Simplifying Linux Under z/VM

Phil Smith III
Linuxcare, Inc.

Agenda

- z/VM hurdles for Linux administrators
- Linux obstacles for z/VM system programmers
- Aspirin & Tylenol for Linux and VM folks
- A solution…

Background

Who Is Linuxcare?

- Founded in 1998 to provide enterprise Linux support
- Deep technical expertise in Linux
  - Recently added z/VM mainframe expertise as well
- Custom technologies 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 enterprise products to the mix!

Linux In the Data Center

- Linux offers excellent business case with good ROI
- Linux on zSeries with z/VM can be even better
  - Linux on z/VM == huge savings for large shops
- Linux overall is a somewhat immature technology
  - Needs care & feeding
- z/VM–Linux combination offers administrative challenges:
  - Network configuration
  - Integration with existing process frameworks
  - Account management, change management, etc.

Linux In the Data Center

More challenges:

- Relatively few, public success examples
  - Winnebago, Boscov's, Korean Air, et al.
  - But do they apply to ***your*** business?
  - Many more known to exist, but not public
- Difficulty generating cost saving projections
  - "Too good to be true" numbers easy to produce

Linux In the Data Center

- Confusion about
  - Applicability
  - Pricing
  - Vendor commitment and direction
- Which applications?
  - DNS, firewall
  - File & print, mail
  - Static web serving, dynamic web serving

Linux In the Data Center

- Lack of enterprise-quality tools
- Skills availability
  - "We don't know z/VM"
  - "We don't know Linux"
- Uniqueness/risk of Linux on z/VM
- Fear and emotion
  - "Server consolidation" = "**layoffs**"?
- Biggest challenge: "culture clash" between mainframe and distributed staff
  - Besides "turf wars", even terminology is different!

Linux In the Data Center

"That's not a data center…"

"Now, ***that's*** a data center!"
— Crocodile Dundee

Linux In the Data Center

- Linux on z/VM requires collaboration and “buy-in” across multiple teams
- “Simple” Samba deployment involves:
  - NT team: Application source & client support
  - Unix team: Application porting, ownership, support
  - Networking: HiperSockets–Ethernet integration
  - Mainframe: z/VM platform support
- Any Linux on z/VM deployment requires 2, 3, and 4

Best Practices: Linux with Data Center RAS

- Tremendous value found in data center processes
  - Disaster recovery/auditability
- Cross-team synergy
  - Managed self-service
  - Better integration of front end/back end systems
- Cost reduction
  - TCO analysis

