Managing Linux Under z/VM

Phil Smith III
Linuxcare, Inc.
SHARE 99
August 2002
Session 9323

Copyright Information

Agenda

- z/VM annoyances for Linux administrators
- Linux annoyances for z/VM system programmers
- Aspirin & Tylenol for Linux admins and VM sysprogs
- The solution…

Background

Who Is Linuxcare?

- Founded in 1998 to provide enterprise Linux support
- Deep technical expertise in Linux
  - Recently added System/390 expertise as well
- Many custom products delivered: IBM, HP, Sun, et al.
- Also provide certification and service: IBM, Dell, etc.
  - Multi-distribution developer support
  - Configuration support for PC Linux customers
  - Security audits
  - etc…
- 2002: Adding products to the mix!

Linux In the Enterprise

- Linux offers an excellent business case with good ROI
- Linux on System/390 (zSeries) with VM is even better
  - Server consolidation=***huge*** savings for large shops
- Linux overall is a somewhat immature technology
  - Needs care & feeding
- Offers administrative challenges:
  - Network configuration
  - Integration with existing process frameworks
  - Account management, change management, etc.
- Biggest challenge: "culture clash" between mainframe and distributed staff
  - Besides "turf wars", terminology etc. are different!

Linux Affects All Platform Teams

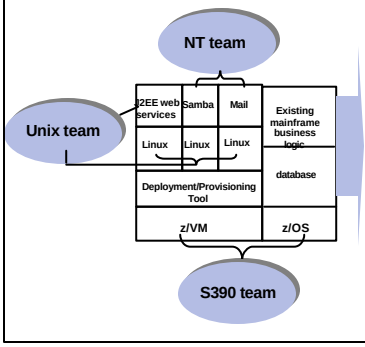

Success is:

1) Realizing the promise of server consolidation

Success requires:

1) Executive sponsorship
2) Mainframe team providing managed self-service to distributed teams

z/VM and Linux Annoyances

z/VM Annoyances for Linux Folks

Itches for the seasoned Linux administrator:

- Strange new IBM terminology
  - DASD == hard drive
  - Core == storage == memory == RAM
  - Storage <> disk space!
  - User == user *or* Linux image?
  - IPL == boot
  - IML == BIOS (more or less)
- Oddly ordered IBM documentation
  - SLSS
  - Bookshelf hard to find on the Web (well, all IBM pages can be hard to find on the Web…)

z/VM Annoyances for Linux Folks

- OCO drivers
  - Means must rely on IBM to fix problems
  - Alien to Linux theology and practice
- Gaining VM expertise is difficult
  - Not very many VM HOWTO documents out there
  - VM Primer manual no longer published
  - Little VM training available
  - No "VM For Dummies" (yet!)
  - Friendly, helpful VM community, however!
  - VMESA-L can be a lifesaver

z/VM Annoyances for Linux Folks

- Hardware is strange and different
  - Boy, that's a big tape drive!
    - And a big tape…yet it only holds ***how*** much?
  - 3215/3270 is very alien
    - Block mode
    - OK, PF keys are Function keys, but PA keys??
  - Wow, a whole laptop as a system console (HMC)!
- Brand new editors (XEDIT)
  - RECFM F, RECFM V, serial numbers…
  - Coupled with 3270 strangeness, ***very*** confusing
- All a part of learning the z/VM theology!

Linux Annoyances for VMers

Itches for the grizzled VM sysprog:

- Case sensitivity
  - This is a surprisingly hard one to learn!
- Vowel shortage
  - VM commands are English; Linux commands are Hrd2Rd
- New and different UNIX terminology
  - "Mount" not "ACCESS"
  - What do you mean, "It's in another file system"?
  - "How do I specify record format?"

Linux Annoyances for VMers

- ASCII
  - "Why can't they use EBCDIC like everyone else??"
- File system fragility
  - Possible data loss after uncontrolled shutdown even after fsck (with default filesystem)
  - Hard to believe in a technology > 30 years old!
- Strange editors…none of which work on 3270s!
  - (Ok, ed and ex…but they're evil)
- "What do you mean, ***anyone*** *can read the source?"*
  - VMers believe in source code—but not for end users
- HELP isn't help
  - HELP is man

Long-Term z/VM Headaches

- DASD management
  - Each new guest is a new install—seems inherently wasteful
- Deploying Linux instances takes time…
  - Resource allocation & z/VM user creation
  - Moving data from .iso to CD-ROM to tape to…
- Networking with z/VM's TCPIP virtual machine
  - CTC, IUCV, Guest LANs – all have issues
  - OBEYFILE has had a lot of problems
- Tuning z/VM & Linux for optimal app performance
  - Poorly documented Linux tuning APIs
  - No existing VM tools for "watching" Linux guests

Long-Term Linux Headaches

- Linux is a poor, yet greedy, guest:
  - Doesn't run in a DCSS
  - Wants to hog the CPU
  - Needs gobs of DASD to be comfortable
  - Likes to have lots of storage
    - And therefore, you end up with lots of paging…
- Hard to manage the configuration of Linux VMs
  - Many customers give up at ~20 production systems
- Difficult, if not impossible, to integrate with ESMs
  - Evolving LDAP support in ESMs may ease this
  - Also general problem of Linux user administration

z/VM Aspirin and Linux Tylenol

Some Quick VM Solutions

- DASD Management
  - Use DISKMAP if not using VM:Secure/DirMaint!
  - Talk to your elder bears
  - Encapsulate DDR sequence in DISKCOPY EXEC
- Route to a Linux instance, instead of the z/VM stack
  - Allows custom firewalling
  - Faster network configuration, with no need to bring down the TCPIP virtual machine
- Networking: Apply recent PTFs to z/VM 4.2
  - Fixed Guest LAN HiperSockets emulation
  - Fixed many OBEYFILE issues

