



IBM

The Value of VM for Linux

*2000 VM/VSE Technical Conference
Session G67*

*Romney White
VM Development*

*romney@vnet.ibm.com
607-755-8276*

© Copyright
IBM 2000

Topics

IBM

- Linux for S/390 under VM
- Performance
- Productivity
- RAS
- Operation and Management
- Application Examples
- References

© Copyright
IBM 2000

Linux under VM

IBM

- Linux runs under VM (!)
 - ▶ Like any other S/390 operating system
 - ▶ Behaves more like CMS than OS/390
- VM environment eased porting effort
 - ▶ Debugging
 - ▶ Boot from virtual card reader
 - ▶ 3215 console driver
 - ▶ Minidisk driver (device-independent)
- Linux exploits VM interfaces
 - ▶ Virtual CTCA driver for IP connectivity
 - ▶ IUCV driver for IP connectivity

© Copyright
IBM 2000

Linux under VM

IBM

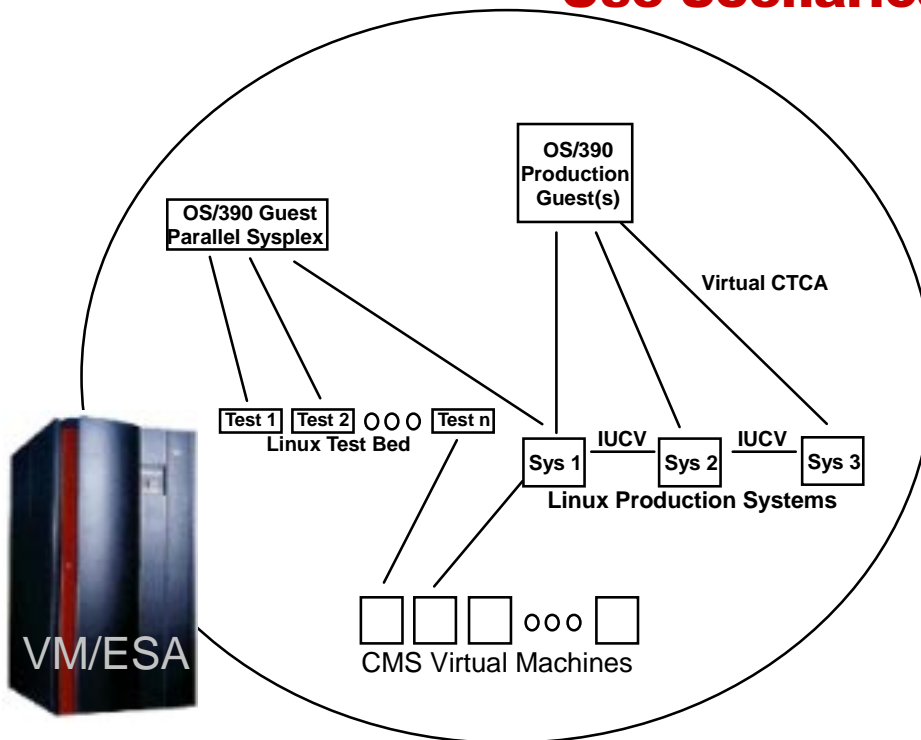
Suggested Minimum Configuration

- 64MB virtual storage
- One 600MB minidisk (CKD, ECKD, FBA)
- Virtual console
 - ▶ Integrated console
 - ▶ 3215 (P/390, standard VM virtual console)
 - Local 3270
 - Telnet
 - SNA
- Network connection
 - ▶ LCS (OSA, 3172)
 - ▶ CTCA (Real, Virtual)
 - ▶ IUCV

© Copyright
IBM 2000

Linux under VM Use Scenarios

IBM



© Copyright
IBM 2000

Linux under VM Rationale

IBM

- Customer unlikely to acquire S/390 just to try Linux
- "Installing" a new virtual machine is trivial
- LPAR option limited
 - ▶ Requires dedicated resources
 - ▶ Difficult to operate and manage
 - ▶ Difficult to use without management permission
- One system is never enough
 - ▶ Test and production
 - ▶ Multiple projects
 - ▶ "One system, one application" mindset