Linux Affects All Platform Teams

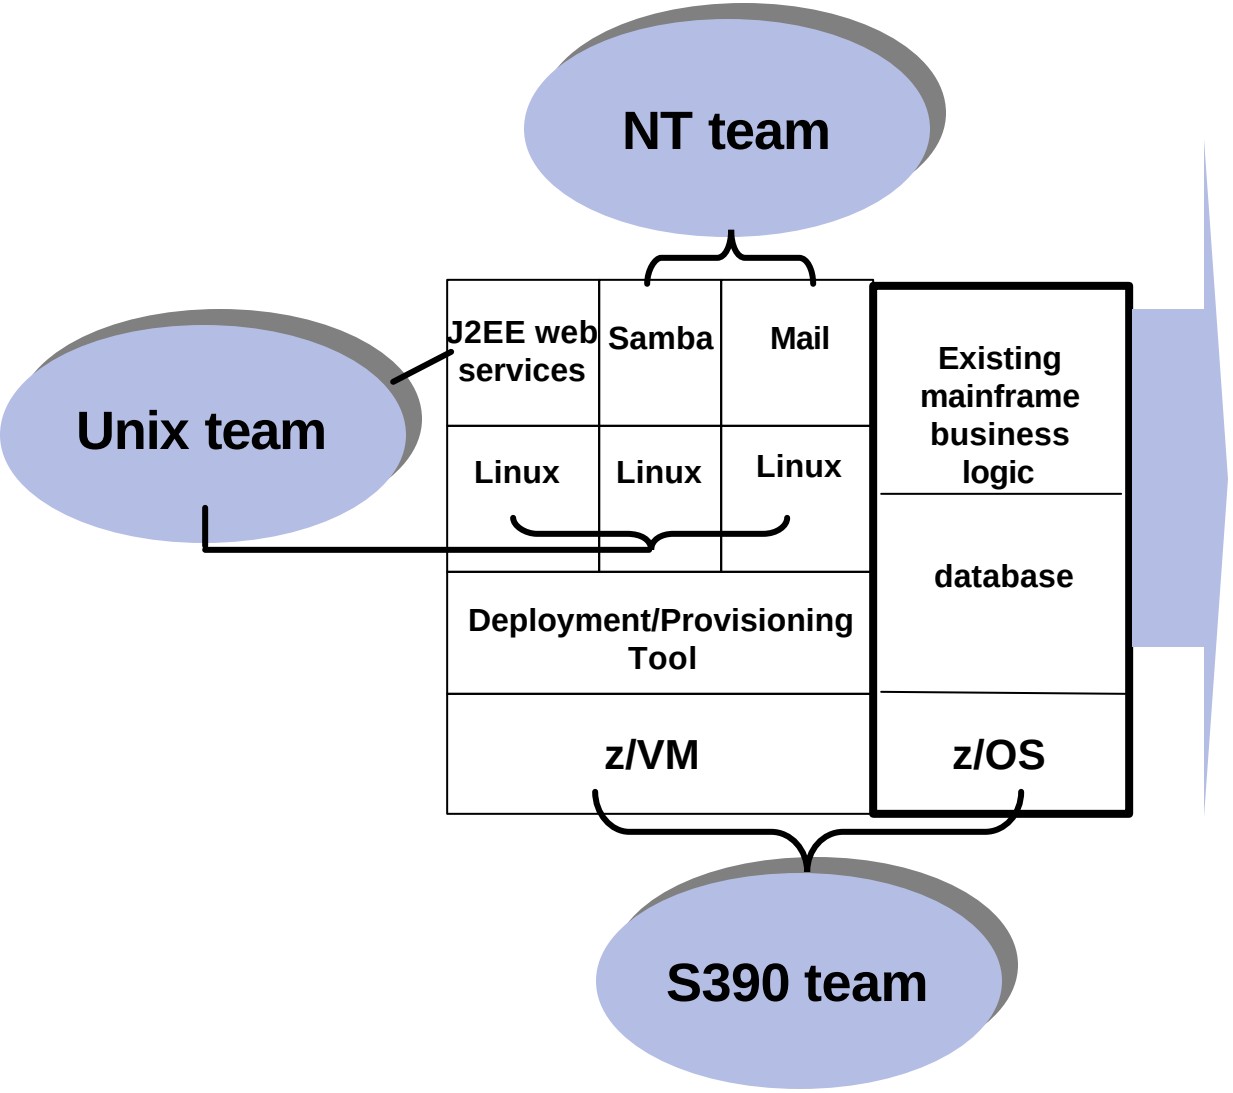

Success is:

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 Hurdles for Linux Folks

Itches for the seasoned Linux admin:

- 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 Hurdles 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 Hurdles 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 Obstacles 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 Obstacles 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; maybe NED)
- “What do you mean, ***anyone*** *can read the source?”*
  - VMers believe in source—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—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 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 (which means lots of paging)
- Hard to manage the configuration of Linux VMs
  - Many folks give up at ~20 production systems
- Difficult to integrate with ESMs
  - Can integrate well with LDAP, NIS, NIS+ ESMs
  - 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
  - UM30495/6 Second DEFINE NIC fails after UM30230

Some Quick Linux Solutions

- Linux is a poorly behaved guest
  - Assumes it owns the entire physical system
  - Uses all storage, lots of CPU, etc.
- Some techniques to minimize the effects:
  - Use the “notimer” patch—it helps a lot!
  - Use VDISK for swap!
  - Consolidate DASD (see next slide)
- Watch paging carefully, especially 2nd level:
  - Disable z/VM minidisk cache if page space short
  - Three levels of paging is Badness – smaller virtual storage may be better

