
# Simplifying Linux Under z/VM

Phil Smith III
Levanta, Inc.

SHARE 103
August 2004
Session 9214





# Copyright Information

SHARE Inc. is hereby granted a non-exclusive license to copy, reproduce or republish this presentation in whole or in part for SHARE activities only, and the further right to permit others to copy, reproduce, or republish this presentation in whole or in part, so long as such permission is consistent with SHARE's By-laws, Canons of Conduct, and other directives of the SHARE Board of Directors

# Agenda


- Linux opportunities and challenges
- Hurdles and obstacles for Linux & VM folks
- A solution: Levanta
- The future…

# Who Is Levanta?


- Founded as Linuxcare in 1998 to provide Linux support
- Deep technical expertise
  - Linux on varied platforms (x86, zSeries, et al.)
  - z/VM and mainframe technology
- Delivered custom technologies, certification, service…
  - IBM, Dell, HP, Sun, et al.
- 2001: Shifted to enterprise software company
- 2004: Changed name to Levanta, Inc.!

# Linux In the Enterprise

- Linux offers excellent business case with good ROI
  - (Almost) zero acquisition/growth costs
  - Lower facilities costs (hardware, etc.)
- Linux on zSeries with VM can be even better
  - Linux + VM = Huge savings!
- Relatively few, public success examples
  - Winnebago, Boscov's, Korean Air, et al.
  - But do they apply to your business?
  - More exist, but are not public

# Linux Challenges

- Applicability
  - Which applications??
- Pricing
  - Is Linux free? How much will it save us?
- Vendor commitment and direction
  - Vendor X does (not?) claim to support it…
- Administration
  - User accounts, network configuration, change management, backups, etc.
- General FUD (SCO lawsuit, MS publicity…)

# Challenge: Applicability

- Which applications?
  - Answer has evolved rapidly
  - Original answer (two years ago): “infrastructure” (DNS, file & print, mail)
  - Evolved to static, then dynamic web serving
  - Now applications of all types
  - Even better value proposition…but makes it harder to pick a starting point!
- Mail, Apache, Samba, WebSphere, DB2 Connect, Oracle, BEA WebLogic are common uses

# Challenge: Cost

- So is Linux free or not?
  - “Just download it”
  - “But you have to pay for support”
  - “I bought Red Hat at CompUSA…”
- How much will it save us?
  - Difficult to generate firm cost savings projections
  - “Too good to be true” numbers easy to produce (even if sometimes accurate!)

# Challenge: Vendor Commitment

- Other mission-critical vendor applications may not openly support Linux
  - Or may support Linux but not on zSeries
  - Pressure vendors for Linux support, or for zSeries support if already on Intel
  - Consider alternatives: cost savings may justify upheaval

# Challenge: Administration


- Network configuration not “the usual”:
  - No (few) wires to plug
  - VM TCP/IP configuration different from Linux (though VM IFCONFIG and IPWIZARD help a lot!)
- Proprietary “hooks” for systems management suites not there (yet?)
- Like any system, housekeeping and maintenance require effort, new tools

# Challenge: FUD

- Lack of enterprise-quality tools
- Skills availability
  - “We don’t know { VM | Linux }”
- Uniqueness/risk of Linux on VM
  - “Who else is doing it?”
- Fear and emotion
  - “Server consolidation” = “layoffs”?
- Biggest challenge: “culture clash” between mainframe and distributed staff
  - “Turf wars”, terminology differences!

