


A Linux for S/390 Primer


SHARE 98

Sessions 5508 / 9204
5509 / 9205



Who I am, and what I do

- Mark Post (mark.post@eds.com)
- Senior Infrastructure Specialist (Systems Programmer)
- EDS, Auburn Hills, Michigan
- Worked with IBM mainframes since entering college.
- Supported MVS and VM at GM Proving Ground.
- EDS Corporate Repository Support since 1993.
- Involved with Linux/390 since May of 2000.
- Co-author of IBM RedBook: Linux for zSeries and S/390: Distributions, SG24-6264



Credits

David Boyes of Sine Nomine Associates

Jim Elliott of IBM USA


Neale Ferguson of SoftwareAG

Simon Williams of IBM Australia

Peter Mattis and Spencer Kimball for the GIMP



And ...




The logo for EDS, featuring a large blue circle on the left containing a white, stylized letter 'E'. To the right of the circle, the letters 'D' and 'S' are written in a blue, serif font. The background is black with a blue and orange gradient on the left side.

EDS

" . . . the recognized global leader in ensuring clients achieve superior value in the Digital Economy."

03/05/2002

4






EDS

" . . . the recognized global leader in ensuring clients achieve superior value in the Digital Economy."

03/05/2002

5



Is it 'real' or 'pure' Linux?

Yes, it is.


IBM provided the architecture-specific patches to the vanilla Linux kernel source and the tool chain (gcc, binutils, glibc, gdb, and strace).

Anyone can download the kernel source, apply the IBM patches, and go through the normal 'make' sequence. The patches are being incorporated into the official source trees.

The same applies to gcc, glibc, etc.

Linux/390 is ASCII, not EBCDIC.

No translations needed.




“not Linux-like, not Linux-lite, it's truly Linux!” - IBM

Note that IBM does not have its own Linux distribution.

They are merely acting as kernel/tool architecture maintainers.

Distributions are (currently) only provided by SuSE, Turbolinux, Caiman, and Red Hat. Thinking Objects in Germany, and Marist College also have one, but they will probably have a very small ‘market share’ in comparison.



```
File Sessions Options Help
glt3901:~ # uname -a
Linux glt3901 2.2.16 #1 SMP Wed Nov 8 10:57:03 GMT 2000 s390 unknown
glt3901:~ #
```

```
9:50am up 13 days, 21:38, 2 users, load average: 1.65, 1.24, 0.94
2 processes: 72 sleeping, 2 running, 7 zombie, 1 stopped
PU states: 42.0% user, 15.2% system, 0.0% nice, 42.7% idle
vm: 127384K av, 122500K used, 4884K free, 77112K shrd, 9292K buff
wap: 230380K av, 380K used, 230000K free 44916K cached

PID USER PRI NI SIZE RSS SHARE STAT LIB %CPU %MEM TIME COMMAND
661 root 17 0 5896 5896 1772 R 0 41.4 4.6 0:02 cc1
0 12.8 0.7 0:00 make
0 7.0 1.0 0:00 cpp
0 5.1 2.3 1:27 gkrellm
0 4.0 2.3 129:31 nmbd
0 4.0 0.8 0:22 top
0 0.7 3.5 0:31 kum
0 0.7 3.4 0:05 konsole
0 0.5 2.1 0:05 x3270
0 0.3 0.5 0:00 as
0 0.1 1.0 0:09 glt3901
0 0.1 0.3 0:00 Mon 16 Apr
0 0.0 0.1 2:21 9:50 52
0 0.0 0.0 0:00
0 0.0 0.0 0:02
```

```
x3270-4 dadnn3k.sys.eds.com
File Options
Display Filter View Print Options Help
-----
SDSF SYSLOG 111.101 SSP1 SSP1 04/16/2001 0w 4,858 COLUMNS 1 80
COMMAND INPUT ==>
D 713 00000090 2: 00000058 3: 0
D 713 00000090 4: 00000001 5: 0
D 713 00000090 6: 0017A038 7: 0
D 713 00000090 8: 00180208 9: 0
D 713 00000090 A: 0000007B B: 8
D 713 00000090 C: 00172AC0 D: 0
D 713 00000090 E: 00174720 F: 0
E 713 00000090 END OF SYMPTOM DUMP
N 00200000 SSP1 01106 07:34:33.65 STC00162 00000090 IEA9931 SYSMDUMP TAKEN T
M 00000000 SSP1 01106 07:34:34.05 STC00162 00000090 IEF4501 MVSNFS MVSNFS -
E 715 00000090 TIME=07.34.34
N 00000000 SSP1 01106 07:34:34.09 STC00162 00000090 IEF4041 MVSNFS - ENDED -
N 40000000 SSP1 01106 07:34:34.09 STC00162 00000090 $HASP395 MVSNFS ENDED
N 00000000 SSP1 01106 07:34:34.25 000000290 IEA9891 SLIP TRAP ID=X33
N 00000000 SSP1 01106 07:34:42.69 STC00161 00000090 IEF4041 FTPSERVE - ENDED
N 40000000 SSP1 01106 07:34:42.69 STC00161 00000090 $HASP395 FTPSERVE ENDED
N 00000000 SSP1 01106 07:34:42.72 000000290 IEA9891 SLIP TRAP ID=X33
N 02000000 SSP1 01106 07:34:42.74 STC00161 00000090 $HASP250 FTPSERVE PURGED
NC00000000 SSP1 01106 07:35:02.07 INSTREAM 00000290 LOGON
NR40000000 SSP1 01106 07:35:26.78 STC00156 00000090 IST5261 ROUTE FAILED FRD
N 40800000 SSP1 01106 07:35:44.66 00000090 IKJ6061 TSOLGDN REJECTE
N 00000000 SSP1 01106 07:35:49.99 TSU00159 00000290 IEA6311 OPERATOR RZTSGK
N 00000000 SSP1 01106 07:35:51.65 TSU00159 00000090 IEF1261 RZTSGK - LOGGED
N 40000000 SSP1 01106 07:35:51.66 TSU00159 00000090 $HASP395 RZTSGK ENDED
N 02000000 SSP1 01106 07:35:51.69 TSU00159 00000090 $HASP250 RZTSGK PURGED -
N 00000000 SSP1 01106 07:35:57.84 000000290 IEA9891 SLIP TRAP ID=X33
N 02000000 SSP1 01106 07:36:03.84 TSU00167 00000090 $HASP100 RZTSGK ON TSO
N 40000000 SSP1 01106 07:36:03.92 TSU00167 00000090 $HASP373 RZTSGK STARTE
N 00000000 SSP1 01106 07:36:03.93 TSU00167 00000090 IEF1251 RZTSGK - LOGGED
N 00000000 SSP1 01106 07:39:07.40 TSU00167 00000090 IEF1261 RZTSGK - LOGGED
N 40000000 SSP1 01106 07:39:07.43 TSU00167 00000090 $HASP395 RZTSGK ENDED
N 02000000 SSP1 01106 07:39:07.53 TSU00167 00000090 $HASP250 RZTSGK PURGED -
N 02000000 SSP1 01106 07:39:10.97 TSU00168 00000090 $HASP100 RZTSGK ON TSO
N 40000000 SSP1 01106 07:39:11.22 TSU00168 00000090 $HASP373 RZTSGK STARTE
N 00000000 SSP1 01106 07:39:11.23 TSU00168 00000090 IEF1251 RZTSGK - LOGGED
```

glt3901
Mon 16 Apr
9:50 52
CPU0
CPU1
Proc
etc0
Mem
Swap
root T
/tmp
/usr T
/usr/erc
/var
/opt T
/Boot
/root
/home
-/-


The History of Linux for S/390

17 December, 1998, Harry Williams of Marist College starts the Linux-390 mailing list.

<http://www.marist.edu/htbin/wlvindex?linux-390>


Originally called "Linux on VM port." Now "Linux for S/390 port."

Some now-well-known names on the mailing list were very early participants: David Boyes, Neale Ferguson, Rick Troth, Rich Smrcina, Ross Patterson.



Very early discussion/recognition of the scope of the project:
David Boyes, 17 Dec, 1998:


“A later poster questions whether we are providing Linux emulation in CMS or a complete Linux system in a virtual machine. If this project is to succeed, I think we have to concentrate on the complete Linux system approach.”



A number of IBM employees were participating in a personal, i.e., unofficial capacity, just for the sheer coolness-factor.

