading /mnt/etc/rc.config and updating the system...
Installing new /etc/HOSTNAME
Installing new /etc/resolv.conf
Installing new /etc/nntpserver
Installing new /etc/inews_mail_gateway
Installing new /var/lib/news/mailname
Installing new /var/lib/news/whoami
Installing new /etc/SuSEconfig/profile
Installing new /etc/SuSEconfig/csh.cshrc

< Continue >
```



SuSE installation - booting new system



[SuSE Instsys] vmlinux5:/root # **shutdown -h now**

...

From VM:

ipl 201 clear

The start of our newly install system is announced by the following initial set of messages:

Linux version 2.2.16 (root@Tape.suse.de) (gcc version 2.95.2 19991024 (release))

#1 SMP Sun May 6 06:15:49 GMT 2001

Command line is: ro dasd=0203,0201,0202 root=/dev/dasdb1 noinitrd

We are running under VM

This machine has an IEEE fpu

Initial ramdisk at: 0x02000000 (16777216 bytes)

Detected device 2928 on subchannel 0000 - PIM = 80, PAM = 80, POM = FF

.....

Welcome to SuSE Linux

Started the SuSE-Configuration Tool.

.....

Finished.



SuSE Installation - booting new system (cont'd)



Now scripts have to be started. They will be started in one minute. You can find a log file under /var/log/Config.bootup.

It will also be printed on console 9.

You can now already use your system. If you shut down the system before the scripts are finished, they are executed again at the next system startup.

Press <RETURN> to continue...

Have a lot of fun!

Your SuSE Team

All installation tasks have been performed when you see the login prompt:

Welcome to SuSE Linux 7.0 (s390) - Kernel 2.2.16 (0).

vmlinux2 login:



Turbolinux installation worksheet



1. VM Guest User ID (if installing under VM)	
2. VM Guest Password (if installing under VM)	
3. Device: 1st DASD/vdev(VM) (/ file system)	
4. Device: 2nd DASD/vdev(VM) (swap space)	
5. Device: Tape Unit *optional (e.g. 0181)	
6. Network Device 1st of Pair (e.g. 292C)	
7. Network Device Port 1st or pair (e.g. 0)	
8. Network Device Type (e.g. OSA-E)	
9. Host name (e.g. vmlinux.itso.ibm.com)	
10. Host IP Address (e.g. 9.12.6.73)	
11. Netmask (e.g. 255.255.255.0)	
12. Network IP Address (e.g. 9.12.6.0)	
13. Broadcast IP Address (e.g. 9.12.6.255)	
14. Gateway IP Address (e.g. 9.12.6.75)	
15. DNS IP Address (e.g. 9.12.14.7)	
16. MTU Size (e.g. 1492)	
17. FTP Server IP Address or hostname	
18. FTP path (e.g. /mnt/cdrom)	
19. FTP User ID (e.g. TOT82)	
20. FTP Password	
21. Linux Root User Password	



Turbolinux - DASD/directory requirements



For the "all servers" (5) installation option:

/bin	5.9 MB
/boot	3.3 MB
/dev	122 KB
/etc	9.8 MB
/lib	6.6 MB
/opt	0 KB
/root	400 KB
/sbin	4.4 MB
/tmp	36 KB
/usr	889 MB
/var	15 MB
total	986 MB



Turbolinux - installation



Boot the installation Linux system

- Complete the network questions
- Ensure the network comes up
- Telnet in to continue with the installation script.

Welcome to
TUBOLINUX
for IBM S/390 and zSeries
Astrodon Release
brought to you by the
FRONTIER TEAM
October 31 2000

Do you want to configure the eth0 device? (y/n) **[y]**:

@server

Turbolinux installation



The installation script - install.pl

login: **root**

[root@vmlinux7 /root]# **install.pl**

The next question will be about which DASD are to be made available to Linux

Which DASD's do you want to have available in Linux?

Example: 192, 194-196

Example: fd02, fd04-fd06

DASD's: **201,203**

insmod dasd_mod dasd=201,203

Using /lib/modules/2.2.16/misc/dasd_mod.o

One or more dasds need low level formatting

	Number	Type	Name	Status
1	0201	ECKD	dasda	n
2	0203	ECKD	dasdb	n

Format disk[# or q]: **1**

dasdfmt -b 4096 -f /dev/dasda

...

Type "yes" to continue, no will leave the disk untouched: **yes**

@server



Turbolinux installation script (cont'd)

Choose a swap device

	Number	Type	Name	Status
1	0201	ECKD	dasda	n
2	0203	ECKD	dasdb	n

Swap_disk[#_or_q]: **2**

mkswap /dev/dasdb1

init_swap: Setting up swapspace version 1, size = 820576256 bytes

chmod 600 /dev/dasdb1

swapon /dev/dasdb1

Which disk should be the root filesystem?

...

Root_disk[#]: **1**

Creat_a_new_filesystem_on_/dev/dasda1?_[y/n]: **y**

Please choose an install method:

1) NFS

2) FTP

3) quit

Which_method? **2**

FTP_hostname: **itsolinux2.itso.ibm.com**

FTP_path: **/mnt/cdrom**

@server



Turbolinux installation script (cont'd)

- 1 Web server
- 2 FTP server
- 3 SMB file server
- 4 Web/Ftp Proxy Server
- 5 All servers
- 6 Development (All-in-one)

Package selection: **6**

1/515: setup-6.--10.noarch.rpm

2/515: filesystem-1.3.5-3.noarch.rpm

3/515: ldconfig-1.9.5-17.s390.rpm

...

515/515: libungif-devel-4.1.0-3.s390.rpm

Please select your time zone

0 - Africa/Abidjan

1 - Africa/Accra

2 - Africa/Addis_Ababa

3 - Africa/Algiers

...

Timezone ([Enter] for more options) :

@server

Turbolinux installation script (cont'd)

Add another user because root cannot telnet by default

Do_you_want_to_add_a_new_user?_(Y/n) **Y**

User_Name:_[neo]: **linuser**

Please set a password for linuser:

New UNIX password: *********

Retype new UNIX password: *********

passwd: all authentication tokens updated successfully

After adding a new user and password there is a prompt for the password to be used for root's password.

Please enter a password for the root user:

Changing password for user root

New UNIX password: *********

Retype new UNIX password: *********

passwd: all authentication tokens updated successfully



Turbolinux installation script (cont'd)

SILO Configuration

	Number	Type	Name	Status	Blksize	Blocks	FS
1	0201	ECKD	dasda	active	4096	601020	/
2	0203	ECKD	dasdb	active	4096	200340	swap

Boot_Disk_[#] **1**

chroot /mnt/root silo -f /boot/image-2.2.16-122-tape -d /dev/dasda -p /boot/parm -b /boot/ipleckd.boot -t2

Silo: o->image set to /boot/image-2.2.16-122-tape

Silo: o->ipldevice set to /dev/dasda

Silo: o->parmfile set to /boot/parm

Silo: o->bootsect set to /boot/ipleckd.boot

Silo: o->Testonly flag is now 0

Silo: o->Testlevel is set to 0

Silo: o->IPL device is: '/dev/dasda'

Silo: o->bootsector is: '/boot/ipleckd.boot' ...ok...

Silo: o->bootmap is: '/boot/image-2.2.16-122-tape' ...ok...

Silo: o->Kernel image is: '/boot/ipleckd.boot' ...ok...

Silo: o->original parameterfile is: '/boot/parm' ...ok...final parameter parm is: '/boot/parm' ...ok...

...





Turbolinux installation - final step

The installation script v6.5 beta final step

Installation is complete.
Type halt to shutdown the system. Do not use any other commands.
Then IPL the root disk.
If you are on VM use #cp i xxx where xxx is boot disk number.

bash-2.04# **shutdown -h now**

@server



RedHat basics

■ **RedHat Beta code**

- ▶ Download from one FTP site:
 - <ftp://ftp.redhat.com/pub/redhat/linux/rawhide/s390/>
 - <ftp://ftp.redhat.de/pub/s390/>
- ▶ Needs about 880MB Disk space without the source RPMs
- ▶ No ISO images available today

■ **Documentation on there FTP site**

- ▶ <ftp://ftp.redhat.de/pub/s390/docu/>

@server



RedHat - space required

Directory	default	full
/bin	5.4 MB	6 MB
/boot	1.7 MB	1.7 MB
/dev	104 KB	104 KB
/etc	8.2 MB	12 MB
/lib	17 KB	17 MB
/opt	0	0
/root	32 KB	44 KB
/sbin	5.1 MB	9 MB
/tmp	140 KB	288 KB
/usr	630 MB	2300 MB
Total:	667 MB	2347 MB



Red Hat parameter file

■ Parameter file inst.parm:

```
root=/dev/ram0 ro ip=off DASD=200-20f
HOST=vmlinux4.itso.ibm.com:eth0:9.12.6.80:1500
NETWORK=9.12.6.0:255.255.255.0:9.12.6.255:9.12.6.75
LCS=0x2926,0
DNS=9.12.14.7
SEARCHDNS=itso.ibm.com:ibm.com
RPMSEVER=ftp://9.12.2.120
MOUNTS=/dev/dasdb1: /, /dev/dasdc1: /usr
SWAP=/dev/dasddl
ROOTPW=root11
INSTALL=default
```



RedHat installation



CP IPL 00C CLEAR

0000003 FILES CHANGED

Linux version 2.2.19-0.07BOOTvrd (laroche@rawhide.redhat.de) (gcc version 2.95. 3 20010319 (prerelease Red Hat Linux S/390)) #1 SMP Tue Apr 24 15:31:32 EDT 2001

```
Command line is: root=/dev/ram0 ro ip=off DASD=200-20f
                HOST=vmlinux4.itso.ibm.com:eth0:9.12.6.80:1492
                NETWORK=9.12.6.0:255.255.255.0:9.12.6.255:9.12.6.75
                LCS=0x2926,0
                DNS=9.12.14.7
                SEARCHDNS=itso.ibm.com:ibm.com
                RPMSERVER=ftp://9.12.2.120/
                MOUNTS=/dev/dasdbl:/dev/dasdc1:/usr
                SWAP=/dev/dasddl
                ROOTPW=root11
                INSTALL=default
```

We are running under VM

This machine has an IEEE fpu

Initial ramdisk at: 0x02000000 (3586960 bytes)

Detected device 2926 on subchannel 0000 - PIM = 80, PAM = 80, POM = FF

Detected device 2927 on subchannel 0001 - PIM = 80, PAM = 80, POM = FF

Detected device 0191 on subchannel 0002 - PIM = F0, PAM = F0, POM = FF

Detected device 0201 on subchannel 0003 - PIM = F0, PAM = F0, POM = FF

...



Red Hat installation



■ If you chose not to completely specify all the parameters in inst.parm/lpar.prm, answer the questions asked by the setup script:

- ▶ Domain name
- ▶ Communication device
 - CTC
 - IUCV
 - eth0
- ▶ IP address
- ▶ Network mask
- ▶ Broadcast address
- ▶ Network address
- ▶ Default gateway IP address



RedHat installation



Once networking is running, you should see something similar to:

```
eth0  Link encap:Ethernet  HWaddr 00:06:29:6C:CB:CE
       inet addr:9.12.6.80  Bcast:9.12.6.255  Mask:255.255.255.0
       UP BROADCAST RUNNING MULTICAST  MTU:1492  Metric:1
       RX packets:0 errors:0 dropped:0 overruns:0 frame:0
       TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
       collisions:0 txqueuelen:100
```

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
127.0.0.1	0.0.0.0	255.255.255.255	UH	0	0	0	lo
9.12.6.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
9.12.6.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
0.0.0.0	9.12.6.75	0.0.0.0	UG	0	0	0	eth0

Starting portmap.

Starting telnetd to allow login over the network.

Please telnet now to 9.12.6.80 and start 'rhsetup'!

Bringing up debugging shell in 3270...

/bin/sh: No controlling tty (open /dev/tty: No such device or address)

/bin/sh: warning: won't have full job control



RedHat installation

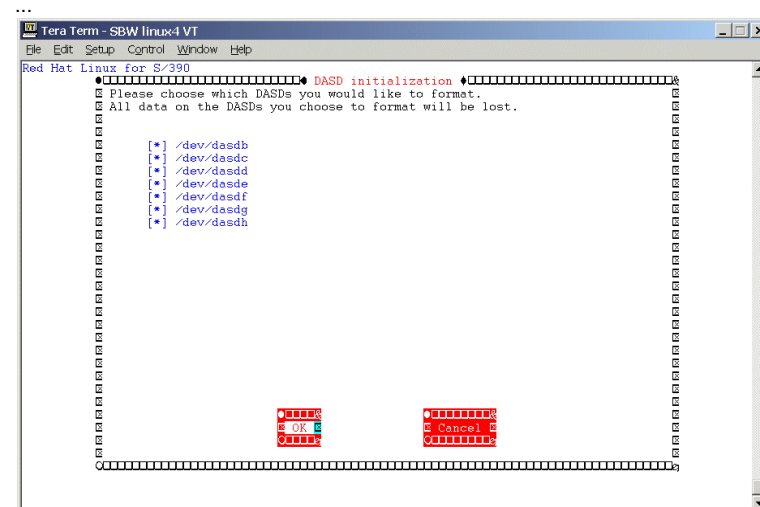


Linux 2.2.19-0.07BOOTvdr ((none)) (tty0)

Welcome to Red Hat Linux on S/390.

Please run 'rhsetup' to start the installation.

rhsetup



RedHat installation



Tera Term - SBW linux4 VT

File Edit Setup Control Window Help

Red Hat Linux for S/390

Mount points

Please choose the mount points for your DASDs.

dasdb1: /_____

dasdc1: /usr_____

dasdd1: _____

dasde1: _____

dasdf1: _____

dasdg1: _____

dasdh1: _____

OK Cancel



RedHat installation



Tera Term - SBW linux4 VT

File Edit Setup Control Window Help

Red Hat Linux for S/390

Package source

Please enter the full path to the Red Hat Linux for S/390 RPMs. This can be an FTP, HTTP or NFS URL.
(Examples: ftp://ftp.local.com/pub/s390, http://www.local.com/s390, nfs.local.com:/mnt/s390)

ftp://9.12.2.120

OK Cancel



Another format: ftp://userid:password@ftp.server/path/to/packages



RedHat installation

- **Installation proceeds - you will see progress of packages as they are installed**
- **Watch the informational messages for errors**
 - zilo does not always succeed
- **When it does succeed, do a "shutdown -h now"**

@server



Notes:

@server



Agenda - section 5 of 9

- **Introduction**
- **Brief history and level set**
- **zSeries platform options**
 - ▶ Break
- **Distributions (SuSE, RedHat, Turbolinux, others)**
- ★ **Possible scenarios and solutions**
 - ▶ Lunch
- **Linux for zSeries in the enterprise**
- **Miscellaneous topics**
 - ▶ Break
- **Sizing and TCO explained**
- **Summary, future, Q&A**



Possible scenarios and solutions

- **From a solution point of view**
 - ▶ Web serving
 - ▶ Mail serving
 - ▶ DNS serving
 - ▶ File and print serving
 - ▶ Database
 - ▶ SAP R3
 - ▶ Firewall
- **From a software point of view**
 - ▶ ISV software
 - ▶ IBM middleware




Solutions - Web serving



■ Apache

- ▶ The most common Web server in the world
- ▶ Most distributions install with Apache running by default
- ▶ For SuSE, /etc/httpd/httpd.conf is the configuration file
- ▶ For SuSE, /usr/local/httpd/htdocs is the default "Server root" (data directory)



Source: Netcraft - <http://www.netcraft.com/survey/>



Solutions - Mail serving



■ Sendmail

- ▶ High-performance solutions that scale as your business grows
 - Corporate e-mail systems are **breaking** due to load
 - Scale to up to 2 million users on a zSeries
 - 80 of the Fortune 100 companies use Sendmail
- ▶ **Sendmail Advanced Message Server** is a complete, rapidly deployed messaging solution that hosts multiple domains and scales comfortably to hundreds of thousands of users on a single server
- ▶ Sendmail's streamlined, efficient approach to Internet messaging architecture gives you a solid foundation to grow and add services - and is backed by Sendmail's expertise
- ▶ An article in *Network World Fusion News*:
 - <http://www.nwfusion.com/news/2001/0910apps.html>
 - "However, companies are now finding the centralized computing model of mainframes more attractive for a variety of applications."





Solutions - Mail serving (cont'd)

■ **Bynari Insight**

- ▶ Cost effective, reliable, simple and scalable messaging and collaboration server
- ▶ Can replace the need for a Microsoft NT/Exchange Server, supporting functionality available in Outlook
 - Shared folders
 - Access to global address books
 - Calendar and meeting management
 - Without the high costs of Client Access Licenses
- ▶ Provides IMAP, POP3 and SMTP mail protocols

■ **"The Truth Behind the Great Server Heist", Stephen E. Harris,**

- ▶ <http://consultingtimes.com/Serverheist.html>
- ▶ Quote: "The Exchange/Intel solution, again, is \$ 8.65 per user per month. The Linux groupware solution now comes out to a mere \$ 2.02 (compared to the original estimate of \$16.73)."



Solutions - DNS serving

■ **Berkeley Internet Name Domain (BIND) v8**

- ▶ Is the de-facto DNS server
- ▶ Is shipped with most distributions

■ **From Linux-390, Sept 21, 2001**

Use bind v8 on a couple of small, separate Intel boxes. We have customers hosting thousands of domains, and it's the backbone of the production Internet. It's on your Linux distribution, and is usually the default nameserver. If you need the MS WINS perversion, couple bind with Samba, which provides WINS function for the MS clients. BIND has DDNS support, so the Win2k people can do their thing (finally a normal name service from MS!).

The small separate machines recommendation is such that for the rare times when you do need to IPL, the service will continue. DNS is probably the most critical non-packet-transport Internet service; pretty much all the rest of the services depend on it to work correctly. By setting up several small machines on different networks, your network can survive most things that require a reboot of the nameserver without a hiccup.





Solutions - File and print serving

■ **Samba**

- ▶ Open implementation of Common Internet File System (CIFS) which is actually NT SMB
- ▶ Uses the NetBIOS over TCP/IP (NBT) protocol
- ▶ Password encryption - maintaining the /etc/smbpasswd file:
 - ─ Adds a level of maintenance
 - ─ Avoids any additional work on the client side
- ▶ Samba Web Admin Tool (SWAT) - Browser interface for system administration
- ▶ Present on multiple Linux and UNIX systems and architectures
- ▶ Samba 2.2
 - ─ Windows 2000 primary domain server support
 - ─ Integration of Samba instances into existing Win95/98, Windows NT, Windows 2000 domains
 - ─ Authenticate Samba users to remote Windows domains
 - ─ Access Control List (ACL) support using only Samba tools, Windows system not required to manage ACL



Solutions - Database

■ **IBM DB2**

- ▶ DB2 Connect
- ▶ DB2 UDB
- ▶ DB2 Intelligent Miner Scoring
- ▶ DB2 Net Search Extender

■ **Oracle - is it coming?**

■ **Open Source**


- ▶ MySQL
- ▶ Postgres SQL



DB2 Connect Solutions - DB2 Connect




Before



Eliminate middle tier with DB2 Connect

After



@server

Solutions - Oracle



■ Oracle9i Database for Linux/390

■ Is it coming?

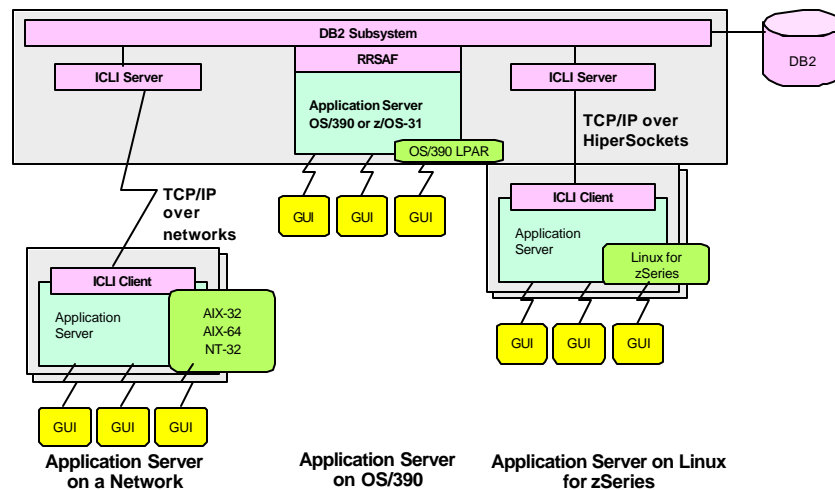
■ From <http://otn.oracle.com/tech/linux/htdocs/devrelease.html>

- ▶ "Oracle has made public its intention to deliver a developer release of Oracle9i Database for the Linux/390 environment."
- ▶ "If you are interested in evaluating Oracle9i Database in this environment and in providing feedback to Oracle, please visit us later to register and complete the accompanying survey."
- ▶ "Should there be a significant demand for Oracle9i Database on Linux/390, Oracle can quickly and easily offer our flagship database at a consistent release level to that available today on Linux/Intel."

@server

Solutions - SAP R3

- Database server on z/OS using DB2
- Application server on Linux for zSeries
- Connection via HiperSockets




Solutions - Firewall

- Packet filtering is a common method of implementing a firewall
 - ▶ Unwanted TCP/IP packets are simply "dropped"
- We will look at different firewall scenarios
 - ▶ Single router/firewall with one server
 - ▶ Single router/firewall with DMZ
 - ▶ Single router/firewall with more subnets
 - ▶ Two layer router/firewall implementation



Packet filtering

- Single router/firewall with one server

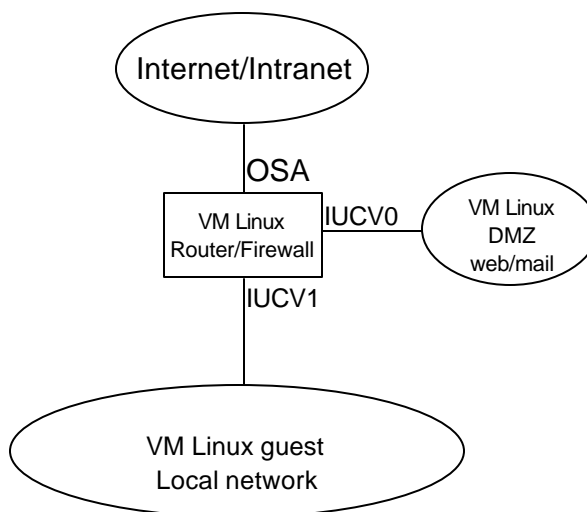


@server



Packet filtering (cont'd)

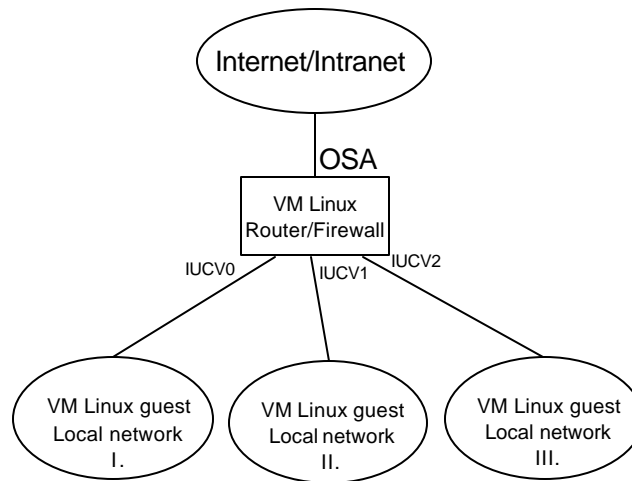
- Single router/firewall with DMZ



@server

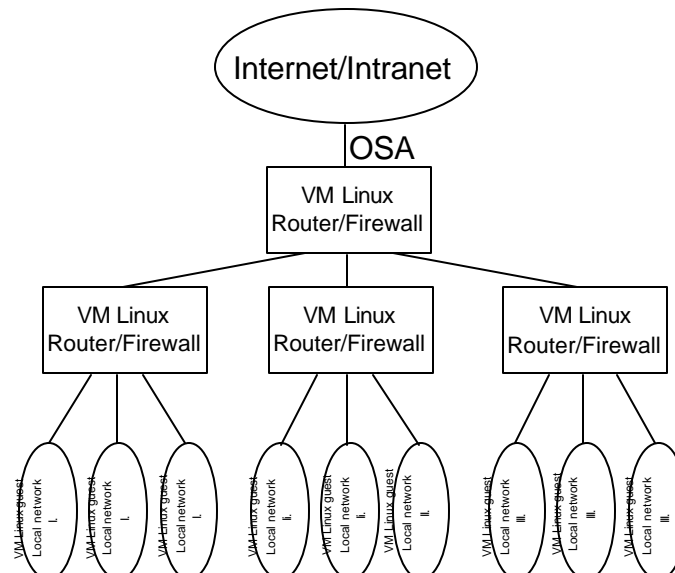
Packet filtering (cont'd)

■ Single router/firewall with more subnets



Packet filtering (cont'd)

■ Two layer router/firewall implementation





Network address translation

■ Definition

- ▶ Network Address Translation (NAT) is used to translate the source or destination address of the IP packets traveling through the gateway
- ▶ Can be used to conserve IP addresses inside a firewall

■ Basic types of Network Address Translation (NAT)

- ▶ Source NAT (SNAT)
 - Source address of the first packet is modified. SNAT is always done after the routing, just before packet is going out of the router/firewall.
- ▶ Destination NAT (DNAT)
 - The destination address of the first packet is modified. DNAT is always done before routing, just after the packet is coming into the router/firewall.
- ▶ Masquerading
 - A specialized version of SNAT.
 - Usually used for dynamically assigned IP addresses, for example dialups.
 - Source address is not needed - the source address of the interface the packet is going out on will be used



Software - ISV applications

■ ACTS

- ▶ Testing tools

■ Aeonware

- ▶ B2B/B2C

■ BMC Software

- ▶ Systems management

■ Bynari

- ▶ Mail and calendar server

■ Computer Associates

- ▶ Systems management

■ Compuware

- ▶ Systems management

■ Dignus

- ▶ Development tools

■ Halcyon Software

- ▶ Windows migration

■ Logics Software

- ▶ Development tools

■ Macro4

- ▶ Print server

■ Rational Software

- ▶ Development tools

■ Rogue Wave Software

- ▶ Development tools



Software - ISV applications (cont'd)



- **RTS Realtime Systems**

- ▶ Stock tracking

- **SAP**

- ▶ ERP

- **Sendmail**

- ▶ Mail server

- **Software AG**

- ▶ XML Database

For a complete list, see "Developer Products for Linux for zSeries and S/390:

<http://www-1.ibm.com/servers/eserver/zseries/solutions/s390da/linuxproduct.html>



Software - IBM Middleware



- **IMS Connect**

- ▶ GA Oct 27, 2000
- ▶ Product information at <http://www-4.ibm.com/software/data/ims/imstoc.html>
- ▶ Priced facility which runs with IMS V5, V6 and V7. Product number 5655-E51 and this is licensed against the OS/390 system (MSU based pricing) on a one time charge.

- **CICS Transaction Gateway V3.1**

- ▶ GA Dec 22, 2000
- ▶ Product information at:
<http://www-4.ibm.com/software/ts/cics/platforms/desktop/>
- ▶ Product available at: <http://www6.software.ibm.com/dl/cctg312/cctg312-p>

- **MQSeries Client**

- ▶ GA Dec 15, 2000
- ▶ Product information and download at:
<http://www-4.ibm.com/software/ts/mqseries/txppacs/mach.html>





Software - IBM Middleware (cont'd)

■ **WebSphere Advanced Edition V3.5 with Java JDK V1.2.2**

- ▶ GA Dec 22, 2000
- ▶ Product information at http://www-4.ibm.com/software/webservers/appserv/linux_ae_v35.html
- ▶ Trial version (60 day) is available at http://www-4.ibm.com/software/webservers/appserv/download_linux_s390.html
- ▶ WAS V4.0 available late 2001?

■ **DB2 UDB Enterprise Edition V7 including DB2 Connect / JDBC**

- ▶ GA Dec 15, 2000
- ▶ Product information at <http://www-4.ibm.com/software/data/db2/linux/s390/>
- ▶ 60 day trial version available at <http://www6.software.ibm.com/dl/db2udbdl/db2udbdl-p>

■ **Tivoli Storage Manager Client**

- ▶ GA Dec 29, 2000
- ▶ Product information at http://www.tivoli.com/products/index/storage_mgr/
- ▶ Customer must have TSM Client licenses for each Linux for S/390 server.



Notes:





Agenda - section 6 of 9

- Introduction
- Brief history and level set
- zSeries platform options
- Distributions (SuSE, RedHat, Turbolinux, others)
 - ▶ Break
- Possible scenarios and solutions
 - ▶ Lunch
- ★ Linux for zSeries in the enterprise
- Miscellaneous topics
 - ▶ Break
- Sizing and TCO explained
- Summary, future, Q&A



Linux for zSeries in the enterprise - Outline

- Enterprise themes
 - ▶ System management
 - ▶ Backup and restore
 - ▶ LDAP
 - ▶ LVM
 - ▶ High availability/failover
 - ▶ Debugging
 - ▶ Performance
 - ▶ Security
- References
- Services





System management - outline

- Basic commands
- YaST (SuSE)
- linuxconf (typically Red Hat)
- Webmin
- Others

@server



System management - commands

■ ps, pstree