© Copyright
IBM 2000

Linux under VM Benefits



- Performance
- Productivity
- Reliability, Availability, and Serviceability
- Operation and Management



© Copyright
IBM 2000

Linux under VM Value Checklist



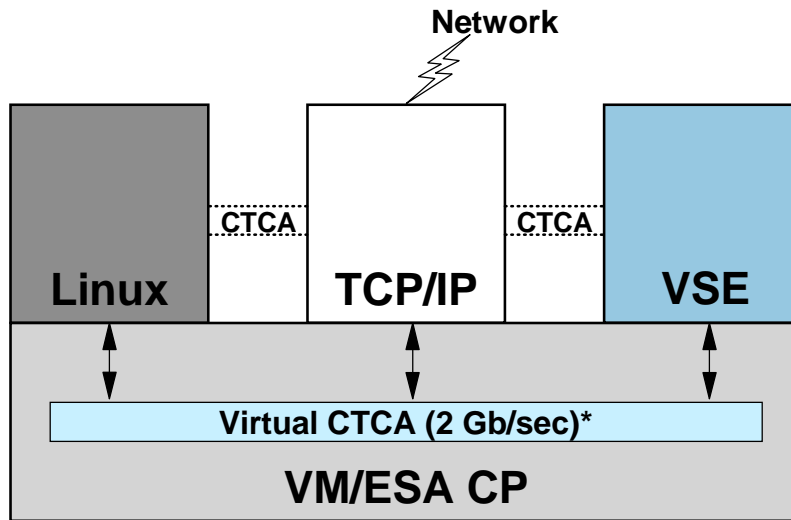
Facilities Provided to Guest OS	OS/390	VSE/ESA	TPF	Linux
Performance				
Virtual Disks in Storage		✓	✓	✓
Minidisk cache		✓	✓	✓
DB2 Guest Sharing (S/390 data space exploitation)		✓	✓	✓
High-performance virtual machine networking	✓	✓	✓	✓
N-way processor exploitation	✓	✓	✓	✓
Page fault handshaking		✓	✓	✓
Transparent S/390 architecture exploitation (native code not req'd)		✓	✓	✓
Fast path CCW support	✓	✓	✓	✓
Productivity				
Temporary Disks (TDISK)	✓	✓	✓	✓
Resource simulation and virtualization	✓	✓	✓	✓
Complex environment testing without duplicating real hardware	✓	✓	✓	✓
Virtual processor support for SMP testing	✓	✓	✓	✓
Resource sharing (DASD, printers, I/O, memory, processors, etc.)	✓	✓	✓	✓
PVMG support for VSE VTAM sessions		✓	✓	✓
Device-independent I/O support		✓	✓	✓
Guest Coupling Facility	✓		✓	✓
RAS				
S/390 error recovery (processors, I/O, etc.)		✓	✓	✓
Dynamic multi-image support for scalability, hot backup, etc.	✓	✓	✓	✓
S/390 trace and debug facilities	✓	✓	✓	✓
Operations				
Virtual machine controls, scheduling, and automation	✓	✓	✓	✓
Performance measurement, reporting, and management	✓	✓	✓	✓
Rapid creation and deletion of guest images	✓	✓	✓	✓
IPL from a VM saved segment		✓	✓	✓
Dynamic I/O reconfiguration, CP Sparing, FICON	✓	✓	✓	✓



© Copyright
IBM 2000

Linux under VM High-speed Networking

IBM




*Turbo G6 projected

© Copyright
IBM 2000

Linux under VM Pure Linux Networking

IBM




*Turbo G6 projected

© Copyright
IBM 2000

Linux under VM Higher-speed Networking

IBM



*Turbo G6 projected

© Copyright
IBM 2000

Linux under VM Performance

IBM

- VM introduces overhead
 - ▶ So does LPAR
 - ▶ Interpretive Execution Facility eliminated traditional sources of VM overhead
 - ▶ Preferred guests run with near-native performance
 - ▶ Non-preferred guests may run with better-than-native performance
- Still, "How much overhead?" is an interesting (and inevitable) question