Some questions as to whether the port should be of Linux or something else: “Should alternatives to Linux be considered? www.netbsd.org, www.freebsd.org, www.openbsd.org” - Richard Corak.


“It seems that the BSD ports are primarily geared toward PCs, whereas Linux seems to have ports for, or porting underway for, a broad range of hardware platforms.”
- Perry Ruiter



Bigfoot Linux

Linus Vepstas formally joins the project, 15 June, 1999, after having been working on his own port to the then-S/370 architecture. His interests are more S/370, S/390 generic, rather than just under VM. His name for the port is ‘Bigfoot Linux,’ or just ‘Bigfoot.’

He winds up doing the bulk of the work on the tool chain to build the Linux/390 kernel: gcc, binutils, etc. He is also instrumental in getting a lot of the modifications accepted into the main source trees for those programs. He spends hundreds of hours of his personal time on the project.




Melinda Varian of the Princeton University Computer Center allows the team significant access to the school's VM resources for development.

Linus and the others have Linux booting 'by hand' to the point of starting user-land processes by 21 June, 1999.

The project is still working on some kind of a boot loader. 6 July 1999, "both Arty [Ecock] & Rick have gotten the Kernel into memory and branched to it, but it then loops."


17 Aug 1999, Neale Ferguson posts a console log of an IPL using "Arty's loader."



Threats and Opportunities

Mid-August, 1999, Amdahl pre-announces 'Dino,' a mainframe version of Linux. Linus Vepstas was both upset and concerned. He wondered where he and the rest of the Open Source community would fit in. The project apparently faded into oblivion.


19 August, 1999, Hercules (S/390 emulator for x86 platform) is mentioned on Linux-390 mailing list as a potential aid to development and testing.



IBM's Christmas Gift of 1999


At the October WAVV Conference. (World Alliance of VSE and VM) “IBM talked about (but did not announce) LINUX/390. The future tech session showed foils about LINUX/390. There will be 2 versions ... 1st is LINUX/390 under VM only and the 2nd will be LINUX/390 on S/390 (bare metal). LINUX/390 under VM is ready now. LINUX/390 on bare metal will take more time (but soon).”
- Jeff Barnard

18 December, 1999, IBM announces it's source modifications to the Linux kernel and tool chain. Linux for S/390 will run under VM, in an LPAR, and 'on the bare metal.'



Even though rumors had been circulating about this for over five months, everyone on the Bigfoot port is a little shocked, including all the IBMers that had been participating. Linas Vepstas in particular is rather bitter because all the time and effort that was invested by the project members was essentially wasted, since IBM's port was done in secret with no sharing of information or resources with the Bigfoot team, and none of the Bigfoot team's work would wind up being merged into the IBM port.


The secrecy on IBM's part was necessary, due to IBM internal politics, but still galling.



S/390 specific code

	Total Lines of Code	S/390 Enabling Code	Percentage of S/390 Code
Linux	2,200,00	45,000	2.00
gcc	1,700,000	9,000	0.50
gdb	1,500,000	8,000	0.50
glibc	1,200,000	5,000	0.40
binutils	800,000	6,000	0.75
strace	41,000	200	0.50

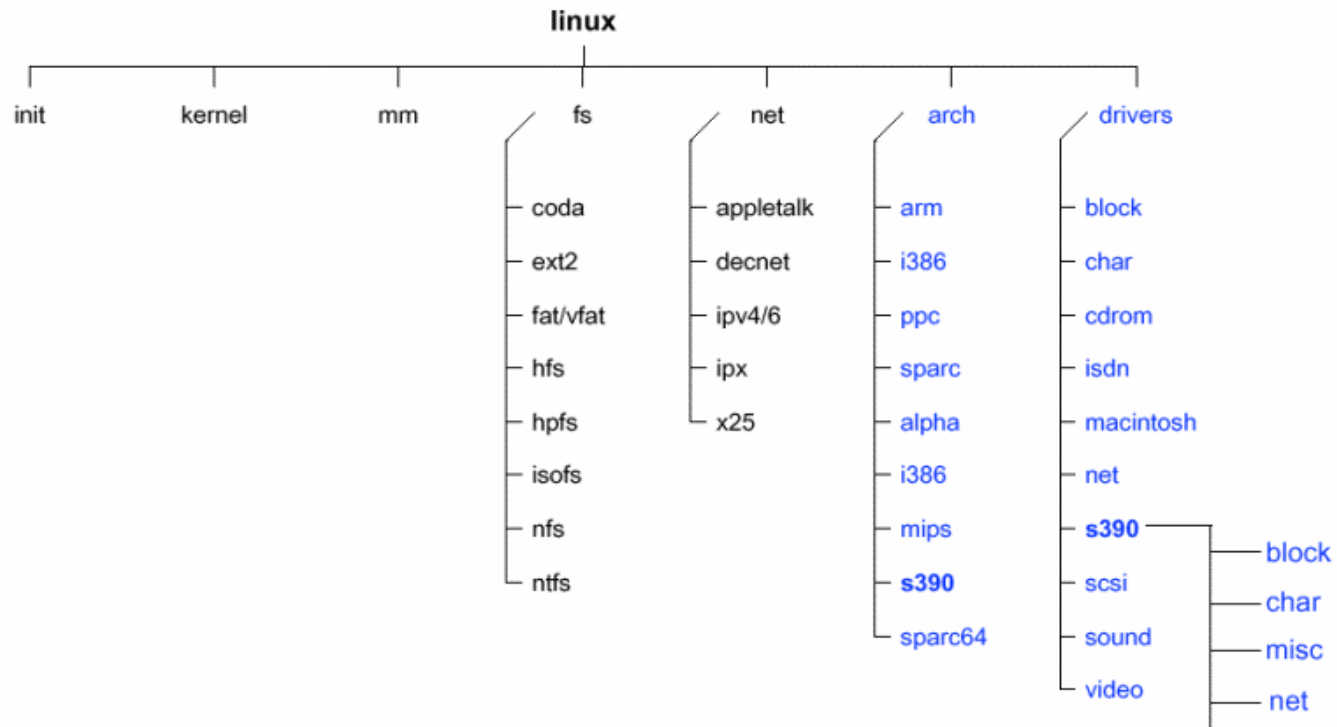
Slide Courtesy of IBM Corporation




Linux Source Tree Changes for S/390

IBM @server zSeries

Linux Source Tree (partial view only):



Slide Courtesy of IBM Corporation




The Marist File System(s)

The IBM group in Boeblingen Germany made a tar backup of a running Linux/390 system to ease initial installation. Loosely based on Red Hat.

There were two 'flavors.' A 'small' file system and a 'large' file system. Both were RPM-based, but the RPM database didn't accurately reflect what was running.

Marist College made these file systems available on their ftp server by 5 January 2000. These file systems have not been maintained over time, and so are very much out of date. Michael Kershaw of Marist College is working on bringing them up to date. (He and I are also working on creating a Slackware/390.)




Development, testing, etc. continued at a rapid pace.

Heavy involvement from the IBM team in Germany.

Numerous bug fixes to kernel, driver modules, gcc, etc. That is still ongoing, but is more directed at moving to the 64-bit architecture and the 2.4.x kernel.

SuSE distribution available first, followed by Turbolinux, Caiman, and then Red Hat.

Slackware, and Debian now in progress.




Running Linux on S/390

IBM e-server zSeries

- Native
- Logically Partitioned (LPAR mode)
- VM/ESA and z/VM
- Virtual Image Facility- Withdrawn
- Integrated Facility for Linux
- z800 - Just announced!

Slide Courtesy of IBM Corporation



Logically Partitioned

IBM @server zSeries

Allows partitioning of *all* resources of a *single* (hardware) System into up to 15 parts (Logical Partitions):

CPU, Storage, I/O-channels, DASD (disks)

Similar to VM/ESA CP-component (Control Program), but not 'visible' as software; integrated as part of hardware system.

'Hypervisor', controlling and managing the partitions.


Allows for *concurrent* usage of different operating systems.