```
# pstree
init--atd
  |-cron---cron
  |-dsmc---dsmc---dsmc
  |-httpd---2*[httpd]
  |-inetd---in.telnetd---login---bash---su---bash---pstree
  |-kflushd
  |-klogd
  |-kmcheck
  |-kpiod
  |-kswapd
  |-kupdate
  |-lpd
  |-mingetty
  |-nmbd
  |-nscd---nscd---5*[nscd]
  |-qeth_softsetup
  |-smbd---7*[smbd]
  `--syslogd
```

@server

System management - commands (cont'd)



■ top

- ▶ Similar info to ps
- ▶ Continuously self-updating
- ▶ Sorted by %CPU time
- ▶ 'q' gets you out
- ▶ Example

```
2:39pm up 26 days, 5:17, 1 user, load average: 0.06, 0.01, 0.00
42 processes: 41 sleeping, 1 running, 0 zombie, 0 stopped
CPU states: 0.3% user, 0.1% system, 0.0% nice, 99.4% idle
Mem: 62544K av, 60152K used, 2392K free, 15524K shrd, 3424K buff
Swap: 143980K av, 9684K used, 134296K free 18984K cached
```

PID	USER	PRI	NI	SIZE	RSS	SHARE	STAT	LIB	%CPU	%MEM	TIME	CMD
25565	mikem	17	0	1432	1232	796	S	0	6.1	1.9	0:31	smbd
28714	root	9	0	1028	1028	864	R	0	0.7	1.6	0:00	top
1	root	0	0	80	60	44	S	0	0.0	0.0	0:00	init
...												



System management - commands (cont'd)



■ /proc file system

- ▶ Virtual file system
- ▶ A peek into the kernels data structures
- ▶ Interesting files and directories
 - /proc/cpuinfo
 - /proc/meminfo
 - /proc/dasd/devices
 - /proc/version
 - /proc/NNN (where NNN is the process number) - directory of a process

■ vmstat

- ▶ gives info about memory, swap space, I/O activity, CPU usage
- ▶ example:

```
# vmstat
procs          memory    swap        io          system          cpu
 r  b  w  swpd  free   buff  cache  si  so  bi  bo  in  cs  us  sy  id
0  0  0   9684  2500   3424  18984  0  0  1  0  0  14  0  0  5
```



System management - commands (cont'd)



■ df

- ▶ Display File systems command
- ▶ -h flag -> human readable
- ▶ Example:

```
# df -h
Filesystem                Size      Used Avail Use% Mounted on
/dev/dasdb1                1.4G    1.1G    158M   88% /
/dev/dasdc1                787M     8.1M    739M    1% /home
/dev/redlvm/redvoll1      20G      14G     5.3G   73% /redbooks
/suse/suse-cd1.iso        390M    390M        0 100% /redbooks/suse/cd1
```

■ du

- ▶ Disk Usage command
- ▶ Output in kilobytes
- ▶ -s flag -> summary
- ▶ Exampe:

```
# cd /redbooks
# du -s
15408184 .
```

@server

System management - commands (cont'd)



■ netstat

- ▶ Displays network statistic information
- ▶ Many options available
- ▶ Example:

```
# netstat -a
Active Internet connections (servers and established)
Proto RecvQ SendQ Local Address           Foreign Address         State
tcp        0      0 lnxpok2.its:netbios-ssn kayinlam.itso.ibm.:1063 ESTAB
tcp        0    138 lnxpok2.itso.ibm:telnet  zvmlnx1.itso.ibm:relief ESTAB
...
```

■ mount

- ▶ Display mounted file systems or mount a file system
- ▶ Example:

```
# mount
/dev/dasdb1 on / type ext2 (rw)
proc on /proc type proc (rw)
/dev/dasdc1 on /home type ext2 (rw)
/dev/redlvm/redvoll1 on /redbooks type ext2 (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=0620)
```

@server



System management - YaST

- **Yet another Setup Tool (YaST)**
- **Written by SuSE**
- **Can be used for:**
 - ▶ Installation,
 - ▶ User administration
 - ▶ System administration
 - ▶ Software package management
 - ▶ Editing the /etc/rc.config file



System management - linuxconf

- **Often used with Red Hat based systems**
- **Can be run many different ways:**
 - ▶ Command line
 - ▶ X Window based
 - ▶ From a browser (linuxconf-web)
 - ▶ As a GNOME application (gnome-linuxconf)
- **Another all-in-one system management tool - can manage:**
 - ▶ Users and groups
 - ▶ System services
 - ▶ Devices



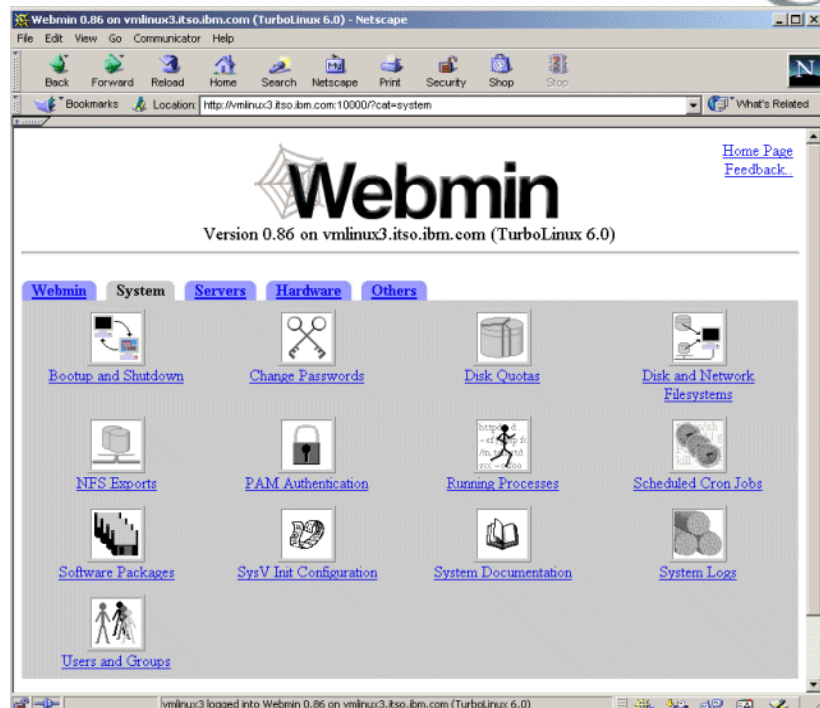
System management - Webmin



- Open source Web based system administration tool
- Includes
 - Large set of Perl scripts
 - Its own Web server - usually on port 10000
- Manages many aspects of a system:
 - Users, groups, passwords
 - Disk quotas
 - NFS exports
 - PAM
 - System logs
- See <http://www.webmin.com/webmin/>
 - Not all distributions ship this package



System management - Webmin screen shot



System management - Others



■ Other system management and system monitoring tools:

- ▶ Scotty/Tkined
- ▶ Big Brother
- ▶ Remstats
- ▶ BMC Patrol
- ▶ NetSaint
- ▶ Big Sister
- ▶ Mon
- ▶ MRTG
- ▶ OpenNMS

@server

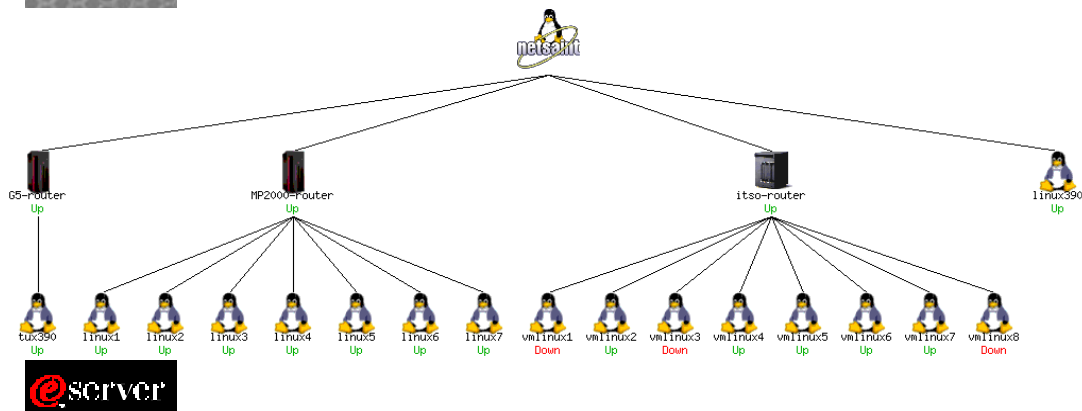
Netsaint screen shot



■ Sample config file /usr/local/netsaint/etc/hosts.cfg

```
host[G5-router]=G5 VM 3.1 Router;9.185.122.219;;check-router-alive;20;60;24x7;1;1;1;  
host[tux390]=tux390 VM Guest;9.185.122.217;G5-router;check-host-alive;10;120;24x7;1;1;1;
```

■ Sample display:




@server

MRTG screen shot



- So you can visualize your traffic

MRTG Index Page



@server

Backup and restore




- Split the task into:
 - ▶ Incremental backup and restore
 - ▶ Disaster recovery
- From Linux
 - ▶ IBM Tivoli Storage Manager (TSM) client - formerly ADSM
 - Is it V4.2 for Linux GA or still beta?
 - <https://www.tivoli.com/secure/support/pi/esp/tsmli390/index.html>
 - ▶ AMANDA - Advanced Maryland Automatic Network Disk Archiver
 - <http://www.amanda.org>
 - ▶ tar/cpio
 - ▶ CA Arcserve
- From z/VM
- From OS/390
- From hardware - Flash copy, PPRC

@server



Backup and restore - AMANDA

- ▶ Configuring a server:
 - Identify the clients and devices to be backed up
 - Establish access controls to read the clients' data
 - Configure disk staging areas on the server
 - Configure tapes and tape schedule
 - Set up and actioning the backup schedule
- ▶ Then use **amdump** in conjunction with **cron**
- ▶ Can also back up SMB shares via the **smbclient** command
- ▶ See chapter 12 of *Linux for zSeries and S/390: ISP/ASP Solutions*, SG24-6299



@server



Backup and restore - tar

- We assume a system with a tape drive and a working `/dev/ntibm0`
 - ▶ But you could also back up to a file and move the file over the network

- To back up:

```
# tar -cpzvf /dev/ntibm0 / --exclude /proc
```

- To restore:

- ▶ You must have a running system to restore over

```
# tar -xpvf /dev/ntibm0
```

@server



Backup and restore - From VM

■ CMS file backup

- ▶ File system must be a CMS reserved file

■ Dasd Dump and Restore (DDR)

- ▶ Example - Linux system on a 201 restored to a 202 minidisk

- ▶ To back up:

- Shutdown Linux and invoke DDR from VM:

```
SYSPRINT CONS // sends system messages to the console
IN 201 3390    // 201 is source DASD - entire root file system
OUT B38 TAPE   // B38 is the physical address of the tape drive
DUMP ALL       // source cylinders could be specified: COPY 0 3338
```

- ▶ To restore:

- Invoke DDR from VM:

```
SYSPRINT CONS // send system messages to the console)
IN B38 TAPE    // B38 is the source tape drive)
OUT 202 3390   // 202 is the minidisk to contain root file system
RESTORE ALL    // dest cyls could also be specified: RESTORE 0 3338
```

@server



Backup and restore - from OS/390

■ On-line DASD

- ▶ Requires:

- New fdasd command and at least one DASD partition
 - New Compatible Disk Layout (-cdl flag to the dasdfmt command)

- ▶ OS/390 can then "see" Linux DASD

- ▶ DFDSS can then be used as normal

■ Off-line DASD

- ▶ Greg Smith's Off-line dump and restore (offlindr) program


- ▶ <ftp://ftp.ox.ac.uk/pub/linux/s390/offlindr>

@server

LDAP

■ What is LDAP (Lightweight Directory Access Protocol)

- ▶ The x500 standard or DAP was too complex and hard to use
- ▶ Simply a hierarchy of information
- ▶ Uses a client/server model
- ▶ Servers are named slapd and slurpd
- ▶ Example directory tree:

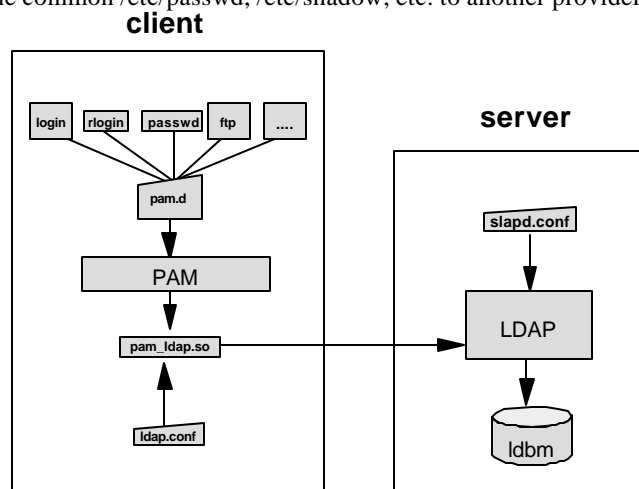


OpenLDAP

■ <http://www.openldap.org/>

■ Requires:


- ▶ Pluggable Authentication Module (PAM) - an authentication API
- ▶ Name Service Switch (NSS) - allows authentication to be redirected away from the common /etc/passwd, /etc/shadow, etc. to another provider





Logical Volume Manager - LVM

- From Sistina Software Inc.
- Kernel patches that allow combining of multiple physical disks into logical volumes
- Space can be dynamically added to a logical volume
- Kernel developers included patches into 2.4 kernel
- SuSE, Turbolinux 6.5, Caiman, ThinkBlue:
 - ▶ Install patches into 2.2 kernel
 - ▶ Include LVM utilities
- Attributes
 - ▶ Flexible disk storage system
 - ▶ Add disks to volume group
 - ▶ Delete disks from volume group
 - ▶ Rearrange contents in volume group

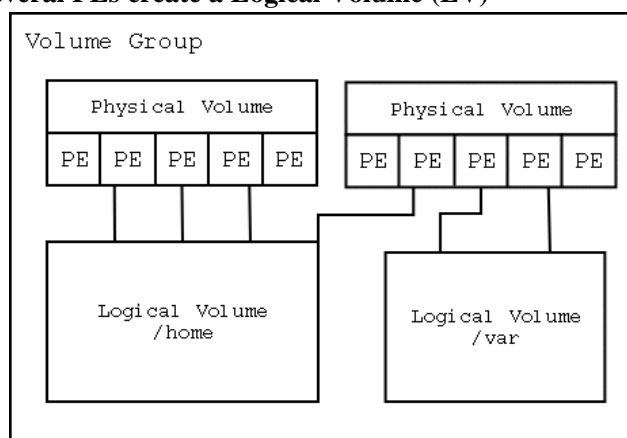


@server



Logical Volume Manager - LVM (cont'd)

- A DASD is a Physical Volume (PV)
- PVs are divided into Physical Extents (PEs) of the same size
- Several PVs create a Volume Group (VG)
- Several PEs create a Logical Volume (LV)



@server

LVM - Sample session