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

DASD Consolidation

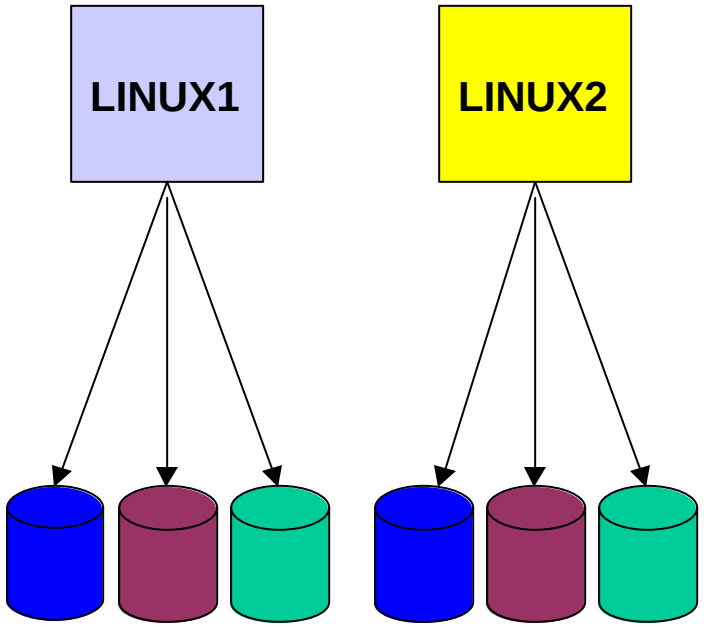

Unshared DASD:
Private copies of
all directories

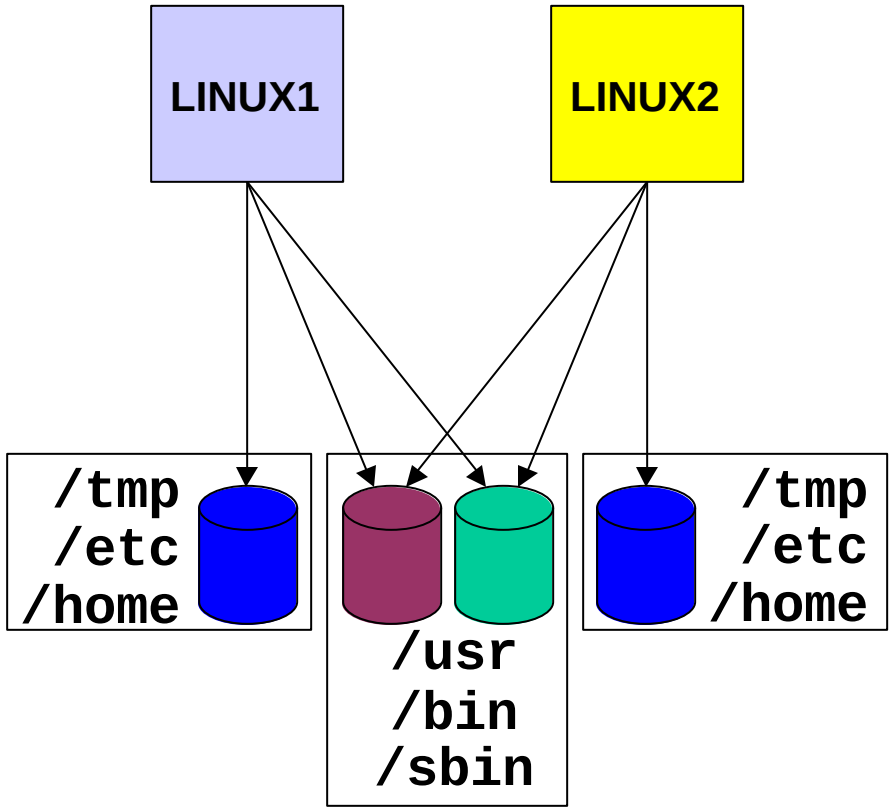

Shared DASD:
Single copy of
common, R/O data

A Solution…

Linuxcare Relieves the Pain…

Levanta – A Linux instance manager for VM

- Completely new product
- Far beyond “cloner” functions of other products
- Created ***by*** Linux and VM people, ***for*** Linux and VM people!
  - Native VM and native Linux components

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; full-screen soon)
- ESM integration (VM:Secure, DirMaint)
  - Maintains existing infrastructure, operations
- Automated DASD sharing and consolidation
  - Read-only Linux binaries placed in shared DASD
  - File server redundancy avoids single point of failure

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
- Execute functions inside an instance (individually controllable by Levanta administrator)