Preferred usage in large systems:

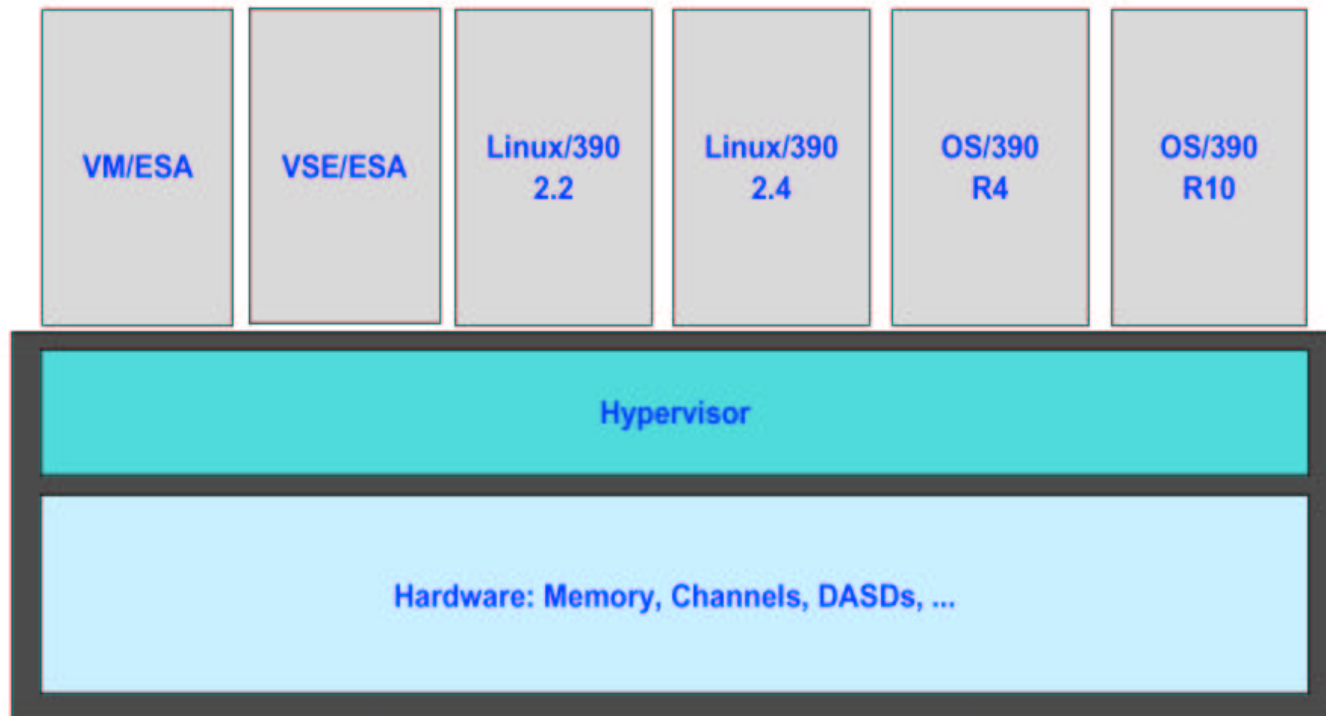
To strictly isolate mixed applications/workloads from each other
(e.g., production / test / development)

To enable 'smooth' software migrations (release to release ...)


Slide Courtesy of IBM Corporation



LPAR



Slide Courtesy of IBM Corporation




What is VM?

Virtual Machine (VM) is an IBM proprietary Operating System.

If you've heard of VMware or plex86, then you have some idea of what VM does (and vice versa).

A Virtual Machine simulates the existence of a dedicated real machine, including processor functions, storage, and input/output resources.

The users/systems/etc. that run 'under' VM are referred to as 'guests.'




VM/ESA and z/VM

Any 'privileged' instructions are intercepted and simulated on behalf of the guests. (Think kernel vs. user space, but at an *instruction* level.)

Generally, VM 'owns' the hardware unless special steps are taken to dedicate something to a guest.

VM allows dynamic creation and destruction of virtual resources, such as disks (DASD), CTC connections, storage, I/O paths, etc.

VM will allow you to control how much CPU, virtual storage, disk space etc., that a given guest will receive, as well as its priority.




David Boyes and Test Plan Charlie

This is probably the most famous story about Linux/390 running under VM.

It is also the reason why so many people see Linux/390 under VM as the preferred implementation.

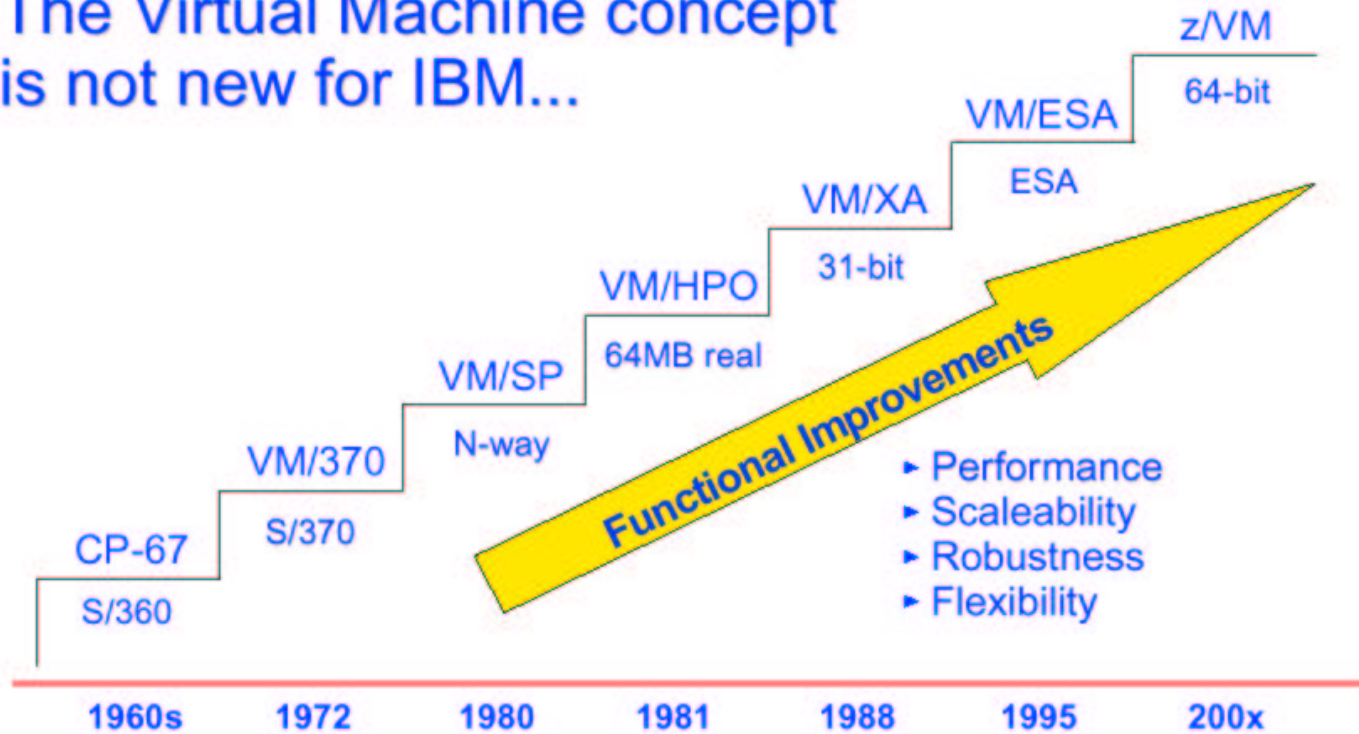
During a consulting engagement, David brought up 41,400 Linux instances under VM before VM ran out of resources. The images were running Apache, serving up static web pages.