© Copyright
IBM 2000

Linux under VM Idling Performance

IBM

- ▶ Based on G4
- ▶ Approximately 17% of cost of an idle OS/390 guest
- ▶ Overhead tied mostly to Linux timer requests, varies slightly with network configuration
- ▶ Projected maximum is 300 to 350 1-way guests per engine

Overhead for Idle Guests




© Copyright IBM 2000

Linux under VM Non-idle Overhead

IBM

- ▶ 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, TPF, Music, or OS/390
- ▶ No surprises - overhead acceptable



© Copyright IBM 2000

Linux under VM Scalability

IBM


- Linux has scalability issues today
 - ▶ Current Linux support for SMP is immature
 - ▶ Locking mechanisms are not performance-friendly
- Virtual machine environment is ideal for handling growing Linux workloads (via server replication)
 - ▶ One application per server is the norm for Linux systems
 - ▶ VM overhead is very low in a multi-Linux environment
 - ▶ VM's support for large n-way maximizes throughput of Linux guest systems
 - ▶ S/390 architecture exploitation in VM further boosts Linux performance
- Industry investment in technology for Linux farms and clustering have particular value for VM

© Copyright
IBM 2000

Linux under VM Scalability

IBM

Compile Test

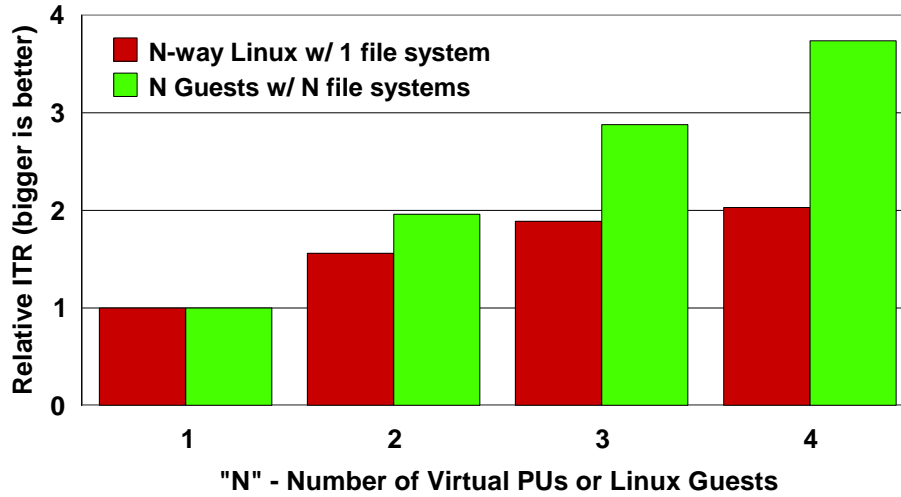


© Copyright
IBM 2000

Linux under VM Scalability



File System Lock Stress Test




© Copyright IBM 2000

Linux under VM Scalability



S/390 Workload Scalability

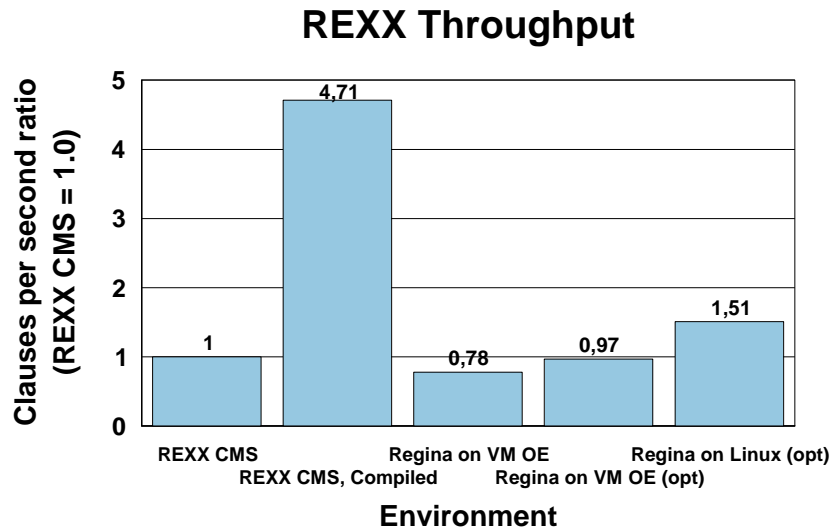


© Copyright IBM 2000

Linux under VM

REXX Performance

IBM



© Copyright
IBM 2000

Linux under VM

Productivity

IBM


- Temporary disk
 - ▶ Device-independent via minidisk driver
 - ▶ Quick and easy way to provide additional temporary space
 - To hold incoming tar files
 - To meet other short-term needs