```
# for i in e f g h i j
> do; dasdfmt -b 4096 -f /dev/dasd${i}; done
# pvcreate /dev/dasd[efghij]1
# vgcreate -s 1m testvg /dev/dasd[efghij]1
# lvcreate -i 6 -I 8 -L 50 testvg -v
# mke2fs -b 4096 -m2 /dev/testvg/lvol1
# mount /dev/testvg/lvol1 /testvg
# cat /proc/lvm
Total: 2 VGs 8 PVs 2 LVs (2 LVs open 2 times)
Global: 16666 bytes vmallocated IOP version: 5 0:44:45 active
VG: rootvg [2 PV, 1 LV/1 open] PE Size: 4096 KB
  Usage [KB/PE]: 1433600 /350 total 1433600 /350 used 0 /0 free
  PVs: [AA] /dev/dasdc1 716800 /175 716800 /175 0 /0
        [AA] /dev/dasdd1 716800 /175 716800 /175 0 /0
  LV: [AWDS2 ] lvoll 1433600 /350 1x open
VG: testvg [6 PV, 1 LV/1 open] PE Size: 1024 KB
  Usage [KB/PE]: 98304 /96 total 55296 /54 used 43008 /42 free
  PVs: [AA] /dev/dasde1 16384 /16 9216 /9 7168 /7
        [AA] /dev/dasdf1 16384 /16 9216 /9 7168 /7
  ...
```

■ Source: <http://pax.gt.owl.de/~higson/LVM/>



High availability (HA)



■ Definition of High Availability (HA)

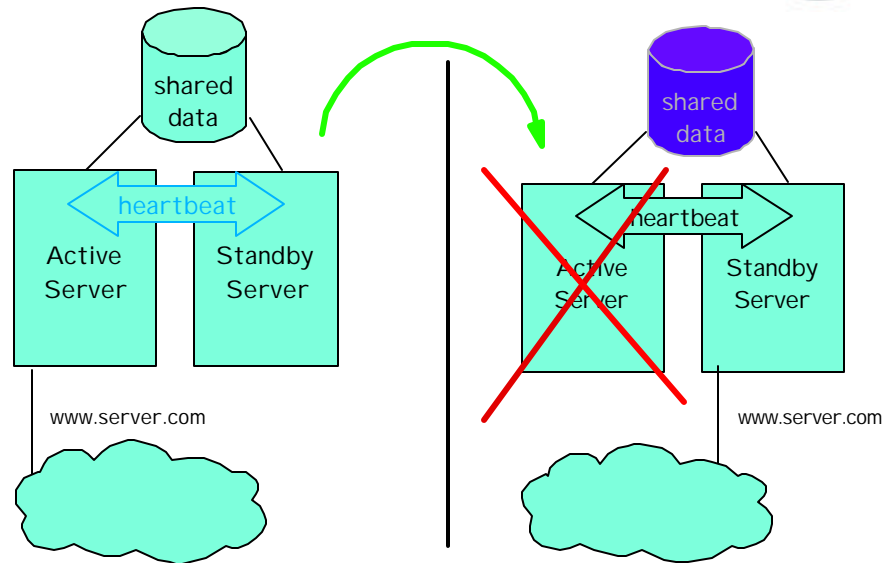
- ▶ Linux cluster
- ▶ Connectivity HA
- ▶ Application HA
- ▶ Data HA

■ Options

- ▶ Linux Virtual Server (LVS)
- ▶ DNS manipulation
- ▶ VIPA
- ▶ Dynamic routing



High availability - LVS



Debugging

■ Linux commands

- ▶ objdump
- ▶ strace
- ▶ ulimit
- ▶ gdp
- ▶ System.map
- ▶ ksymoops
- ▶ lcrash - new tool from LaPlace
- ▶ log files

■ VM diagnostics

- ▶ CP INDicate
- ▶ CP TRace
- ▶ CP Query
- ▶ many more



Performance tools

■ **Linux performance tools**

- ▶ vmstat
- ▶ top
- ▶ netsnmp

■ **z/VM performance tools**

- ▶ Real Time Monitor (RTM)
- ▶ VM Performance Reporting Facility (VMPRF)
- ▶ FCON/ESA

■ **Vendor tools**

- ▶ Velocity Software ESALPS

@server



Security

■ **Linux basics**

- ▶ Disable unneeded services
- ▶ Use secure Shell for remote access
- ▶ Use shadow password utilities
- ▶ Use the tcpd wrapper
- ▶ Use Pluggable Authentication Module (PAM)
- ▶ Use hardening tools
- ▶ Monitor security news and alerts

■ **Secure z/VM**

@server



Security (cont'd)

■ Disable unneeded services - you probably don't need all these:

- ▶ telnet
- ▶ ftp
- ▶ smtp
- ▶ time
- ▶ finger
- ▶ http
- ▶ pop-3
- ▶ login
- ▶ shell
- ▶ printer
- ▶ nfs

@server



Security (cont'd)

■ Use secure Shell for remote access

- ▶ SSH protects you from:
 - IP spoofing
 - DNS spoofing
 - Interception of cleartext passwords
- ▶ Ways of authentication for SSH
 - By password
 - By key authentication
 - Host authentication
 - Kerberos

■ Use shadow password utilities

- ▶ /etc/shadow does not have world-read permission
- ▶ Additional rules for strong passwords, expiration date, etc.

■ Use the tcp wrapper (usually in inetd)

@server



Security (cont'd)

■ **Use the Pluggable Authentication Module (PAM)**

- ▶ Keeps all the security code in one place
- ▶ Configured in /etc/pam.d directory

■ **Monitor security news and alerts**

■ **Use LDAP servers for authentication**

■ **Use firewall tools**

■ **Monitor log files**

■ **Use hardening tools**

- ▶ IPchains
- ▶ Tripwire

@server




Security

■ **z/VM security**

- ▶ CP access limit the privileged users
- ▶ Restrict access to resources (mdisk, spool, tape)
- ▶ DIRMAINT for directory administration
- ▶ Use RACF for most security
- ▶ TCP/IP z/VM stack disable unused services
- ▶ SSL, Kerberos
- ▶ IUCV specify the userid not '*'

@server

References - Banco Mercantil, Venezuela



Linux provides OS/390 users with the ability to extend and enhance their S/390 infrastructure by providing utility functions and rapidly deployable applications that run on S/390

S/390 provides Linux with


-) Fast connection to OS/390 managed data
-) Reliability and availability
-) An economical alternative to stand alone servers
-) Simplified systems management

Venezuelan Banking institution - 375 branches
 Linux for S/390 replaced 30 Microsoft NT servers
 File Servers
 DNS Servers
 Firewall
 Customer on-line account inquiry
 Uses SuSE distribution

"Linux on the IBM mainframe allows us to consolidate in a very cost effective way. In addition to these applications running on the mainframe, we have higher availability and reliability, and better performance. Consolidation also means that we have far fewer individual servers to administer." - Isaac Arismendi, IT Infrastructure Manager, Banco Mercantil.



References - Winnebago Industries



In distributed application environments Linux provides:

-) Highly stable Open Source operating system
-) Large selection of low cost, high quality applications
-) Large numbers of trained administrators & programmers

S/390 provides Linux with:

-) The ability to run many Linux instances
-) Secure isolation of user workload and data
-) Upgradability without repurchase
-) Reliability and availability

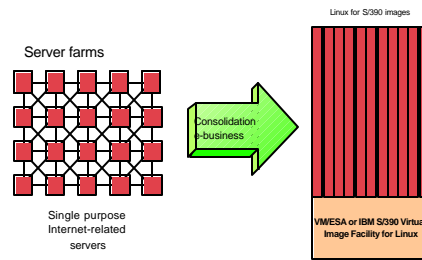
Recreational Vehicle Manufacturer
 Replaced Microsoft Mail and Novell servers
 with Bynari Insight Server
 Users can still use Microsoft Outlook e-mail client
 Runs on Multiprise 3000

"Upgrading to Microsoft Exchange or Novell GroupWise would cost Winnebago about \$100,000. We expect Bynari Insight Server to cost a third or even a quarter of that, because Winnebago will be able to make use of excess capacity on the company's IBM Multiprise 3000 mainframe. The mainframe is our industrial-strength system that stays up, and that's where our expertise is."

- Dave Ennen - Technical Support Manager - Winnebago Industries



References - Korean Air



In distributed application environments Linux provides:

-) Highly stable Open Source operating system
-) Large selection of low cost, high quality apps
-) Large numbers of trained administrators & programmers


S/390 provides Linux with

-) The ability to run many Linux instances
-) Secure isolation of user workload and data
-) Upgradability without repurchase
-) Reliability and availability

"A single IBM ^ z900 running Linux can do the work of an entire server farm. Multiple copies of Linux can run side by side on a server allowing for highly scalable and manageable environments that can handle unpredictable spikes in internet activity"
Yong-Seung Hwang - Chief Information Officer - Korean Air

@server

References - Grede Foundries



In VM and VSE environments Linux provides:

-) A wide range of applications not available on VM/ESA or VSE/ESA
-) The ability to leverage existing applications and data in e-business scenarios

S/390 provides Linux with:

-) Fast connection to VSE/ESA managed data
-) Reliability and availability
-) An economical alternative to stand alone servers

Parts supplier to the automotive industry

VM/ESA - VSE/ESA user

Initial use of Linux on S/390 to complement:

-) VM and VSE environment
-) File/Print serving
-) Domain name serving


Participated in UDB and DB2 Connect beta program

I got the request to bring up the DNS and had it running quickly enough to solve our business need almost instantly".

-Rich Smrcina, Systems Software Specialist Grede Foundries

@server

References - Telia AB



In ASP/ISP environments Linux provides:

-) The operating system of choice for many new dot com's and Internet startups
-) Highly stable Open Source operating system
-) Large selection of applications
-) Large numbers of trained administrators & programmers

S/390 provides Linux with:

-) The ability to run many Linux instances
-) Secure isolation of user workload and data
-) Upgradability without repurchase
-) Reliability and availability

"This new S/390 ZZ7 running Linux allows us to rethink our total pricing structure for Internet services and to offer customers a more affordable web application service than ever before"

Henrik Wulff Riedl, CFO, TeliaNet."



References - Telia, the true story?



From the linux-390 list server:

(Almost) our entire Usenet service is running on the G6 we bought. It runs on Diablo (<http://www.openusenet.org/diablo/>), **this adds up to apx. 15 servers**. Only our spool servers are still running on other equipment, but they will be moved as well.

We run some of our DNS-servers using Bind (all the primary ones)

We are moving all our web and mail services within a couple of weeks (running on Apache, Postfix (Procmail), Pure-ftpd and some others).

Kind regards,

--

Allan Jørgensen (AJ1382-RIPE), Newsmaster/Senior Unix technician
Telia Net, Internet Production, Data Operation Center East
Ejby Industrivej 135, 2600 Glostrup, Denmark



References - Telia, the true story? (cont'd)



From the linux-390 list server:

Date: Tue, 18 Sep 2001 09:32:31 -0400
From: David Boyes <dboyes@sinenomine.net>
Subject: Telia reports increased profitability....

From the International Herald-Tribune:

"TELIA SEES PROFIT GROWTH: The Swedish telecom operator Telia AB said that its underlying profit margin would grow 25 percent to 30 percent by late 2002 or early 2003 from 23 percent in the first half of 2001.

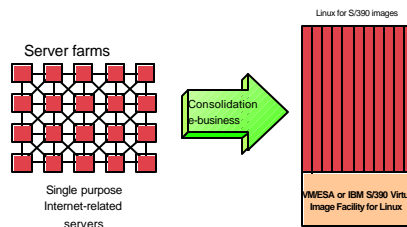
The rise in profitability will come thanks to recent restructuring measures and from the return to profitability of its loss-making international carrier and Internet service units, its chief executive, Marianne Nivert, said. Annual sales growth of 8 percent to 10 percent was expected in the coming years.

Compiled by Victoria Shannon from AP, AFP, AFX, Bloomberg and Reuters reports

<http://tm0.com/IHT/sbct.cgi?s=155663576&i=391626&d=1785256>"

@server

References - DreamBall, Korea



In ASP/ISP environments Linux provides:

-) The operating system of choice for many new dot com's and Internet startups
-) Highly stable Open Source operating system
-) Large selection of applications
-) Many trained administrators & programmers

S/390 provides Linux with:

-) The ability to run many Linux instances
-) Secure isolation of user workload and data
-) Upgradability without repurchase
-) Reliability and availability

Dreamball - Korean on-line games provider

On-line soccer game which coincides with World Cup

Supports large numbers of Internet users

Each user requires multiple servers

Dreamball chose S/390 for:

Economy vs. stand alone servers

Reliability, availability, scalability, performance

Time to market

"What is crucial to on-line games is whether they can stably accommodate a large number of users at the same time. We found the IBM ^ zSeries best suited for such crucial missions. We chose IBM as we required high performance and availability, reliability, scalability and serviceability. By deploying a single IBM mainframe rather than a bunch of other servers we have benefited in terms of lower costs of office space and system management resources."

Won-Young Bu - CEO DreamBall

@server



References - Credit Index

■ Credit Index

- ▶ A large risk modeling company for the direct response (mail order) industry
- ▶ Traditionally has run business entirely on the mainframe

■ Needed large file support

- ▶ Routinely deal with files many GB in size
- ▶ Co-reqs the linux 2.4 kernel

■ Needed tape support

- ▶ These large files are often stored on 3490 tapes

■ Turbolinux v6.5 enabled all of this

- ▶ With custom services

■ See <http://eltoday.com/article.php3?itsn=2001-08-15-001-14-PS>

- "Running Linux on Intel-based servers as well as the mainframe has led to greater than anticipated costs savings for the Credit Index because it has become much easier for programmers to develop and test applications for the mainframe. Programmers can write and test applications on Linux servers and then deploy them on the mainframe."

@server



References - ITSO Poughkeepsie

■ IBM organization producing redbooks and workshops

- ▶ Need shared file system accessible to Windows clients
- ▶ Current solution is OS/2 - adding servers and hard drives is a hassle
- ▶ Samba allows all maintenance to be done on the server side
- ▶ Linux on S/390 allows:
 - S/390 reliability, backup and restore experience
 - Flexibility of VM minidisks
 - Linux telnet, FTP, etc. access to clients
- ▶ LVM with three 3390-9s give us a 20GB file system
- ▶ Administration (Samba, users, groups, adding volumes) is not easy for personnel not trained in UNIX
- ▶ Performance has not been an issue

"Serving redbooks off of a Linux for S/390 server has allowed me to provide a large amount of shared file space, easily add new users, and feel confident the share is being backed up consistently"

Michael MacIsaac - IBM

@server



Services - Overview

- **SuSE**
- **Turbolinux**
- **IBM**
- **Others**

@server



Services - SuSE

- <http://www.suse.com/en/produkte/susesoft/s390/services.html>
- **Attributes**
 - ▶ Installation planning, services and verification
 - ▶ Individual consulting services
 - ▶ Education - S/390 specific, general Linux
 - ▶ Pilot project
 - ▶ Setting up a Web Server
 - ▶ Setting up mail services
 - ▶ Firewall planning and installation
 - ▶ Performance measurement and tuning

@server



Services - Turbolinux

- **z/Linux(sm) Service Delivery Option (SDO)**
- See <http://www.turbolinux.com/products/s390/sdo.html>
- **Attributes:**
 - ▶ 24 x 7 Help Desk support for defect support. Defects supported include functional problems with any RPM-packaged component, performance problems, and product currency issues. For additional support plans please see our TurboCare pages.
 - ▶ Fixes to known problems supplied as requested.
 - ▶ System refresh tapes available every six months, shipped upon request.
 - ▶ Unlimited Linux images under VM.
 - ▶ Pre-delivery consultation(s) with the customer.
 - ▶ One free pass to our 5 day Linux course taught at Turbolinux headquarters in Brisbane, California.



Services - IBM Global Services

Consulting

- ▶ High performance clusters consulting

Implementation

- ▶ High Performance Cluster Integration & Installation
- ▶ Linux Solutions for e-business on e-server
- ▶ DB2 Migration on zSeries
- ▶ MQSeries SmoothStartJ
- ▶ xSeries Hosting
- ▶ Server Consolidation (Q301)
- ▶ Web Clusters (Q401)
- ▶ Platform Migrations from OS/2, & Solaris (Q401)

Technical Support

- ▶ Support Line
 - One-stop (24x7) enterprise level support
 - Support for major Linux distributions
 - Toll-free Phone & electronic access
- ▶ Account Advocate
- ▶ Advanced Support

Education

- ▶ Class room courses
- ▶ Web-based courses
- ▶ Redbooks (free)

- ▶ General: http://www-1.ibm.com/services/e-business/linux_2.html
- ▶ New service: http://www.ibmmlink.ibm.com/usalets&parms=H_601-031





Services - Others

■ Sine Nomine

- ▶ <http://www.sinenomine.net/>
- ▶ "We are shaping the cutting edge of the Linux for System/390 revolution. Two of the founders of Sine Nomine are deeply involved in the development and testing of the Linux port for the System/390. We have found Linux for System/390 to be a reliable and stable solution for mission-critical Internet support services.

■ Equant

- ▶ <http://www.equant.com>
- ▶ Interview with Jimmy Lee of Equant: <http://consultingtimes.com/equant.html>

■ VM Resources Ltd.

- ▶ <http://209.167.83.8/vmrnew/services.htm>
- ▶ VM Linux Manager (VM Linman) in conjunction with Glass House Systems
 - <http://www.glasshousesystems.com/home.html>



Notes:





Agenda - section 7 of 9

- **Introduction**
- **Brief history and level set**
- **zSeries platform options**
- **Distributions (SuSE, RedHat, Turbolinux, others)**
 - Break
- **Possible scenarios and solutions**
 - Lunch
- **Linux for zSeries and the enterprise**
- ★ **Miscellaneous topics**
 - Break
- **Sizing and TCO explained**
- **Summary, future, Q&A**



Miscellaneous - outline

- **Run levels**
- **File systems**
 - types
- **Networking**
 - Starting
 - inetd
- **Desktop**
- **Tape driver**





Run levels

■ A Run level is a state of the Linux operating system.

- ▶ Some are transitory - halt and reboot
- ▶ Some are common - Multi-user mode, with X11
- ▶ Can be queried with the command:

```
# runlevel
N 2
```

- ▶ Can be set with the command:

```
# telinit 3
```

- ▶ To start a run level, the "S*" commands in the directory /etc/rc.d/rcN.d are run

```
# cd /etc/rc.d/rc3.d
# ls S*
S01kernelcd      S08portmap      S17nfsserver    S20rstatd      S21nsd
S04dummy         S09nfs          S17pcnfsd      S20rusersd     S21smbfs
S04firewall_init S09syslog       S19userproc    S20rwhod       S30xdm
S05dhclient      S10routed       S20apache      S20sendmail    S99final
S05network       S13random       S20at          S20smb
...
```

- ▶ To leave a run level, the "K*" commands are run



Run levels (cont'd)

Run level	SuSE7.0	Red Hat, many others
0 (reserved)	Halt	Halt
S	Single user mode	N/A
1 (reserved)	Multi-user without network	Single user mode
2	Multi-user with network	Multi-user, without NFS
3	Multi-user with network and xdm	Full multi-user mode
4,7,8,9	Unused	Unused
5	Unused	X11
6 (reserved)	Reboot	Reboot





File systems - Types

- **Non-journalled - ext2**
- **Journalled**
 - ▶ ext3
 - ▶ ReiserFS
 - ▶ IBM JFS
 - ▶ SGI XFS
- **Global or networked**
 - ▶ NFS
 - ▶ Sistina GFS
 - ▶ AFS
 - ▶ IBM GPFS
 - ▶ CIFS (Samba)
- **Others**
 - ▶ Swap
 - ▶ Virtual (/proc, /devfs in 2.4 kernel)



File systems - Why journalled?

- **Non journalled file systems**
 - ▶ Require multiple media I/Os
 - ▶ Can be in an inconsistent state after a crash
 - ▶ Result in a file system check (fsck) if suspected to be inconsisent
 - ▶ Recovery can take hours
- **Journalled file systems**
 - ▶ Use atomic transactions to keep track of meta-data changes
 - ▶ After a crash, runs a "replay log" to repair a file system
 - ▶ Recovery is usually in seconds





Journalled file systems - ext3

- Caiman has added
- LaPlace is using
- Millenux has added
- Can switch between non-journal and journal mode (ext2 and ext3)
- May ride the popularity of ext2 into the future




Journalled file systems - IBM JFS

- Download from <http://www-124.ibm.com/jfs/>
- IBM has contributed it to the open source community
- Attributes
 - 64-bit
 - Journalling of meta-data only in B+ trees
 - Restart after crash in seconds
- Recently, serious problems with SMP machines were encountered
 - This bug is purportedly fixed



File Systems - Sistina's GFS



@server

Sistina's GFS (cont'd)



- From Linux Weekly News, September 6th, 2001, see <http://www.lwn.net/2001/0906/>:

- ▶ "GFS is no longer free software. The Global Filesystem (GFS) is a clustered filesystem developed by Sistina. It is meant for the implementation of high-performance, high-availability filesystems on "storage area networks." It has long been available under the GPL, and was considered as a candidate for inclusion into the Linux kernel if and when the 2.5 series comes into existence.
- ▶ That was until version 4.2 came out under the new "Sistina Public License." This license looks somewhat like a free software license, in that source is available. The similarity ends there, however. Redistribution requires that a license fee be paid to Sistina; one must also pay if GFS is used to offer a commercial service, even if the software is not redistributed. The SPL is certainly not a free software license. It has more of a "shared source" smell to it.

- From Alan Cox, see <http://lwn.net/2001/0906/a/ac-gfs.php3>:

- ▶ "> I just saw that GFS is now being released under a new license: SPL
- ▶ I've already sent them formal notice that their current 4.2 tree mixed with Linux appears to be violating my GPL copyright. I've been ignored so far. They will be getting a registered letter in a few days if they don't bother to respond."

@server

Global file systems - AFS



■ Andrew File System (AFS)

- ▶ Solid code - proven as a commercial product
- ▶ OpenAFS is a contribution to the open source community by IBM

■ OpenAFS.org

- ▶ Site for coordination and distribution of ongoing OpenAFS development.
- ▶ <http://www.openafs.org/>

■ Running on Linux for zSeries and S/390

- ▶ OpenAFS 1.1.1 built cleanly for us
- ▶ See section 15.4 of *Linux for S/390 and zSeries: ISP/ASP Solutions*, SG24-6299, redpiece on the Web at:
<http://www.redbooks.ibm.com/abstracts/sg246299.html>
- ▶ OpenAFS 1.1.2 is now out: <http://www.openafs.org/release/openafs-1.2.0.html>



Global file systems - CIFS/Samba



■ Samba executables: (usually in /usr/sbin)

- ▶ `smbd` Server Message Block Daemon - main workhorse
- ▶ `nmbd` NetBIOS Daemon - browse list daemon
- ▶ `smbclient` Connects to SMB shares, ftp-like syntax
- ▶ `testparm` Tests validity of `/etc/smb.conf` file
- ▶ `swat` Samba Web Administration Tool

■ Samba configuration file: (usually `/etc/smb.conf`)

- ▶ Special sections are `[global]`, `[homes]` and `[printers]`
- ▶ Some important values are:
 - The "homes" section - makes each user's home directory a "share"