David has long since surpassed that record by bringing up 97,943 images on a dedicated 9672-ZZ7. (Test Plan Omega)




VM Evolution

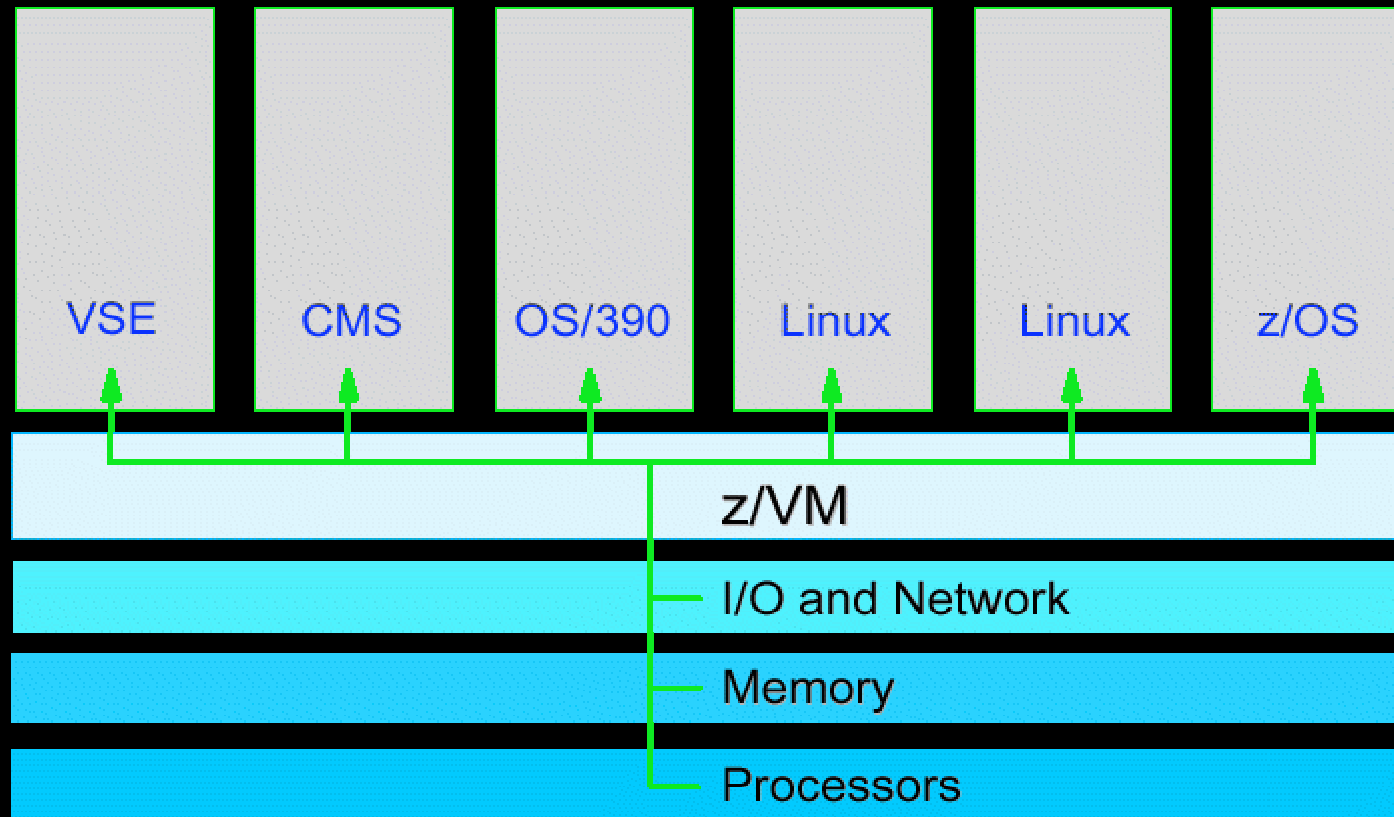
The Virtual Machine concept
is not new for IBM...




Slide Courtesy of IBM Corporation

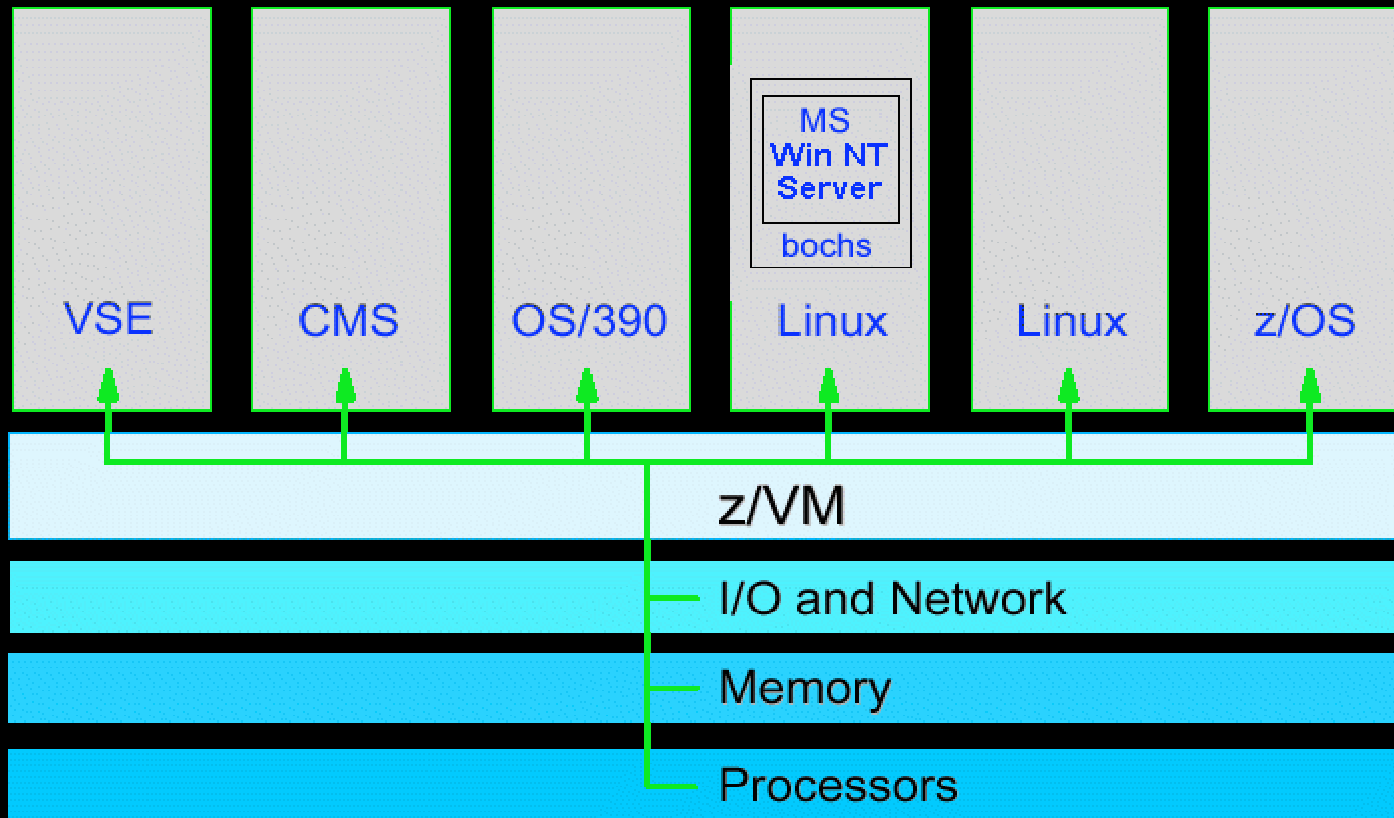


VM Partitioning (Guests)



Slide Courtesy of IBM Corporation





Virtual Image Facility for Linux


Enables multiple Linux images to run on a single S/390 system

- VIF runs natively or in an LPAR
- Does not require VM skills, but in reality is a stripped-down version of VM.
- Supports ten to hundreds of images, depending on workload
- Less capable in terms of management ability than VM, with less granularity of control.

Complements the S/390 Integrated Facility for Linux

Much more attractive engine-based pricing terms and conditions than VM.

Only runs Linux for S/390.



Integrated Facility For Linux


Additional G5 / G6 / MP3000 / z800 / z900 engines dedicated to Linux workloads (at roughly 50% of the cost of “normal” CPUs)

Can run a single Linux/390 instance or many, under z/VM 4 (no prior VM, VSE, or OS/390).

Begin deployment or consolidation of UNIX/Linux workloads to S/390 immediately

Traditional S/390 software charges unaffected


- IBM S/390 software and middleware
- Independent Software Vendor product pricing
 - BMC, Candle, Compuware committed



Linux/390 'Oddities'

What's different about running Linux/390?


- No VGA/SVGA monitor, or graphics card
- No console keyboard, sound cards, etc.
- No access to RTC (yet)
- Not all machines have an IEEE FPU (G5 or higher)
- Point-to-point connections require a workaround for Samba to work
 - Channel-to-Channel connections are point-to-point
 - IUCV connections (under VM) are point-to-point
- No floppy drives
- Not all systems have CD drives.



Can you run X, Gnome, KDE (or X, KDE, Gnome) ?

- Yes, but it doesn't make much sense
- An X server manages the local monitor, keyboard, mouse, which is on your local workstation, not the S/390
- Running KDE or Gnome on an S/390 is not a good use of an S/390
- X clients can run on an S/390, but that may or not be a good thing to do.
- IEEE Floating Point Unit will improve performance.

There are far fewer kernel options to worry about.
No support needed for a LOT of hardware types.