- Resource sharing
 - ▶ Minidisks
 - ▶ Read-only LINKs
 - High-performance data sharing
 - Augmented by minidisk cache
 - Version control and central upgrade

© Copyright
IBM 2000

Linux under VM Resource Sharing

IBM



© Copyright
IBM 2000

Linux under VM RAS

IBM

- Debugging
 - ▶ Legendary VM strength
 - ▶ TRACE, DISPLAY, DUMP, STORE, VMDUMP
 - ▶ APAR VM62480
 - ASCII support for DISPLAY, DUMP, STORE, LOCATE, LOCATEVM
 - ▶ The only (sane) way to debug kernel and device driver enhancements

© Copyright
IBM 2000

Linux under VM

Operation and Management

IBM

- Virtual machine controls, scheduling, and automation
 - ▶ Shares and share limits
 - ▶ AUTOLOG and FORCE
 - ▶ Secondary Console Interface Facility (SCIF)
 - SET SECUSER
 - SEND
 - Programmable Operator (PROP)

S/390

© Copyright
IBM 2000

Linux under VM

A Real Life Example

IBM

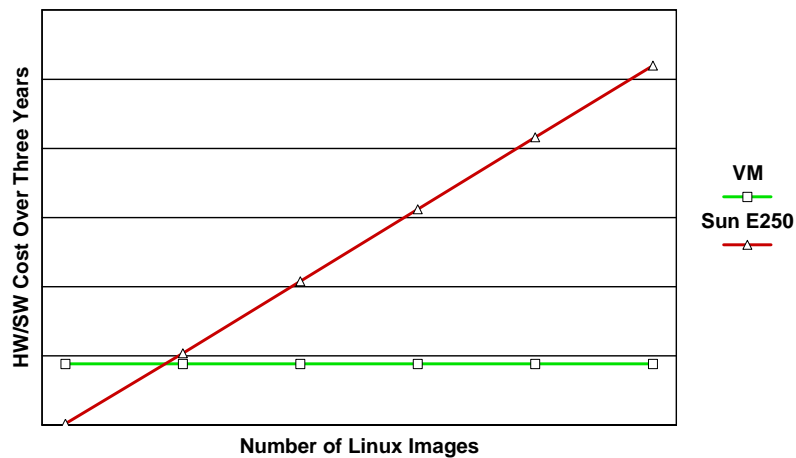
- Large telecom provider wanted to offer managed services
 - ▶ Domain Name System
 - ▶ Internet News Network
- Already running S/390
 - ▶ IBM 9672
 - ▶ VM/ESA 2.3.0 native
 - ▶ Two OS/390 guests
- New workload planned for deployment on 250 Sun E250 servers
- Instead, deployed 250 S/390 Linux guests on VM
- Linux workload
 - ▶ 200 INN servers (about 11K hits per server per day)
 - ▶ 50 DNS servers (averaging 2K hits per minute)

S/390

© Copyright
IBM 2000

Linux under VM 3-year Cost Trend

IBM



Note: 3-year lease
Multiprise 3000 + VM V2 + TCP/IP
Sun E250

© Copyright
IBM 2000

Linux under VM Possibilities

IBM

- VM provides an interesting model for server consolidation and cooperation
 - ▶ David Boyes experiment ran 41,400 Linux virtual machines in an LPAR with two engines
 - ▶ Consolidation without combination has advantages

- Linux may offer an alternative way to obtain new and more modern versions of applications on VM
 - ▶ DB2 UDB
 - ▶ Lotus Domino
 - ▶ SAP
 - ▶ Oracle
 - ▶ WebSphere
 - ▶ Java
 - ▶ ...