Levanta Functions

- Instance groups
  - Startup/shutdown Linux instances en masse
  - Add/remove/modify packages
  - Schedule start-up of instances
- Package pools
  - Logical definition of a set of Linux programs
  - Can be manipulated as a single object
  - Aid in defining/controlling instances

Levanta Functions

- Templates
  - Essential concept in understanding Levanta
  - “Cookie cutters” to create Linux instances
  - Template + change log define Linux configuration
  - Change log can be “harvested” into new template
  - Untracked instance changes can be “refreshed” into change log
  - Defined in XML documents: editable, sharable

Levanta User Interface

- Users can use interface of choice
  - Unix, z/VM, and Web UIs
  - Improves collaboration, cooperation
  - Operate against single abstraction layer
- Granular functionality by user type
  - Three administrator skill/permission levels
  - Linux functions for Linux administrators
  - VM functions for VM administrators

Levanta User Interface

- 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
- Command line clients mean all functions scriptable

Levanta User Interface

Levanta User Interface

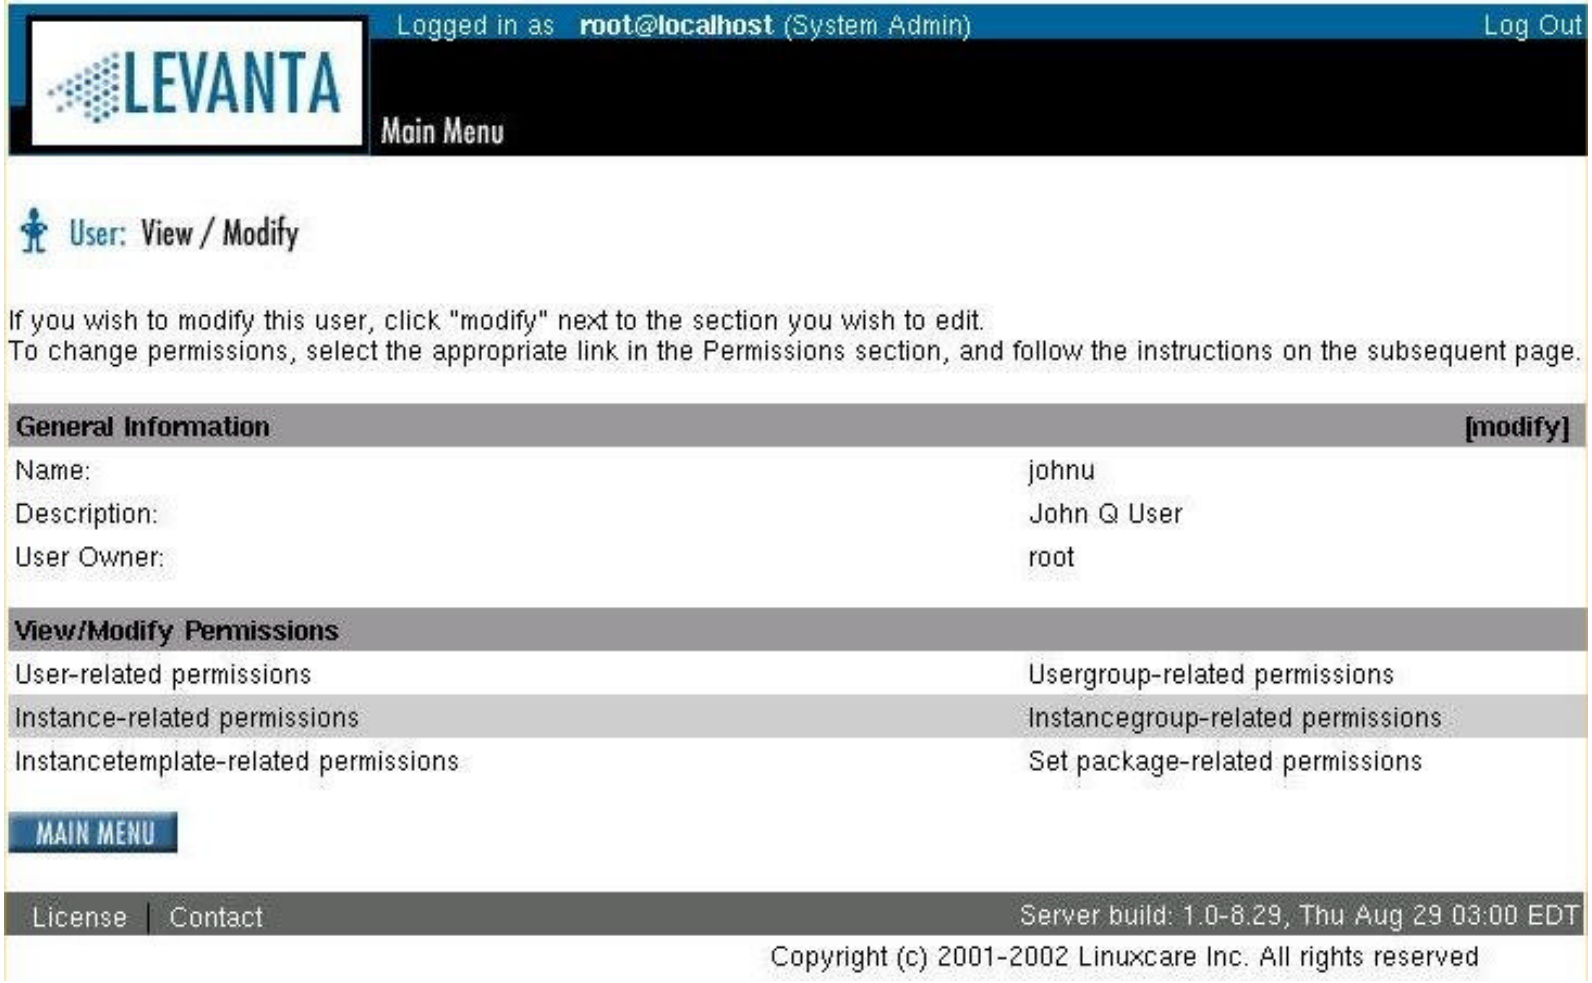

Levanta Architecture

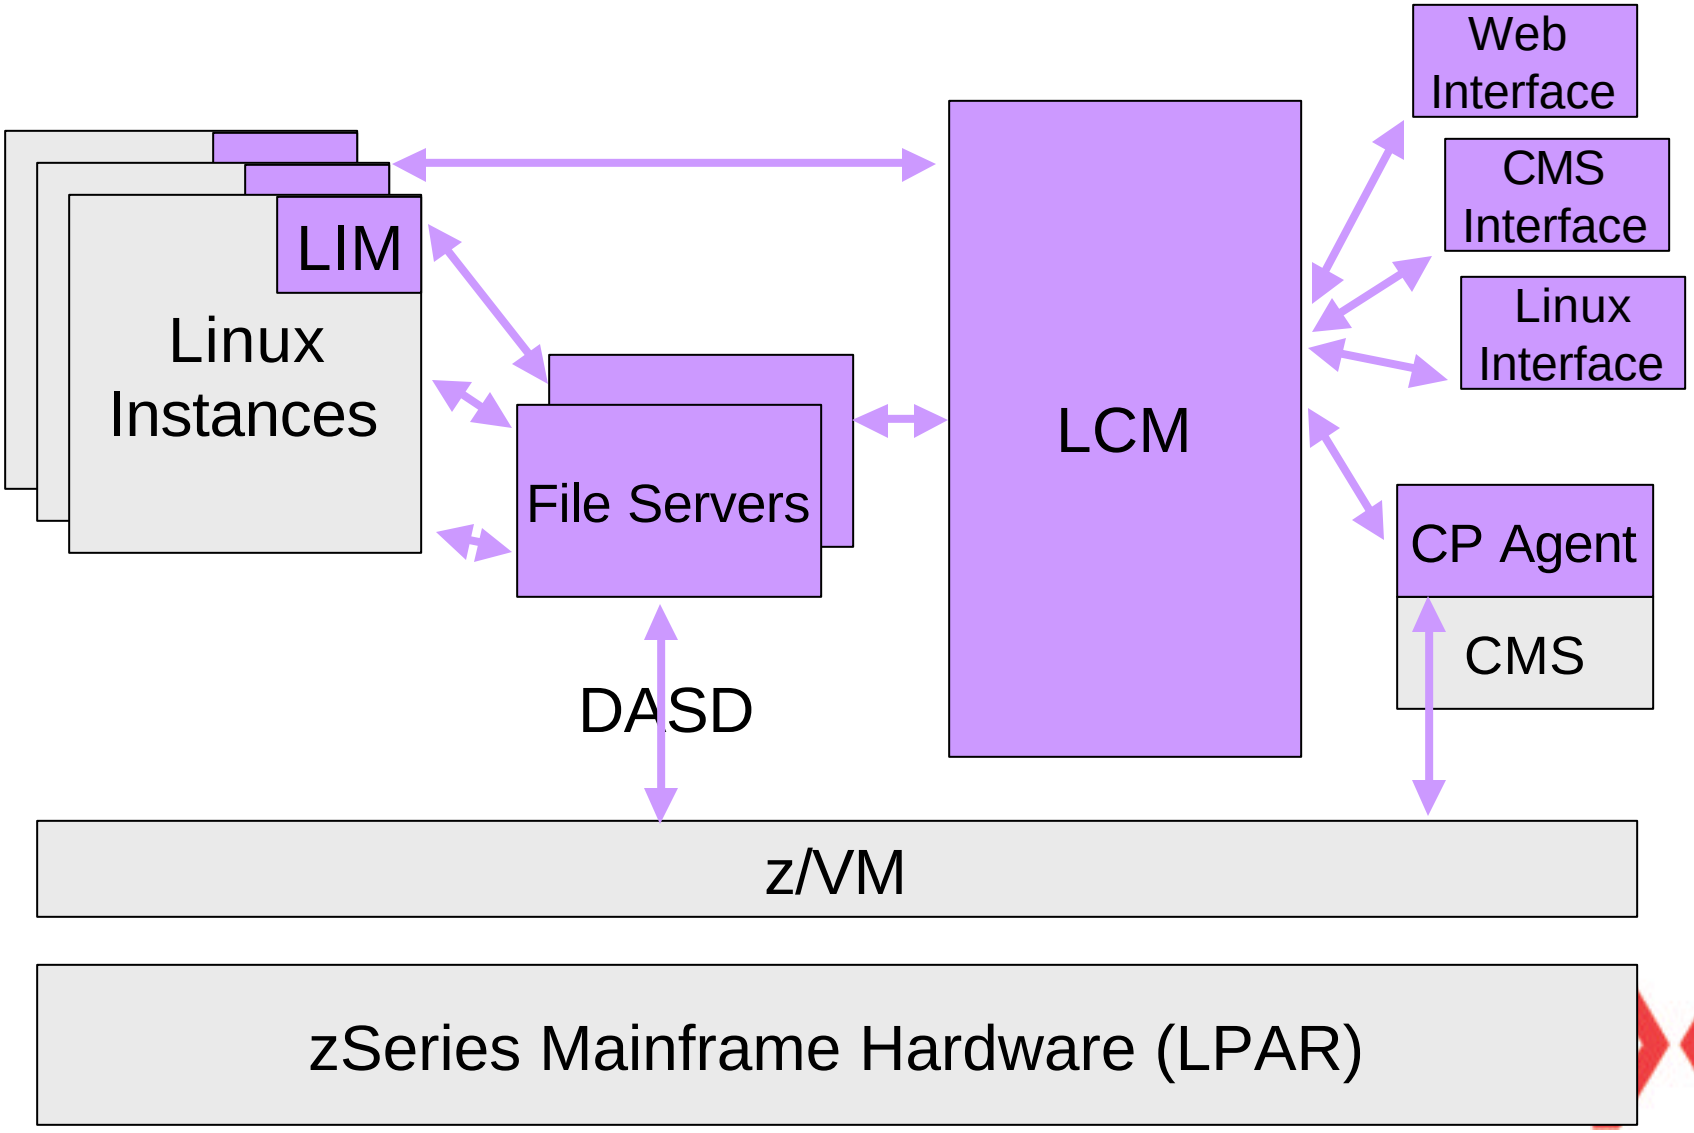

Levanta DASD Consolidation

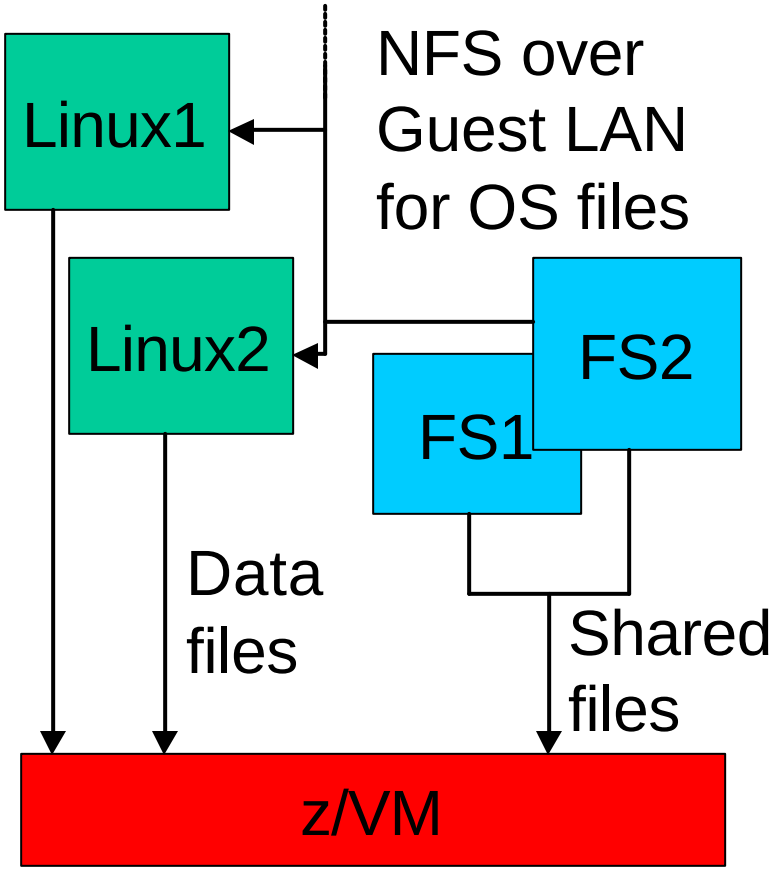

Remaps files: Specific files/copies accessible only by owning instances

Is redundant: Can failover to ensure instances stay up

Bypasses shared R/O minidisk limitations:

- Need not reIPL to update filesystems
- Enables fast file deployment
- Each file exists *once* no matter how often used
- Far more manageable as usage grows

Private vs. Consolidated DASD

- Situation: 20 web servers, each with 20MB of data and configuration files, 300MB of operating system files

	Individual	Consolidated
Data	20*20MB = 400MB	20*20MB = 400MB
OS files	20*300MB = 6GB	1*300MB = 300MB
Total	**6.4GB**	**700MB**

Levanta Benefits: Force Multiplier

- **Provides “force multiplier” capability: Can double system administrators’ effectiveness**