So, What's This Thing Look Like?


```
Ready;  
ipl 100 clear  
Linux version 2.2.16 (root@Tape.suse.de) (gcc version 2.95.2 19991024 (release))  
#1 SMP Mon Jul 24 18:48:07 GMT 2000  
Command line is: dasd=100-10f root=/dev/dasda1 ro noinitrd
```

We are running under VM


This machine has no IEEE fpu

Initial ramdisk at: 0x02000000 (8388608 bytes)


```
Detected device 0009 on subchannel 0000 - PIM = 80, PAM = 80, POM = FF  
Detected device 000C on subchannel 0001 - PIM = 80, PAM = 80, POM = FF  
Detected device 000D on subchannel 0002 - PIM = 80, PAM = 80, POM = FF  
Detected device 000E on subchannel 0003 - PIM = 80, PAM = 80, POM = FF  
Detected device 0190 on subchannel 0004 - PIM = F0, PAM = F0, POM = FF  
Detected device 019E on subchannel 0005 - PIM = F0, PAM = F0, POM = FF  
Detected device 0319 on subchannel 0006 - PIM = F0, PAM = F0, POM = FF  
Detected device 0E00 on subchannel 0007 - PIM = 80, PAM = 80, POM = FF  
Detected device 0E01 on subchannel 0008 - PIM = 80, PAM = 80, POM = FF  
Detected device 0191 on subchannel 0009 - PIM = F0, PAM = F0, POM = FF  
Detected device 0101 on subchannel 000A - PIM = F0, PAM = F0, POM = FF  
Detected device 0102 on subchannel 000B - PIM = F0, PAM = F0, POM = FF  
Detected device 0103 on subchannel 000C - PIM = F0, PAM = F0, POM = FF  
Detected device 0104 on subchannel 000D - PIM = F0, PAM = F0, POM = FF  
Detected device 0105 on subchannel 000E - PIM = F0, PAM = F0, POM = FF  
Detected device 0106 on subchannel 000F - PIM = F0, PAM = F0, POM = FF  
Detected device 0107 on subchannel 0010 - PIM = F0, PAM = F0, POM = FF  
Detected device 0108 on subchannel 0011 - PIM = F0, PAM = F0, POM = FF  
Detected device 0109 on subchannel 0012 - PIM = F0, PAM = F0, POM = FF  
Detected device 010F on subchannel 0013 - PIM = F0, PAM = F0, POM = FF  
Detected device 0100 on subchannel 0014 - PIM = F0, PAM = F0, POM = FF
```




```
Detected device 0300 on subchannel 0015 - PIM = F0, PAM = F0, POM = FF
Highest subchannel number detected: 22
SenseID : device 0009 reports: Dev Type/Mod = 3215/00
SenseID : device 000C reports: Dev Type/Mod = 2540/00
SenseID : device 000D reports: Dev Type/Mod = 2540/00
SenseID : device 000E reports: Dev Type/Mod = 1403/00
SenseID : device 0190 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 019E reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0319 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0E00 reports: Dev Type/Mod = 3088/08
SenseID : device 0E01 reports: Dev Type/Mod = 3088/08
SenseID : device 0191 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0101 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0102 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0103 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0104 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0105 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0106 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0107 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0108 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0109 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 010F reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0100 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
SenseID : device 0300 reports: CU Type/Mod = 3990/E9, Dev Type/Mod = 3390/0A
Calibrating delay loop... 89.91 BogomIPS
Memory: 119328k/131072k available (1116k kernel code, 4k reserved, 2432k data, 0k init)
Dentry hash table entries: 16384 (order 5, 128k)
Buffer cache hash table entries: 131072 (order 7, 512k)
Page cache hash table entries: 32768 (order 5, 128k)
16 areas reserved for debugging information
reserved 4 areas of 4 pages for debugging ccwcache
```




```
VFS: Diskquotas version dquot_6.4.0 initialized
POSIX conformance testing by UNIFIX
Detected 2 CPU's
Boot cpu address 0
cpu 0 phys_idx=0 vers=FF ident=043797 machine=9672 unused=0000
cpu 1 phys_idx=1 vers=FF ident=043797 machine=9672 unused=0000
Linux NET4.0 for Linux 2.2
Based upon Swansea University Computer Society NET3.039
NET4: Unix domain sockets 1.0 for Linux NET4.0.
NET4: Linux TCP/IP 1.0 for NET4.0
IP Protocols: ICMP, UDP, TCP
TCP: Hash tables configured (ehash 131072 bhash 65536)
Starting kswapd v 1.5
pty: 256 Unix98 ptys configured
RAM disk driver initialized: 16 RAM disks of 32768K size
loop: registered device at major 7
md driver 0.36.6 MAX_MD_DEV=4, MAX_REAL=8
dasd:initializing...
dasd:Registered successfully to major no 94
dasd(eckd):ECKD discipline initializing
dasd:Registered ECKD discipline successfully
dasd(fba):FBA discipline initializing
dasd:Registered FBA discipline successfully
dasd(eckd):0100 on sch 20: 3390/0A(CU:3990/01) Cyl:50 Head:15 Sec:224
dasd(eckd):0100 on sch 20: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0100 on subchannel 20 (ECKD) is /dev/dasda (94:0)
dasd(eckd):0101 on sch 10: 3390/0A(CU:3990/01) Cyl:15 Head:15 Sec:224
dasd(eckd):0101 on sch 10: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0101 on subchannel 10 (ECKD) is /dev/dasdb (94:4)
```



```
dasd(eckd):0102 on sch 11: 3390/0A(CU:3990/01) Cyl:35 Head:15 Sec:224
dasd(eckd):0102 on sch 11: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0102 on subchannel 11 (ECKD) is /dev/dasdc (94:8)
dasd(eckd):0103 on sch 12: 3390/0A(CU:3990/01) Cyl:25 Head:15 Sec:224
dasd(eckd):0103 on sch 12: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0103 on subchannel 12 (ECKD) is /dev/dasdd (94:12)
dasd(eckd):0104 on sch 13: 3390/0A(CU:3990/01) Cyl:150 Head:15 Sec:224
dasd(eckd):0104 on sch 13: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0104 on subchannel 13 (ECKD) is /dev/dasde (94:16)
dasd(eckd):0105 on sch 14: 3390/0A(CU:3990/01) Cyl:75 Head:15 Sec:224
dasd(eckd):0105 on sch 14: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0105 on subchannel 14 (ECKD) is /dev/dasdf (94:20)
dasd(eckd):0106 on sch 15: 3390/0A(CU:3990/01) Cyl:1500 Head:15 Sec:224
dasd(eckd):0106 on sch 15: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0106 on subchannel 15 (ECKD) is /dev/dasdg (94:24)
dasd(eckd):0107 on sch 16: 3390/0A(CU:3990/01) Cyl:1500 Head:15 Sec:224
dasd(eckd):0107 on sch 16: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0107 on subchannel 16 (ECKD) is /dev/dasdh (94:28)
dasd(eckd):0108 on sch 17: 3390/0A(CU:3990/01) Cyl:1300 Head:15 Sec:224
dasd(eckd):0108 on sch 17: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0108 on subchannel 17 (ECKD) is /dev/dasdi (94:32)
dasd(eckd):0109 on sch 18: 3390/0A(CU:3990/01) Cyl:80 Head:15 Sec:224
dasd(eckd):0109 on sch 18: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x0109 on subchannel 18 (ECKD) is /dev/dasdj (94:36)
dasd(eckd):010F on sch 19: 3390/0A(CU:3990/01) Cyl:320 Head:15 Sec:224
dasd(eckd):010F on sch 19: 3390/0A (CU: 3990/01): Configuration data read
dasd: devno 0x010F on subchannel 19 (ECKD) is /dev/dasdp (94:60)
dasd:waiting for responses...
dasd(eckd):/dev/dasda (0100): capacity (4kB blks): 36000kB at 48kB/trk
dasd(eckd):/dev/dasdb (0101): capacity (4kB blks): 10800kB at 48kB/trk
```



```
dasd(eckd) :/dev/dasdc (0102): capacity (4kB blks): 25200kB at 48kB/trk
dasd(eckd) :/dev/dasdd (0103): capacity (4kB blks): 18000kB at 48kB/trk
dasd(eckd) :/dev/dasde (0104): capacity (4kB blks): 108000kB at 48kB/trk
dasd(eckd) :/dev/dasdf (0105): capacity (4kB blks): 54000kB at 48kB/trk
dasd(eckd) :/dev/dasdg (0106): capacity (4kB blks): 1080000kB at 48kB/trk
dasd(eckd) :/dev/dasdh (0107): capacity (4kB blks): 1080000kB at 48kB/trk
dasd(eckd) :/dev/dasdi (0108): capacity (4kB blks): 936000kB at 48kB/trk
dasd(eckd) :/dev/dasdj (0109): capacity (4kB blks): 57600kB at 48kB/trk
dasd(eckd) :/dev/dasdp (010F): capacity (4kB blks): 230400kB at 48kB/trk
dasd:initialization completed
xpraminfo:initializing:
xpraminfo: number of devices (partitions): 1
xpraminfo: size of partition 0 to be set automatically
xpraminfo: hardsector size: 4096B
xpramerror:No or not enough expanded memory available
xpramerror:Giving up xpram
channel: 2 Parallel channel found - 0 ESCON channel found
ctc0: read dev: 0e00 irq: 0007 - write dev: 0e01 irq: 0008
Partition check:
  dasda:(nonl)/      : dasda dasda1
  dasdb:(nonl)/      : dasdb dasdb1
  dasdc:(nonl)/      : dasdc dasdc1
  dasdd:(nonl)/      : dasdd dasdd1
  dasde:(nonl)/      : dasde dasde1
  dasdf:(nonl)/      : dasdf dasdf1
  dasdg:(nonl)/      : dasdg dasdg1
  dasdh:(nonl)/      : dasdh dasdh1
  dasdi:(nonl)/      : dasdi dasdi1
  dasdj:(nonl)/      : dasdj dasdj1
  dasdp:(LNX1)/ x7fff: dasdp dasdp1
```



```
VFS: Mounted root (ext2 filesystem) readonly.
```

```
Freeing unused kernel memory: 0k freed
```

```
INIT: version 2.78 booting
```

```
Running /sbin/init.d/boot
```

```
Mounting /proc device done
```

```
Activating swap-devices in /etc/fstab...
```

```
Adding Swap: 230380k swap-space (priority -1)  
done
```

```
Checking file systems...
```

```
Parallelizing fsck version 1.18a (11-Nov-1999)
```

```
/dev/dasda1: clean, 1680/9024 files, 8377/8997 blocks
```

```
/dev/dasdc1: clean, 64/6304 files, 274/6297 blocks
```

```
/dev/dasdd1: clean, 144/4512 files, 3439/4497 blocks
```

```
/dev/dasde1: clean, 593/27008 files, 4661/26997 blocks
```

```
/dev/dasdi1: clean, 5855/117248 files, 31712/233997 blocks
```

```
/dev/dasdh1: clean, 58116/135072 files, 237082/269997 blocks
```

```
/dev/dasdg1: clean, 16/135072 files, 4256/269997 blocks
```

```
/dev/dasdj1: clean, 1402/14400 files, 13318/14397 blocks
```

```
/dev/dasdf1: clean, 25/13504 files, 449/13497 blocks
```

```
done
```

```
Mounting local file systems...
```

```
proc on /proc type proc (rw)
```

```
/dev/dasdc1 on /home type ext2 (rw)
```


```
/dev/dasdd1 on /lib type ext2 (rw)
```

```
/dev/dasde1 on /opt type ext2 (rw)
```

```
/dev/dasdf1 on /root type ext2 (rw)
```

```
/dev/dasdg1 on /tmp type ext2 (rw)
```


```
/dev/dasdh1 on /usr type ext2 (rw)
```




```
/dev/dasdi1 on /usr/src type ext2 (rw)
/dev/dasdj1 on /var type ext2 (rw)
done
Mounting /dev/pts.
done
Setting up timezone data done
Setting up loopback device done
Setting up hostname done
Setting up the CMOS clock
Cannot access the Hardware Clock via any known method.
Use the --debug option to see the details of our search
for an access method.
Done
Running /sbin/init.d/boot.local
failed
Creating /var/log/boot.msg
done

Enabling syn flood protection done
Disabling IP forwarding done
INIT: Entering runlevel: 2

Master Resource Control: previous runlevel: N, switching
to runlevel: 2
Setting up network device ctc0
ctc0: connected with remote side
done
Setting up routing (using /etc/route.conf)
done
Starting RPC portmap daemon
done
```




```
Starting scanlogd
done
Starting syslog services
done
Starting NIS+ services:
  keyserv done
Initializing random number generator done
Starting NFS server
done
Starting service httpd
done
Starting service at daemon:
done
Starting INET services (inetd)
done
Starting lpd
done
Initializing SMTP port. (sendmail)
done
Starting SMB services:
done
Starting snmpd:
done
Starting xntpd
done
Starting CRON daemon
done
Starting Name Service Cache Daemon
done
```