# 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 can be a “disruptive technology”, especially to folks who have strong investments in other technologies
- Linux on z/VM requires collaboration and “buy-in” across multiple teams
- “Simple” Samba deployment involves:
  1. NT team: App source and client support
  2. Unix team: App porting, ownership, support
  3. Networking: Routing, IP assignment
  4. Mainframe: hardware platform support
- Any Linux on VM deployment requires (at least) 2, 3, 4

# Best Practices:
# Linux with Data Center RAS


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


# z/VM and Linux Annoyances





# z/VM Hurdles for Linux Folks

- Strange new IBM terminology
  - DASD = hard drive
  - Core = storage = memory = RAM
  - Storage <> disk space!
  - User = Linux user or Linux guest virtual machine?
  - IPL = boot
  - IML = BIOS (more or less)
  - MIPS = MHz ??
- 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 (recently rectified — way to go, IBM!)
  - Means must rely on IBM to fix problems
  - Alien to Linux theology and practice
- Gaining VM expertise is difficult
  - Few VM HOWTO documents out there
  - VM Primer manual no longer published (buy used!)
  - Little VM training available (coming back gradually)
  - No “VM For Dummies” (yet)
  - Friendly, helpful VM community, however!
  - VMESA-L, LINUX-390 can be lifesavers

# 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

- 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 filesystem”?
  - “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 ext2 filesystems)
  - Hard to believe in a technology > 30 years old!
- Strange editors…which mostly don't work on 3270s!
- "What do you mean, anyone can read the source?"
  - Source code is good — but for end-users?!
- HELP isn't really help
  - HELP is man

# Long-Term z/VM Headaches

- DASD management
  - Each new guest is a new install — wasteful, in time and resources
- 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
  - CTCAs, IUCV, Guest LANs — all have had issues
- Tuning z/VM & Linux for optimal app performance
  - Few existing VM tools for “watching” Linux guests (Velocity, PerfMan…)
  - Poorly documented Linux tuning APIs

# Long-Term Linux Headaches


- One or two Linux guests are easy to manage
  - Complexity grows rapidly with number of guests
  - Many folks give up at ~20 production systems
- Software administration headaches are ongoing
  - New packages must be installed
  - Installed packages must be updated
  - Backups take a ***long*** time
  - Mistakes get made, backups must be restored

# Linux Guest Performance

- Standard Linux is a greedy guest:
  - Assumes it owns the entire physical system
  - Uses all storage, lots of CPU, etc.
  - Needs large amounts of DASD to be comfortable
  - Likes to have lots of storage (which means lots of paging)
- Some techniques to minimize the effects:
  - Use the “notimer” patch — it helps a lot!
  - Use VDISK for swap

# Paging and E-lists

- Typical problem: “Things just stop”
  - CP INDICATE et al. don’t show any problems
  - But users are in Eligible lists (typically E3)
- Large virtual storage and large working sets combine:
  - CP does resource calculation
  - Realizes it will overcommit, forms E-lists instead
  - Things look OK, but not much is happening!
- Watch paging carefully, especially second level:
  - Disable z/VM minidisk cache if page space short
  - Three levels of paging is Badness — ***smaller*** virtual storage is often ***better***


# z/VM Aspirin

## and

# Linux Tylenol®





# Levanta Relieves the Pain…


- LEVANTA™ – A Linux instance manager

  Manages entire Linux server lifecycle in enterprise data centers
  - Far beyond “cloner” functions of other products
  - Provision/deployment time reduced to minutes!
- Uses “operating system” approach
  - Reframes server lifecycle as state-based problem, solved by managing state transitions
- Capability-based versus procedural management
  - Saving/switching among states vs. automating operations
- Available on z/VM since 2002; available on Intel now!

# Levanta Is the Solution


## With Levanta

- 10 minutes to provision 10 servers
- Under 10 minutes to recover to last known stable state
- 1 sysadmin for every 60 servers
- Move virtual servers to physical hardware in less than 5 minutes

## Without Levanta

- 4 hours to provision 1 server (scripted)
- 4-hour restore from tape
- 1 sysadmin per 20 – 30 servers
- Days of labor for a virtual to physical hardware migration

**Levanta manages the *entire* lifecycle of Linux servers in enterprise data centers**

# Levanta Features

- Change and configuration management
  - Checkpoint and rollback
  - Rollback files or entire instance to prior stable version
  - (Re)apply changes, to same or other Linux instance
  - Support/test can have complete system copy in minutes!
- ESM integration (VM:Secure, DirMaint)
  - Maintains existing infrastructure, operations