  - Create Linux instances quickly and accurately
  - Create Linux server templates to leverage scarce skills and speed provisioning
  - Meet urgent requests for new servers by creating instances from templates in minutes
  - Group instances to leverage administration across several instances
  - Maximize administrative operational expense savings

Levanta Benefits: Best Practices

- **Enables data center best practices with Linux: Alleviates deployment and operational issues**
  - Improve change management processes with change logging/rollback
  - Codify best practices templates and groups
  - Control server configurations throughout solution development cycle
  - Standardize software deployment via package pools
  - Tailor access through user permissions
  - Integrate with existing management frameworks

Levanta Benefits: Collaboration

- **Team collaboration shortens pilot process**
  - Encourage team collaboration among Linux, mainframe, NT, networking, and other experts
  - Enable mainframe team to provide “managed self-service” access
  - Reduce cultural barriers via multiple, familiar interfaces
  - Provide safe, reasonable flexibility with granular access permissions

Levanta Timeline and Futures

- **Limited availability:** Now
- **General availability:** October 15, 2002

Possible enhancement areas:

- Console management
  - Both active (operational) and spool files
- Integration with system management frameworks
  - Tivoli Enterprise Console interface, BMC Patrol…
- Performance management
  - Quick “What are my Linuxes doing?”
  - Control if (when!) one is being a hog

Summary

Conclusions

- Moving to Linux on z/VM can be very beneficial, even fun, but also causes pain for both Linux and z/VM folks
  - The combination is a new platform
  - New challenges to both communities
- The climb is worth the view, especially with some planning and tools
- 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?

Phil Smith III

LINUXCARE

703.568.6662
psmith@linuxcare.com
www.linuxcare.com


SERVERS RUNNING LINUX