Master Resource Control: runlevel 2 has been reached

Give root password to login:

xxxxxxxx

bash-2.04# ps ax

PID	TTY	STAT	TIME	COMMAND
1	?	S	0:00	init
2	?	SW	0:00	[kflushd]
3	?	SW	0:00	[kupdate]
4	?	SW	0:00	[kpiod]
5	?	SW	0:00	[kswapd]
123	?	S	0:00	/sbin/portmap
135	?	S	0:00	/usr/sbin/scanlogd
144	?	S	0:00	/usr/sbin/syslogd
148	?	S	0:01	/usr/sbin/klogd -c 1
161	?	S	0:00	/usr/sbin/keyserv
204	?	S	0:00	/usr/sbin/rpc.mountd
205	?	S	0:00	/usr/sbin/rpc.nfsd
221	?	S	0:04	/usr/sbin/httpd -f /etc/httpd/httpd.conf -D MODULES -
230	?	S	0:00	/usr/sbin/atd
237	?	S	0:00	/usr/sbin/inetd
248	?	S	0:00	/usr/sbin/lpd
263	?	S	0:00	/usr/sbin/fcgi- -f /etc/httpd/httpd.conf -D MODULES -
269	?	S	0:00	/usr/sbin/httpd -f /etc/httpd/httpd.conf -D MODULES -
285	?	S	0:00	sendmail: accepting connections
287	?	S	0:00	/usr/sbin/smbd -D
288	?	S	0:00	/usr/sbin/nmbd -D
290	?	S	0:00	/usr/sbin/snmpd -f
309	?	SL	0:00	/usr/sbin/xntpd
316	?	S	0:00	/usr/sbin/cron
328	?	S	0:00	/usr/sbin/nscd
329	?	S	0:00	/usr/sbin/nscd



```
330 ?      S      0:00 /usr/sbin/nscd
331 ?      S      0:00 /usr/sbin/nscd
332 ?      S      0:00 /usr/sbin/nscd
333 ?      S      0:00 /usr/sbin/nscd
334 ?      S      0:00 /usr/sbin/nscd
362 console S      0:00 bash
363 console R      0:00 ps ax
```


```
bash-2.04# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/dasda1	34M	32M	684k	98%	/
/dev/dasdc1	24M	292k	22M	1%	/home
/dev/dasdd1	17M	13M	3.3M	80%	/lib
/dev/dasde1	102M	15M	82M	15%	/opt
/dev/dasdf1	51M	92k	48M	0%	/root
/dev/dasdg1	1.0G	28k	985M	0%	/tmp
/dev/dasdh1	1.0G	909M	76M	92%	/usr
/dev/dasdi1	900M	109M	744M	13%	/usr/src
/dev/dasdj1	54M	50M	1.5M	97%	/var