```
[homes]
  browseable = no
  read only = no
  create mode = 0750
```
 - Authentication parameters:

```
security = user                # UNIX authentication - recommended
# security = server            # to use NT-Server-Name authentication
# password server = <NT-Server-Name> # goes with previous line
encrypt passwords = yes        # no registry changes needed
smb passwd file = /etc/smbpasswd # goes with previous line
```



Global file systems - CIFS/Samba (cont'd)



■ Samba configuration file: (cont'd)

► Other global parameters:

```
log file = /var/log/samba/log.%m    # recommended
max log size = 100                  # max file size in KB
socket options = TCP_NODELAY        # usually better performance
```

► Creating a share: just add a new [section] with a "path = /dir" parameter - example:

```
[usrsrc]                             # share name
path = /usr/src                      # file system path of share
writable = yes
force group = +sg246299
create mask = 0664
directory mask = 0774
oplock = no                          # possible performance bug
```

► Test changes to smb.conf with **testparm** command

@server

Networking - overview



■ Important files

► /etc/modules.conf - lists modules and "devices" they are associated with:

```
alias eth0 lcs
alias eth1      off
alias tr0       off
```

— lcs can mean the module lcs.o is being used

```
# locate lcs.o
/lib/modules/2.2.16/net/lcs.o
# lsmod
Module          Size  Used by
lcs              14888   1 (autoclean)
```

► /etc/rc.config (SuSE only) - a list of all your system's parameters

```
NETCONFIG="_0"
IPADDR_0="9.12.17.172"
NETDEV_0="eth0"
IFCONFIG_0="9.12.17.172 broadcast 9.12.17.255
netmask 255.255.255.0 mtu 1492 up"
```

► /etc/hosts - IP addresses are associated with "host names"

```
# cat /etc/hosts
...
9.12.6.53      lnxpok2.itso.ibm.com    lnxpok2
```

► /etc/route.conf - static routes

```
# cat /etc/route.conf
...
9.12.6.0       0.0.0.0      255.255.255.0    eth0
default        9.12.6.75
```

@server



Networking - overview (cont'd)

■ Important files


- ▶ /etc/resolv.conf - DNS server search order - we are not using DNS, but an example would be:
 - search itso.ibm.com
 - nameserver 9.12.14.7
- ▶ /etc/rc.d/dhclient - script that sets you up as a DHCP client
- ▶ /etc/rc.d/route - script that sets up static routes
- ▶ /etc/rc.d/inet - start the inet super daemon
- ▶ /etc/rc.d/network - script that brings up networking

@server



Networking - inetd

- ▶ Many daemons listen on a "well-known port", and fork() themselves when a request comes in.
- ▶ The inetd daemon is a "generic listener" that listens on many well-known ports and forks the appropriate daemon
- ▶ List of "services" and port numbers is in /etc/services
- ▶ List of daemons and "services" is in /etc/inetd.conf



@server



Networking - inetd

- ▶ Example of looking for SWAT service

```
# grep ^swat services inetd.conf
services:swat    901/tcp      # XXX Samba Web Admin Tool
inetd.conf:swat  stream tcp    nowait.400 root /usr/sbin/swat  swat
```

- ▶ tcpd is another "wrapper" - tcpd: this provides a generic method of authenticating incoming hosts

- ▶ After the inetd.conf file is modified, inetd can be forced to re-read it:

```
# ps -ef | grep inetd
root      467      1   0 10:40 ?        00:00:00 /usr/sbin/inetd
root     13550  2393   0 15:17 pts/2    00:00:00 grep  inetd
# kill -HUP 467
```

- **xinetd is a new generic listener with one configuration file per service**


- ▶ Better security, more options
- ▶ Red Hat 7.1 uses this



Networking - virtual IP addressing

- **VIPA (Virtual IP Addressing)**

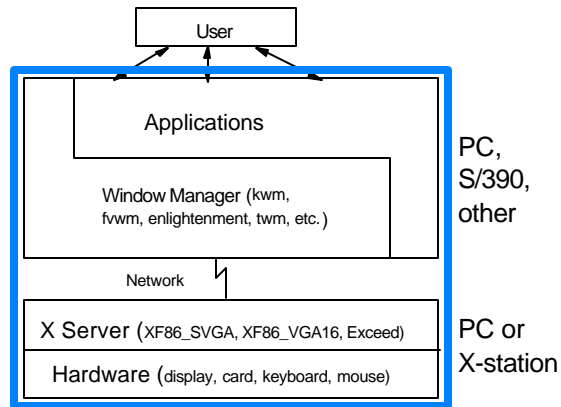
- ▶ Isolate applications from the failure of an IP interface
- ▶ Provided through z/VM





Desktop

- OK to run small X Window apps on S/390 displayed to an X server
- It is a questionable use of cycles to run a desktop (e.g. KDE) on S/390
- You cannot run an X Server on S/390



@server



Tape driver

- Distributions that include it:
 - ▶ Turbolinux v6.5.1
 - ▶ Caiman
 - ▶ Millenux
- Still basic by S/390 standards

@server



Notes:

@server



Agenda - section 8 of 9

- Introduction
- Brief history and level set
- zSeries platform options
- Distributions (SuSE, RedHat, Turbolinux, others)
 - ▶ Break
- Possible scenarios and solutions
 - ▶ Lunch
- Linux for zSeries in the enterprise
- Miscellaneous topics
 - ▶ Break
- ★ Sizing and TCO explained
- Summary, future, Q&A

@server



Sizing and TCO explained - outline

- **Processor Design**
- **z/OS Base**
- **Linux Impact**
- **Capacity Management**
- **Networking**
- **Application Scalability**
- **Summary/Conclusions**



Benchmark definition


A reference point.

Originally: a mark on a workbench used to compare the lengths of pieces so as to determine whether one was longer or shorter than desired.

For computers: a "benchmark" is a test, or set of tests, designed to compare the performance of one computer system against the performance of others.



Some performance design points




@server

Processor Design



Four Components of Capacity

- ▶ Processor
- ▶ Memory Hierarchy
- ▶ I/O Structure
- ▶ Management

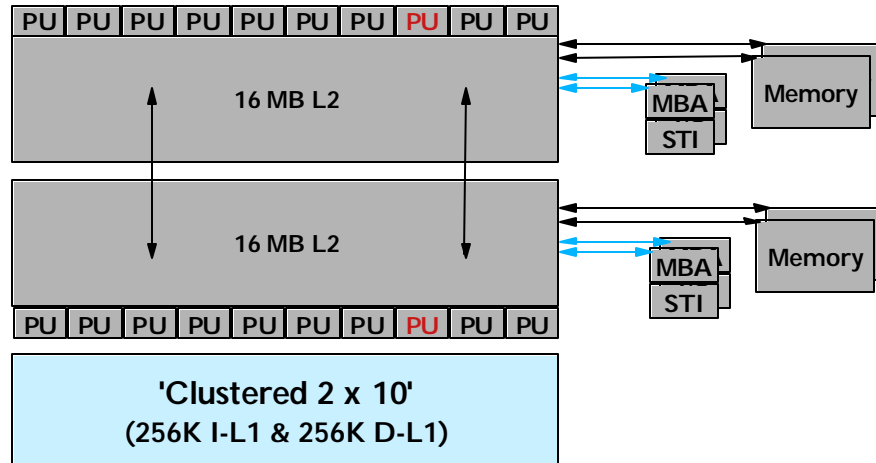


Processor, Memory and I/O time vary greatly with application

Processor designs vary greatly in the balance of capacity across components which results in a wide range of relative capacity.

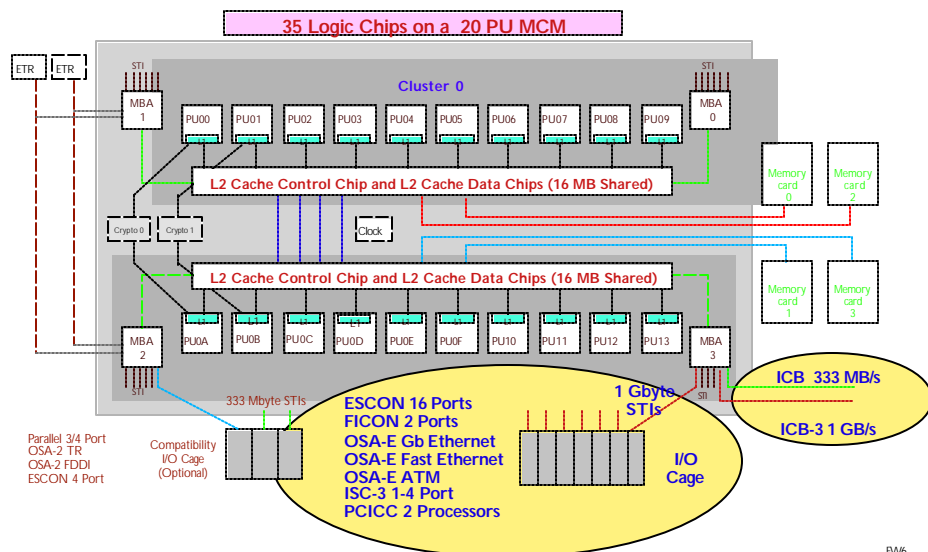
@server

z900 System-Structure ('Logical View')



@server

z900 20 PU System




PW6

@server




The diagram illustrates the relative execution times of CPU, Memory, and I/O operations for two different architectures. The top row, labeled 'S/390', shows a large orange bar for 'CPU Time', a medium green bar for 'Memory Time', and a small red bar for 'I/O Time'. The bottom row, labeled 'Others', shows a small orange bar for 'CPU Time', a medium green bar for 'Memory Time', and a large red bar for 'I/O Time'. Arrows indicate the sequence of operations: CPU Time, then Memory Time, then I/O Time. The 'CPU Busy' label is placed above the CPU Time bars. The 'I/O Busy' label is placed above the I/O Time bars. The 'Memory Busy' label is placed above the Memory Time bars.

Data Intense work such as BI, Very Large Database, Classical OLTP or "cache killer" workloads (Object oriented code or context switching) will potentially run much better on S/390




* I/O scaling also varies: S/390 is better at 4k blocks (major ISVs tune to this) while the others are better at single byte movement (mostly RYO code).


$$\text{Power} = \text{Geo Mean}(\text{Utilization, Clock, SPEC, tpm})$$

$$\text{effBW} = \text{Geo Mean}(\text{Cache, BW, Latency, Memory, IO BW})$$


Processor Performance impacts capacity planning but use of CPU execution alone is misleading

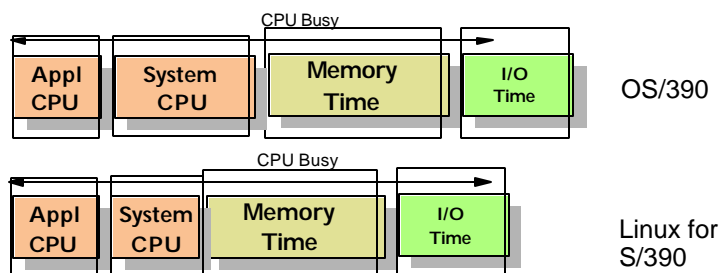


- ★ S/390 is not for "Deep Computing"
- ★ Most commercial applications range between CPU and Data intensity
- ★ S/390 is very good at "data intense" work for many users.

@server

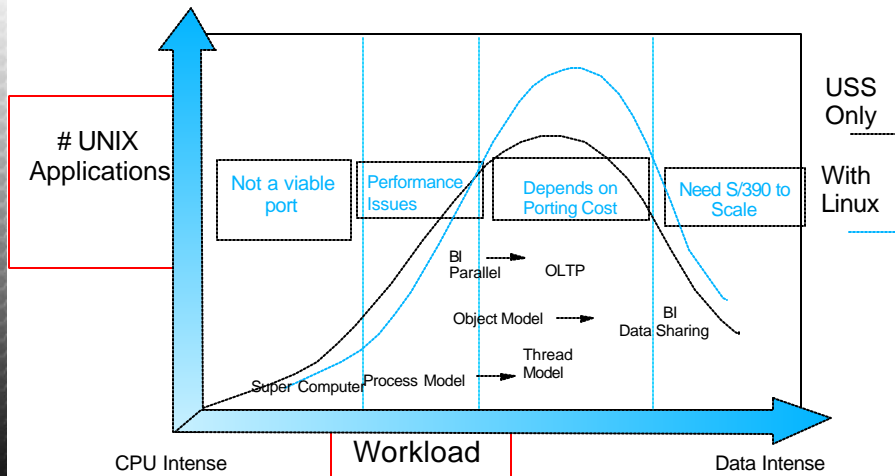
Impact for LINUX on S/390 Performance

Shorter Linux system path lengths will yield improvements for *some* applications



@server


Impact of Linux on S/390 Positioning



AND....Linux capability reduces porting costs, removes scaling issues in some applications, and reduces cost of acquisition. Therefore, applications move to S/390 more easily.



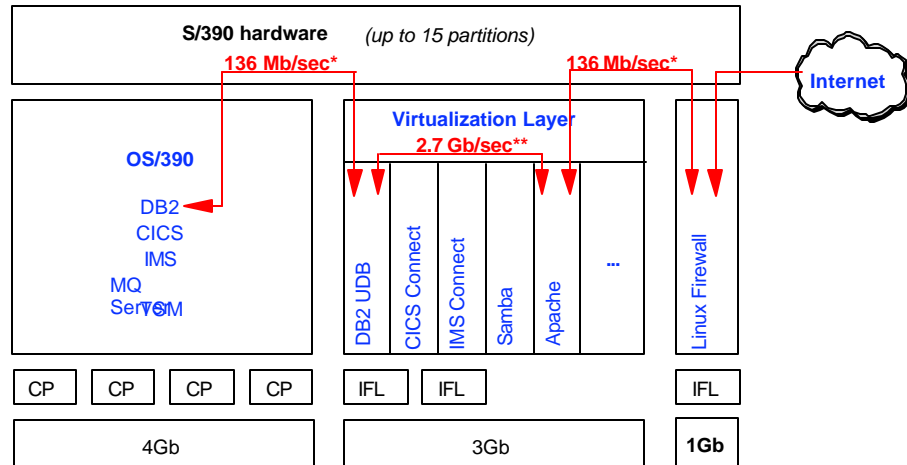
Linux Capacity Management at the Box Level



Virtual Server Technology increases useful capacity for customers. The combination of LPAR/VM/Virtual Image Facility and Linux is extremely powerful to manage "peaks" across Linux images within a single S/390.



Linux for S/390 for an existing S/390 customer



* ESCON CTC


** G6 Turbo



Linux under VM/Virtual Image Facility

Non-idle Overhead

- Proportional to VM function requests
- Depends on configuration
- Typical of other guests and CMS (somewhere between them)
 - ▶ Diagnose 250 for minidisk I/O and IUCV make it like CMS
 - ▶ Timers, real I/O, SMP, etc. make it like VSE, Music, TPF, or OS/390
- No surprises - overhead acceptable





S/390 Capacity Messages

- **The range of relative capacity for a particular application is workload and hardware architecture dependent**
- **The relative capacity range is approximately the same for Linux for zSeries and z/OS**
- **Server Consolidation or mixed workload configurations can exploit S/390 functions for more efficient use of hardware resources**
 - ▶ OS/390 and z/OS: WLM management
 - ▶ Linux for S/390: S/390 unique functions (LPAR, VM and Virtual Image Facility, z/OS Intelligent Resource Director)
- **Certain applications will not scale without re-engineering, regardless of platform**
- **Certain UNIX applications can perform well under Linux for little or none of the re-engineering required to exploit OS/390 efficiently**



Notes:





Agenda - section 9 of 9

- **Introduction**
- **Brief history and level set**
- **zSeries platform options**
- **Distributions (SuSE, RedHat, TurboLinux, others)**
 - ▶ Break
- **Possible scenarios and solutions**
 - ▶ Lunch
- **Linux for zSeries in the enterprise**
- **Miscellaneous topics**
 - ▶ Break
- **Sizing and TCO explained**
- ★ **Summary, future, Q&A**



Summary

- **IBM is seriously committed to Linux**
- **Linux runs natively on zSeries and S/390**
 - ▶ On "bare iron"
 - ▶ In an LPAR
 - ▶ Under VM
 - ▶ VIF is deprecated
- **The value of Linux on zSeries and S/390 is:**
 - ▶ Linux portfolio will increase the applications available to the S/390 customer
 - ▶ Linux-knowledgeable employees will continue to grow in Universities
 - ▶ The open source development model is changing the way software is developed
 - ▶ The flexibility and openness of Linux combined with qualities of service of S/390 result in an industrial strength Linux environment
 - ▶ Total cost of ownership (TCO) is initially higher on zSeries, but can become very low






Summary (cont'd)

- Look to replace servers with low utilization
- Let the distributors create the kernel and system for you
 - Unless you have a good reason to Roll Your Own
- 2.4 kernel is coming but not quite here yet
- Journalled file systems are coming but not quite here yet



Future directions - Hipersockets



- Connects Linux-Linux, Linux-z/OS, z/OS-z/OS
- TCP/IP based, high bandwidth, low latency
- Link layer reliable and secure transmission implemented in hardware
- Standard socket interface
- Operating system configuration for standard network interface
- Can be used by all TCP/IP based software, e.g. NFS, DB2Connect, MQSeries, Connectors





Linux for zSeries - future directions

■ **Code Availability in 2001 / 2002**

- ▶ Refreshes of Current Distributions
 - no enhancements, fixes only
- ▶ New Distributions (3Q or 4Q/01)
 - 2.4.x kernel - 31 bit
 - 2.4.x kernel - 64 bit
 - gcc 3.0
 - glibc 2.2
- ▶ IBM Middleware
 - Updates, Matching Distribution Levels
- ▶ Software (Updates)
 - e.g. SAP R/3 Application Server for 64 bit Linux

■ **Linux Enhancements in 2001**

- ▶ HiperSockets (announced Oct 2001)
- ▶ IPv6
- ▶ High Speed Token Ring
- ▶ Cryptographic coprocessor device driver (announced Oct. 2001)



Linux for zSeries - future directions (cont'd)

■ **z/VM enhancements**

- ▶ z/VM 4.1 - July 2001
 - Will run on IFLs
- ▶ z/VM 4.2
 - System Administration Facility (SAF) - similar function to VIF
 - A new virtual machine named VMADMIN
 - A VIF migration command named VIFMGR
 - VMADMIN commands can be run from **either** Linux or CMS
 - VM LAN
 - New CP function - DEFINE LAN and DEFINE NIC commands
 - Not a point to point connection which caused issues with vCTC and IUCV
 - Requires hipersockets





Web Sites

- ▶ This presentation in PDF form (IBMers only - others upon request):
 - ─ <http://wtscpok.itso.ibm.com/~zseries/linux/index.html>
- ▶ Mark Post's Linux for S/390 site:
 - ─ <http://linuxvm.org/>
- ▶ Jim Elliott's Linux for zSeries Home Page:
 - ─ <http://www.vm.ibm.com/devpages/jelliott/linux.html>
- ▶ Linux for zSeries Home Page:
 - ─ <http://www-1.ibm.com/servers/eserver/zseries/os/linux/>
- ▶ Velocity Software's Linux under VM performance site:
 - ─ <http://www.linuxvm.com/>

@server



A perspective from Gartner

When and when not to consider Linux on the mainframe:

"A growing number of Gartner's clients are asking about the feasibility of running Linux on S/390 and zSeries mainframes, rather than on the more usual Intel platforms... We believe there are only two major environments where Linux on mainframes should be considered: where there is a need to run hundreds to thousands of low-utilization Linux servers, and in enterprises currently using G5, G6 or zSeries machines and needing to run more than 30 to 50 utility servers providing facilities for applications such as Web serving, e-mail serving, or print and file services... Many specialized edge servers run at low utilization. This is the area where Linux for S/390 or zSeries can provide an advantage, because, by using z/VM, all the Linux images can share CPU, central storage and I/O resources without wasting them... The major TCO advantage that IBM is predicting is in staff costs, which can be significant when compared with other costs. One IBM reference account believes that it will be able to increase the number of virtual servers from fewer than 100 to 1,500 without increasing its Linux administration staff (there are currently four)... the burden of proof is on IBM."

Source: Gartner, June 19, 2001, John Phelps

@server



A perspective from Giga

Entry-level S/390 systems offer improved pricing for UNIX/NT server farm consolidations:

"The drive to make entry-level, System/390 pricing more appealing to IT professionals continues in high gear with recent IBM offerings. As part of this year's major commitment to advance Linux within the enterprise, IBM is offering significant price reductions for new shipments of as well as upgrades to the high-end model of its entry-level Multiprise 3000 family.... With these two offerings, IBM is providing IT professionals with substantially improved pricing in the under 700-MIPS computing range -- precisely the target for enterprises considering mid to large-scale Unix and NT server farm consolidations. The remaining ingredient for success here will be applications from the Linux community. Although the full suite of enterprise Linux applications has yet to fully materialize, those applications handling Internet and network infrastructure functions (i.e., Web page serving, DNS, firewall, file/print serving) and most recently mail, messaging and workgroup collaboration software products are now beginning to populate the enterprise Linux landscape."

Source: Giga Information Group, Jun. 20, 2001, David Mastrobattista



Questions and answers ???





vi - cheat sheet

■ vi works in three modes

- ▶ (1) input mode ('i' and other commands put you in this mode)
- ▶ (2) Command mode - 'Esc' brings you there
 - i brings you back to input mode
 - dd deletes a line and puts it in the buffer
 - <n>dd delete <n> lines
 - x delete a character
 - dw delete a word
 - p add the buffer past the current location
 - P add the buffer before the current location
 - o add a line and go into insert mode
 - /string search for string
 - n do the last command again
 - jkl; cursor movement
 - A add text at the end of the line
 - <line>G go to the line <line>
 - G go to the last line in the file

@server



vi - cheat sheet (cont'd)

- ▶ (3) Command line mode
 - : go to a command prompt
 - :wq save (write & quit)
 - :q! quit and discard changes
 - :<num> go to line number <num>
 - :r <file> read <file> into the current file
 - :1,\$s/find/replace/g global <replace> of <find>
 - :help give help
- ▶ There is a vi tutor (pretty basic). To use, go to a r/w directory and:
 - # cp /usr/share/vim/vim56/tutor/tutor .
 - # vi tutor

@server