© Copyright
IBM 2000

Linux under VM

First In The Location Offering

IBM

- Install VM in a location where it is not already installed
- Pay no VM license charges for 6 months + 2 month test period
- Can be combined with other offerings (e.g., New Workload)
- Use to experiment with Linux for S/390

© Copyright
IBM 2000

Linux Workspace Offering

IBM

VTAM	CICS	DB2
VSE	VSE	VSE
S/390		

Example 1

VTAM	CICS	DB2	
VSE	VSE	VSE	LINUX
S/390			

Example 2

VTAM	CICS	DB2	MQ DB2 Conn Samba Apache
VSE	VSE	VSE	LINUX
S/390			

Example 3

Example 1: 3 engines running traditional workload
Business-as-usual pricing

Example 2: Add 1 engine to example 1 dedicated to Linux for S/390 with commercial Linux and open source software

Same traditional workload as in Example 1 – same BAU pricing – additional engine dedicated to Linux for S/390 does not affect existing IBM software charges

Additional costs

Hardware – additional engine dedicated to Linux
Software – cost of Linux distribution

Example 3: Same configuration as example 2 with IBM middleware for Linux for S/390 added

Same traditional workload as in example 1 – same BAU pricing – additional engine dedicated to Linux for S/390 does not affect existing IBM software charges

Additional costs

Hardware – additional engine dedicated to Linux
Software – cost of Linux distribution
Cost of IBM middleware for Linux for S/390 – DB2 Connect and MQSeries
IBM middleware for Linux for S/390 priced based on UNIX pricing model

© Copyright
IBM 2000

Linux for S/390 Extended Environment

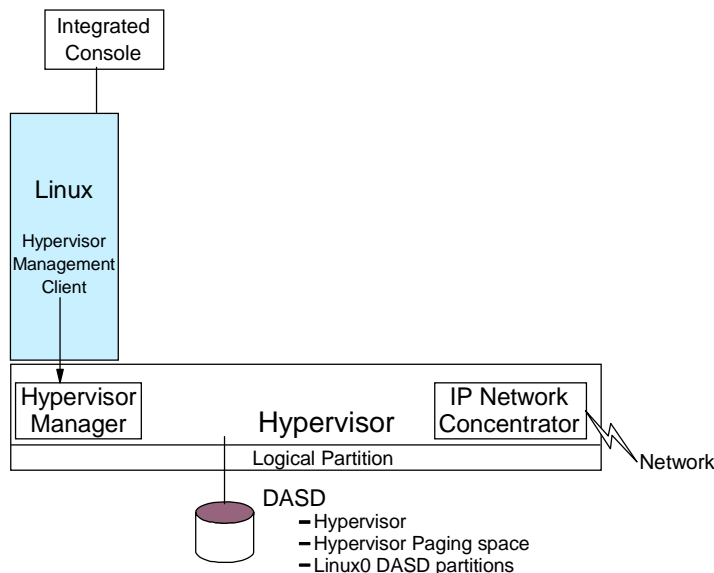
IBM

- Part of Workspace Offering special bid
- Simplified hypervisor layer to provide multiple Linux images
 - ▶ Native or under LPAR
 - ▶ Less function than VM
 - No image debugging facilities
 - No support of other operating systems
 - Controlled exposure of hypervisor interface
 - ▶ Attractive to customers without VM skills
- Linux-based administration client