```
bash-2.04#
```

```
exit
```

```
Give root password to login:
```




Current Status

Three largest Linux distributors have distributions generally available (GA). Debian.org is working on its port.

Over 400 Open Source packages ported, many with no modifications needed to configuration scripts, and very few requiring source modifications. The most common changes required are to config.sub and config.guess from libtools. More and more these have already been done.


A number of vendors have announced the availability of their products on Linux/390.



A number of companies have Linux/390 in production. Many more are testing/evaluating it.

IBM has already contributed the source patches to support the new zSeries hardware, which is a 64-bit architecture. This is currently undergoing testing by a number of mailing list participants. Neale Ferguson of SoftwareAG is particularly active here. The patches are based on the 2.4.x kernel. New distributions are expected sometime in the next 6 months.

IBM *has* created a fix for the 100Hz timer pop problem.



What's Missing?

- * CLAW support (Cisco routers, etc.)
- * IBM 3480/3490 tape drive support
- * Ability to vary Linux/390 volumes online to OS/390.


LILO equivalent (being discussed)

File system support for native OS/390 volumes


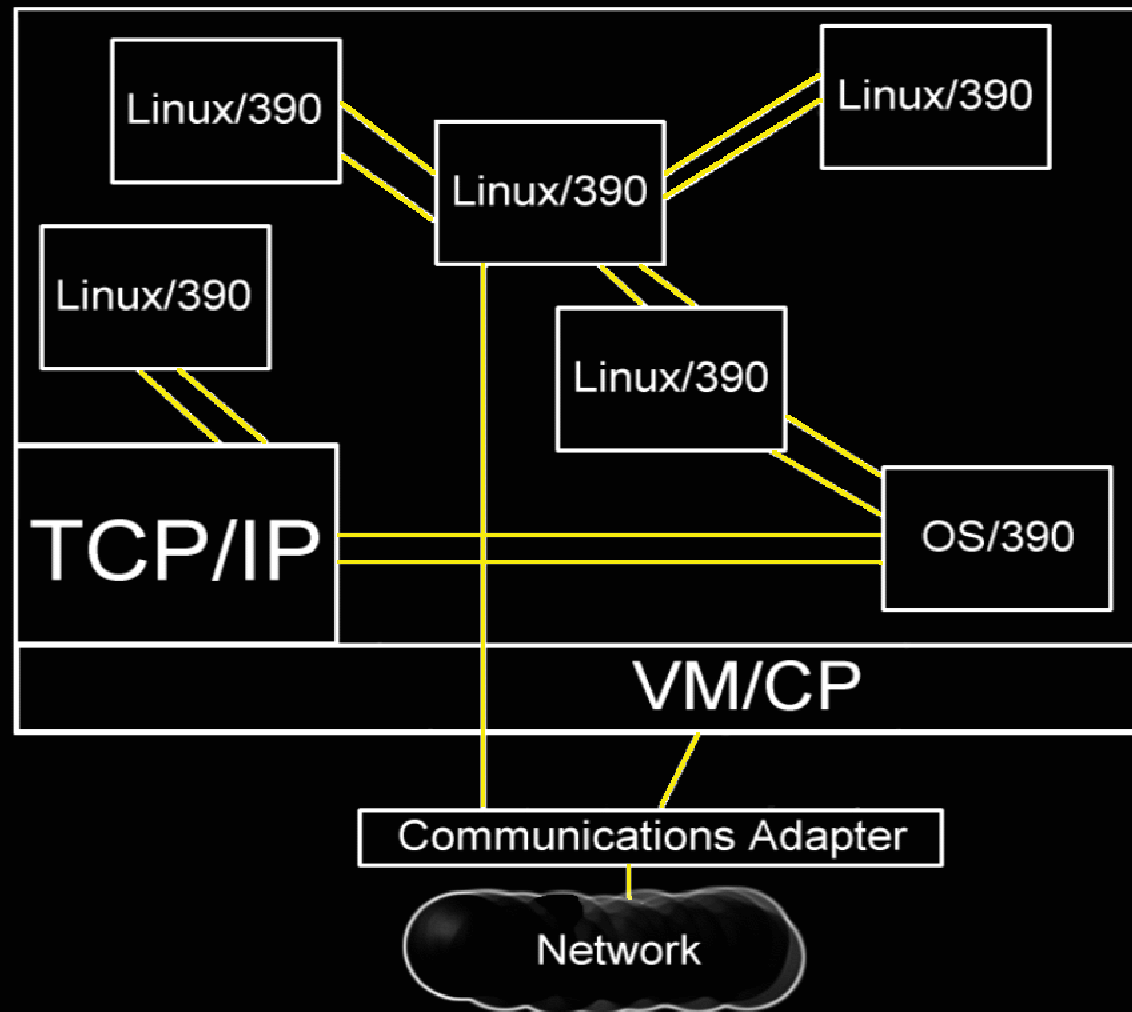
- * (CMS minidisk file system support.)

OS/390 and VM support of ext2 formatted disks.

* = Now available

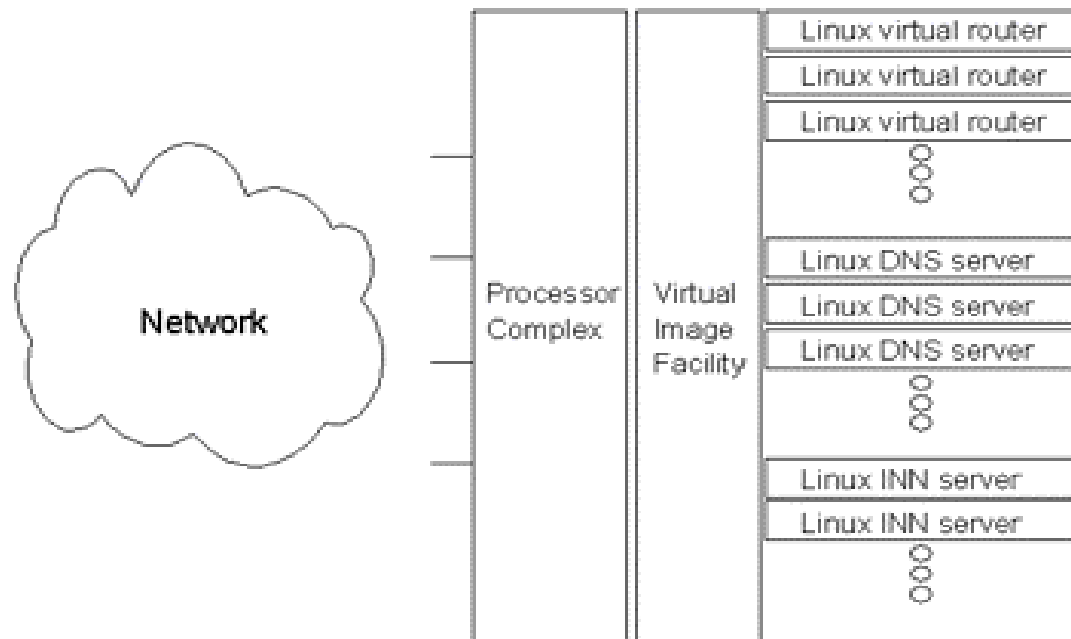


How Does This Thing (net)Work?




Virtual Linux Server Environment

Running on S/390® Virtual Image Facility™ for LINUX®



Slide Courtesy of IBM Corporation



Linux/390 servers for \$500*


*Price claim based on z900 configuration for Linux (\$ 1,200,000) running 2500 separate instances of Linux. Storage not included. The applications running in the 2500 Linux server farm are a mix of InterNetNews (INN) servers and Domain Name Servers (DNS) running on the Virtual Image Facility for Linux. There is also a small number of Linux images providing virtual routing functions in the configuration.

The specific hardware configuration used in this example is IBM zSeries 900 Model 101 with 3 Integrated Linux Facilities. The following Feature Codes are used in the configuration:

1051	1 Way processor
0997	Integrated Linux Facility (3 features)
2324	Escon ports
8051	5 GB memory
2365	OSA Express
2023	I/O Cage
0061	Hardware Master Console
0083	Service Element

Specific IBM software licenses for this configuration include S/390 Virtual Image Facility for Linux with one year service and support

<http://www.ibm.com/servers/eserver/zseries/linuxconfig/>



S/390 Strengths


Up until recently, S/390 CPUs were the fastest *CISC* processors available. Still pretty decent, and IBM will continue to improve them. Up to 16/system.

Hardware Reliability/Redundancy

‘Spare’ CPUs, Channels, Power, Cooling
MTBF Measured in Decades

5GB - 64GB Real Storage

Capacity Upgrade on Demand (Non-Disruptive)




I/O Throughput

- Up to 256 Channels (what's a channel?)
- 17 MB/s - 70 MB/s.
- 1,000 - 4,800 SIOs/s (Start I/O Commands)
- 9Km - 100Km to Devices

262,144 - 1,572,864 Devices (Maximum)


Multiple Paths to Devices

Caching Controllers (4GB, 64MB Non-Volatile)



Who's Using it?


- An unnamed “East Coast telephone company”
- Grede Foundries
- Telia ISP
- Sonera Entrum ISP
- Infocrossing (US ASP)
- coreFusion (Canadian ASP)
- Boscov's Department Stores Inc.
- Gerling AG, Insurance
- Korean Airlines
- Winnebago Industries
- NYSE (Security Industry Automation Corp. (SIAC))
- Newell Rubbermaid
- DreamBall Co.
- Banco Mercantile of Venezuela
- The Credit Index
- sendmail.com



Who's Using it? (2)

- Mutual of Omaha
- Lexis Nexis
- Ford Motor Company
- United Parcel Service
- RiteAid Pharmacies
- CSK Network Systems Corp. (Japan)
- Payless Shoes
- Safeway
- Colorado State University
- Marist College
- TimeWarner Cable
- Dunn and Bradstreet
- Merrill Lynch

Others not talking because of competitive advantage.




Who's Interested in Linux/390?

Colleges

- California Polytechnic State University (Cal Poly)
- Columbia University (New York)
- Harvard
- Mayo Clinic
- MIT
- Princeton
- Rutgers
- Stanford
- University of Oxford (UK)

Linux Distributors

- SuSE
- TurboLinux
- Caiman (Korea)
- RedHat
- Thinking Objects / Millenux (Germany)
- Debian.org




Government Agencies

- Air Force Arnold Engineering Development Center
- Army (Pentagon)
- City of Buenos Aires
- California Business, Transportation and Housing Agency
- Defense Information Systems Agency
- DOE
- EPA
- FAA
- Hampshire County Council, UK
- Idaho Natl. Engineering and Environmental Lab
- IRS
- NASA
- NIH
- Orange County, California
- Reserve Bank of Australia
- SEC
- Smithsonian Museum
- Treasury Department
- Washington State Courts

Software Companies

- Action Software
- Attachmate
- BMC Software
- Candle
- Computer Associates (CA)
- Compuware
- Dignus Systems
- FDR/Innovation Software
- Macro4
- SAS
- Software AG
- Velocity Software




Banks / Financial Services

- ABN AMRO (Standard Federal Bank)
- Bank of America
- Bank West (Australia)
- Citicorp
- Colonial Bank
- Dresdner Bank (Germany)
- Fidelity Investments
- Goldman Sachs
- Lexis-Nexis
- Merrill Lynch
- Norges Bank (Norway)
- Royal Bank (Canada)
- Wells Fargo

Manufacturers

- 3M
- Amdahl
- Boeing
- Daimler-Chrysler
- EMC
- Ford Motor Company
- Fujitsu
- Lucent
- Motorola
- StorageTek



Communication Companies


- Ameritech
- Alltel
- Cisco
- WorldCom

Healthcare Companies

- Blue Cross Blue Shield of South Carolina
- Delta Dental of New Jersey
- Quest Diagnostics

Insurance Companies

- Erie Insurance Group
- Heritage Insurance
- Insurance Corporation of British Columbia
- Mutual of Omaha
- Northwestern Mutual
- Oregon Mutual
- Provident Mutual




Services Companies

- Cap Gemini Ernst & Young