# Levanta Features

- Can aggregate instances into groups
  - Remotely startup/shutdown Linux instances *en masse*
  - Deploy updates across servers simultaneously
  - Add/remove/modify packages to one instance or an entire group
- Can aggregate installable software into groups
  - Logical definition of a set of Linux programs
  - Can be manipulated as a single object
  - Aid in defining/controlling instances

# Levanta Features

- Templates
  - “Cookie cutters” to create Linux instances
  - Changed instance can be “harvested” into new template
- Users can select interface of choice
  - Command line and Web UIs available
  - All functions scriptable
- Granular functionality by user type
  - Multiple administrator skill/permission levels
  - Systems functions for systems staff, application functions for application staff, etc.

# Levanta DASD Consolidation

- Automated DASD sharing and consolidation
  - Read-only Linux binaries placed in shared DASD
- Guest-specific data stored on private R/W “overlay” disk
- Bypasses shared R/O minidisk limitations:
  - Need not reIPL to update filesystems
  - Enables fast file deployment
  - Shared files exist once no matter how often used
  - Still manageable as usage grows

# The Server Lifecycle

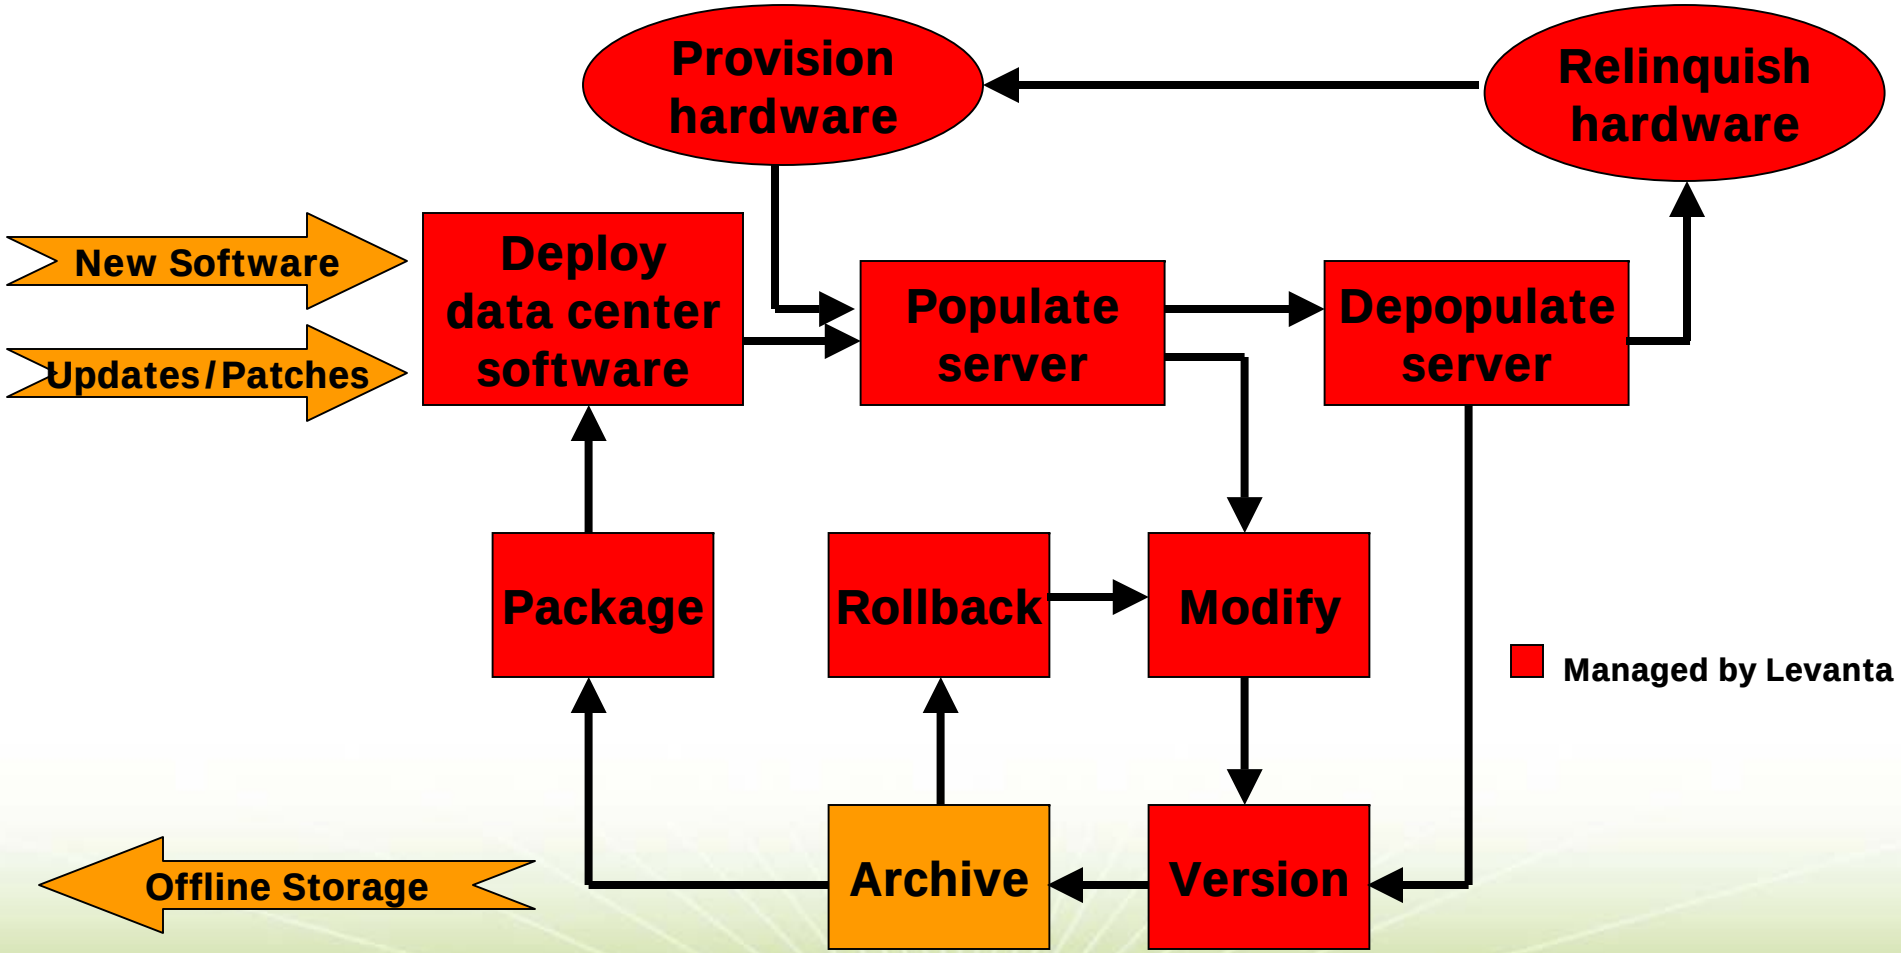

# Levanta Repositories

- A Levanta Repository is a mountable volume containing one or more installation packages
  - May contain multiple RPMs, but can be treated as a single entity
  - Repositories contains exploded files, pre- and post-install operations, and dependency information
- Repositories are read-only and managed by Levanta
  - Are thus authoritative in terms of server content
- Servers are populated from Repositories
  - Can add and remove repositories from Levanta-managed servers