© Copyright
IBM 2000

Linux for S/390 Extended Environment

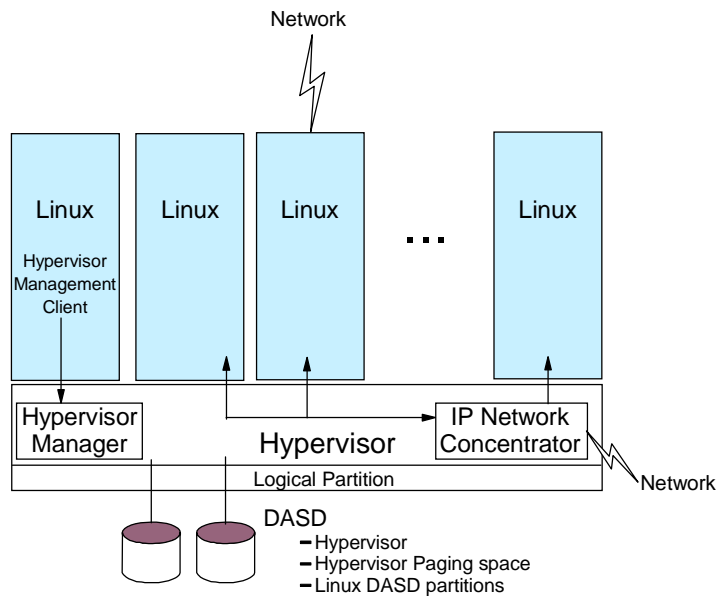
IBM



© Copyright
IBM 2000

Linux for S/390 Customized Extended Environment

IBM



© Copyright
IBM 2000

Observations

IBM

- Linux for S/390 changes the landscape
 - ▶ Software portability
 - ▶ Application environments
- VM technology adds value to Linux for S/390
 - ▶ Performance
 - ▶ Productivity
 - ▶ RAS
 - ▶ Operation and Management

© Copyright
IBM 2000

References

IBM

- "Linux for S/390"
 - ▶ IBM Developer Works
 - ▶ <http://oss.software.ibm.com/developerworks/opensource/linux390/>
- "LS-J003 Linux on S/390"
 - ▶ IBM ITSO
 - ▶ <http://itsopok.itso.ibm.com/internsh.nsf/IntNumber/LS-J003>
- "Linux on S/390 Project"
 - ▶ Marist College
 - ▶ <http://vm.marist.edu/linuxvm/>
- "Linux/390 at Princeton University"
 - ▶ Princeton University
 - ▶ <http://penguinvm.princeton.edu/>

© Copyright
IBM 2000

References *(continued)*

IBM

- "Linux for S/390"
 - ▶ Thinking Objects Software
 - ▶ <http://linux.s390.org/>
- "Listserv lists for www.marist.edu"
 - ▶ Marist College
 - ▶ <http://www.marist.edu/htbin/wlvlists/>
- "S/390: The Linux Dream Machine"
 - ▶ Linux Today
 - ▶ <http://www.linuxplanet.com/linuxplanet/reports/1532/1/>
- "VM/ESA and Linux Resources"
 - ▶ IBM
 - ▶ <http://www.vm.ibm.com/linux/>

© Copyright
IBM 2000

References (continued)

IBM

- "An Introduction to Linux"
 - Neale Ferguson
 - <http://pucc.princeton.edu/~neale/index.htm>

- "Java, threads, and scheduling in Linux"
 - IBM
 - <http://www-4.ibm.com/software/developer/library/java2/index.html>

- "Linux on the System/390"
 - Adam Thornton
 - <http://www.linux390.com/>

- "Linux for Big Iron"
 - User community Linux/VM Project
 - <http://www.linuxvm.org/>

© Copyright
IBM 2000

References (continued)

IBM

- "Linux on Big Iron"
 - open source IT
 - http://www.opensourceit.com/tools/000327_mainframe.html

- "IBM Unveils Way to Make Mainframes Act Like Hundreds of Smaller Servers"
 - Bloomberg (Technology News)
 - <http://www.bloomberg.com/bbn/technology.html>

- "IBM kicks off Linux developer program for its S/390 mainframes"
 - InfoWorld
 - <http://www.infoworld.com/articles/pi/xml/00/03/31/000331piibmlinux.xml>

- "Has the Linux 'dream machine' arrived?"
 - ZDNet News
 - <http://www.zdnet.com/zdnn/stories/news/0,4586,2491800,00.html>

© Copyright
IBM 2000

References (continued)

IBM

- "IBM Goes Countercultural With Linux"
 - ▶ NY Times
 - ▶ <http://www.nytimes.com/library/tech/00/03/biztech/articles/20soft.html>

- "Linux livin' large on mainframe"
 - ▶ Network World
 - ▶ <http://www.nwfusion.com/news/2000/0403linuxibm.html>

- "Linux vs Microsoft: Who solves security problems faster"
 - ▶ Security Portal Weekly Newsletter
 - ▶ <http://www.securityportal.com/cover/coverstory20000117.html>

© Copyright
IBM 2000