- Computer Sciences Corporation (CSC)
- EDS
- IBM Global Services (IGS)

Media / Publishing

- Atlanta Journal-Constitution
- Financial Times
- McGraw-Hill




Retail Sales Companies

- Nordstrom's
- Payless Shoes
- Rite Aid
- Safeway

Space Industries

- Thiokol
- United Space Alliance




Who's **NOT** Interested?

Informix - acquired by IBM last year.

Oracle - Maybe, maybe not

Sybase


And many more. Sigh.



IBM

Why is IBM doing this?

- **Money.** IBM is a for-profit corporation.
- IBM sells hardware. Software drives hardware sales.
- **Money.**
- Linux/390 will sell more hardware (and has already).
- **Money.**
- It positions them nicely against Microsoft, Sun, HP, etc.
- Did I mention **money**?
- They get the same advantages of Open Source as everyone else.
- It upsets Microsoft.
- They get to do things they couldn't do otherwise, and it saves them **money**.




How committed is IBM to Linux?

IBM budgeted 1 billion dollars for Linux development (on all platforms, not just S/390) and support for last year alone.

IBM has budgeted another 300 million dollars over the next three years on top of this for creating 'Linux services.' Presumably this includes the new Linux developer lab, etc.

Has committed to 'Linux enable' all their hardware platforms.

Has said that AIX may go away in favor of Linux.



Thinkpad

- Certification on 5 commercial models
- Pre-install available

Desktop

- Select models certified

Intellistation

- All 2D models certified

Network Station

- Red Hat toolkit on web for 2200 & 2800

NetVista

- Internet appliance device

IBM Java Virtual Machine 1.2.2

- All 4 distributions

IBM VisualAge for Java Pro

- with JDK 1.3

xSeries

- 90-day no-charge start-up support
- Support from 4 major distributions
- Pre-install available
- 125+ ServerProven applications

zSeries

- Runs native, LPAR, on z/VM
- Support from 3 major distributions (SuSE, Red Hat, TurboLinux)

iSeries^{SCD}

- LPAR
- Support from 3 major distributions

pSeries

- Select models
- Support from 3 major distributions (SuSe, RedHat, TurboLinux)

Linux Cluster

- up to 64 Netfinity 4500R nodes

IBM MQSeries 5.1

- Technology Release

IBM DB2 UDB 7.1

- All 4 distributions

Lotus Domino 5.03

- All 4 distributions

IBM WebSphere Commerce Suite


- Red Hat Now; SuSE, TurboLinux, Caldera soon

IBM WebSphere Application Server 3.02

- Standard & advanced on RedHat & Caldera
- HomePage Builder V4



Slide Courtesy of IBM Corporation



Where to Get More Information

<http://linuxvm.org>

The 'official' web page of the Linux-390 mailing list. I'm the webmaster, so I'm biased.

<http://linuxvm.org/Info/l390link.html>


Over 780 URLs (and growing) relating to Linux and Linux/390.

<http://www.marist.edu/htbin/wlvindex?linux-390>

Linux/390 mailing list archive and subscription form.

<http://linuxtoday.com/stories/11028.html>

12 October, 1999; first known article about Linux/390. David Boyes' corrections to and comments about it can be found there or in the mailing list archives.




Where to Get More Information

<http://www.ibm.com/servers/eserver/zseries/os/linux/>
IBM's official entry point for Linux for S/390 and Linux for zSeries

<http://www10.software.ibm.com/developerworks/opensource/linux390/index.shtml>

IBM's developer information, includes all the source patches to the kernel, gcc, etc., necessary to build your own Linux/390.




Audience Submitted Questions

For a particular Linux image, what kind of peak performance do you see? Compare the CPU to, say, an Athlon. If I write a program that just uses CPU, how fast will it run relative to a current PC? What are the bottlenecks?

There are ports for RH and Suse LINUX for the S/390...why multiple distro's? SuSE seems to be farther along on the ports.

What about 3rd party software? How source-code compatible is/are the port(s)? Is anything being done to convince vendors to port their ix86 software to the S/390 (like Sun/Java, iPlanet/Netscape servers, Oracle, etc.). Would/are you considering providing S/390 RPMs to something like RPMFIND.NET?




I read that sendmail.com uses an S/390 to do their development for their advanced message server... At what point does an S/390 become cost effective over doing individual machines?

Disk Storage: Does the S/390 VM provide SCSI disk, IDE disk, or a custom disk driver?

Portability: Is/Are the port(s) to VM easier without that pesky 'PC' architecture? (VGA consoles, MMX/MMX-like extensions, etc.)


Speaking of consoles... How does the console work?



How does performance of a particular image tend to degrade as more images are added? Assuming that all the images are being actually used, not sitting idle, how close is it to $1/N$, where N is the number of images?

Say you set up a couple thousand servers on an S/390, each one of them serving up a website or something. How are they addressed by the net? Can each have a separate IP address? If so, do you really have 2,000+ NICs or what?

Can you allocate/control the CPU available to the LINUX VM? If so, when you allocate capacity beyond one processor on the S/390, does it require a SMP kernel, or does the VM still appear as a uniprocessor? Is there a choice? Is there a benefit?



Linux for zSeries

reliability and flexibility



the best of both worlds

Linux for S/390



Powered by **S/390**



A Genuinely Open System for S/390



Open up your enterprise
to a world of new applications
with Linux for S/390



Linux Vollwaschmittel 7 kg-Packung




10,99

€ 5,62

(1 kg = DM 1,57)

für 60 Waschen.

http://www.plus.de/aktion/food2/produkte_linux.htm



IBM


LINUX FOR S/390...



...IT'S COOL.

IBM @business servers. Technology. Innovation. Magic.

To get the magic out of the box and into your business, visit our website at ibm.com/s390/linux



Questions?



The mainframe is back! (Not that it was ever “gone.”)