Apply Required VM Service

- TCP/IP PTFs:
  - UQ61461 Guest LAN fixes
- CP PTFs:
  - UM30225 Guest LAN fixes
  - UM30230 Hard CP loop when short on real CHPIDs
  - UM30290 HiperSockets/OSAExpress QDIO input queue stall
- Still open:
  - VM63210 Second DEFINE NIC fails after UM30230

Some Quick Linux Solutions

- Linux is a bad guest
  - Use the "notimer" patch—it helps a lot!
  - Use VDISK!
  - Consolidate DASD (see next slide)
- Watch paging carefully, especially if running 2nd level:
  - Disable z/VM minidisk cache if short on page space
  - Three levels of paging is Badness – smaller virtual storage may be ***better***
- Configuration management
  - Check out the Linux community – many people have solutions, one may be right for your site
- Linux can integrate well with LDAP, NIS, NIS+ ESMs

DASD Consolidation

- Means using shared data to save space
- Common issue for Linux on S/390
  - Minimal Linux install at least 700 3390 cylinders
- Consolidation candidates include anything which is read-only and identical across multiple Linux instances
  - Theoretically all except /home, /var, /etc, and /tmp
  - Can be complicated, using RAMdisks, etc.
- Consolidation can save lots of DASD quickly, but…
  - Makes production upgrades more complex
- Implement using NFS or shared read-only DASD
  - Both have advantages/disadvantages

The Solution…

Linuxcare Relieves the Pain…

Levanta – A Linux image manager for VM

- Completely new product
- Goes beyond "cloner" functions of other products
- Created ***by*** Linux and VM people, ***for*** Linux and VM people!
  - Native VM and native Linux components
- Available 3Q2002

Linuxcare Relieves the Pain…

Levanta Version 1.0

- Rapid Linux instance creation and cloning: z/VM virtual machines, Linux, applications
- Configuration change management (apply, rollback)
- Three administration interfaces
  - Web, Linux, CMS (command line and full-screen)
- ESM integration (VM:Secure, DirMaint)
- DASD sharing and consolidation
  - Read-only Linux binaries placed in shared DASD

Levanta Functions

- Manage instance groups with an intuitive tool
  - Provision/deployment time reduced to 2 minutes
  - Deploy updates across servers simultaneously
  - Cycle instances remotely
- Enable change & configuration management
  - Change management
  - Rollback to prior stable version
  - Configuration templates
  - Support can have complete system copy
- Interact with multiple interfaces
  - Unix, z/VM, or Web interface

Levanta Functions

- Template capabilities
  - Template and change log define the Linux instance
  - Change log can be "harvested" into a template
  - Untracked instance changes can be "refreshed" into the change log
- Grouping instances
  - Startup/shutdown Linux instance virtual machine
  - Add/remove/modify packages
  - Schedule start-up of instances
- Execute functions inside an instance (individually controllable by Levanta administrator)
- Install via tape or ftp
- File server redundancy

Levanta User Interface

- Individual user interfaces
  - Three admin skill/permission levels
  - Operate against single abstraction layer
- Granular functionality by user type
  - Linux functions for Linux administrators
  - VM functions for VM administrators
- Internationalization fully enabled
  - American English at GA, other languages to follow
- Use case-driven design and development
- Web interface portable and high-performance
  - All HTML, no Java etc.
- Full change log search capabilities

Levanta Architecture

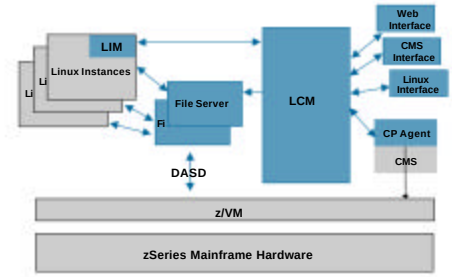

Levanta Benefits

- Minimize capital & operating expenses
  - Fewer physical boxes to maintain
  - Avoid new staff by reducing VM skills required
  - Add new servers without adding admin headaches
  - Each admin can manage more servers
  - Grow servers without growing headcount
- Provide managed self-service IT
  - Distributed teams can provision virtual servers while maintaining centralized discipline and control
  - Server templates, access controls, and change logs help avoid mistakes while preserving flexibility
- Approach zero-defect IT
  - Higher availability, more robust backup & DR

Levanta Futures

Possible areas for enhancements:

- Console management
  - Both active (operational) and spool files
- Integration with management monitoring frameworks
  - Tivoli Enterprise Console interface
  - Tivoli Partner??
  - BMC Patrol
- Performance management
  - Quick "What are my Linuxes doing?"
  - Control if (when!) one is being a hog
- Clustered instances on multiple systems (failover)

Summary

Conclusions

- Moving to Linux on System/390 causes pain for both Linux and z/VM folks
  - New platform
  - New challenges to both communities
- None of the pain is insurmountable
- Work with your fellow admins & sysprogs
  - Subscribe to LINUX-390@marist.edu and VMESA-L@listserv.uark.edu
  - Scan the list archives
- Linuxcare and Levanta can help!
  - "Operators are standing by…"

Questions?

Contact Info

Phil Smith III

LINUXCARE

703.568.6662
psmith@linuxcare .com
www.linuxcare.com