# Data Center Software Deployment

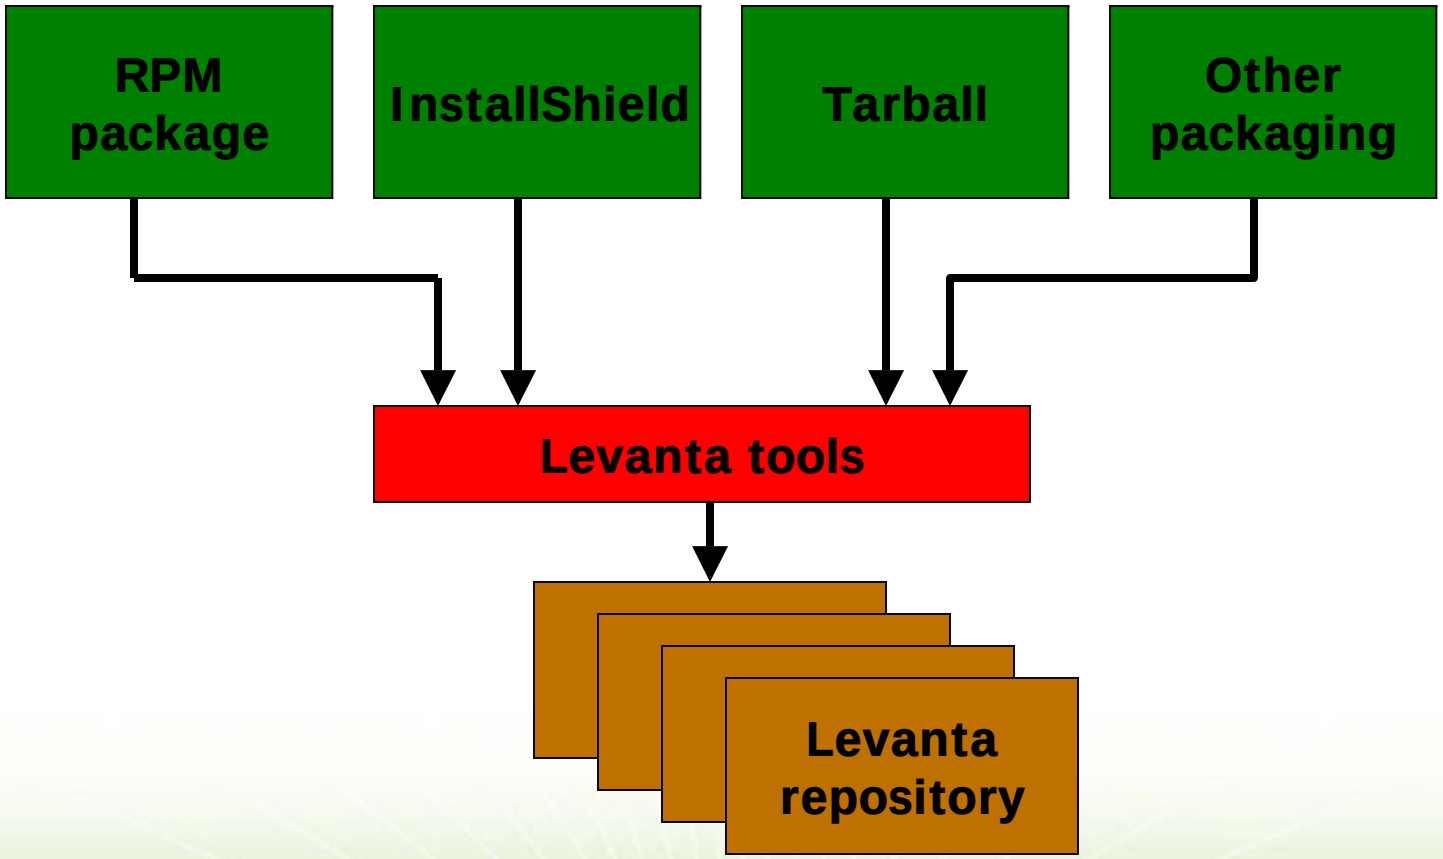

# Platform Provisioning

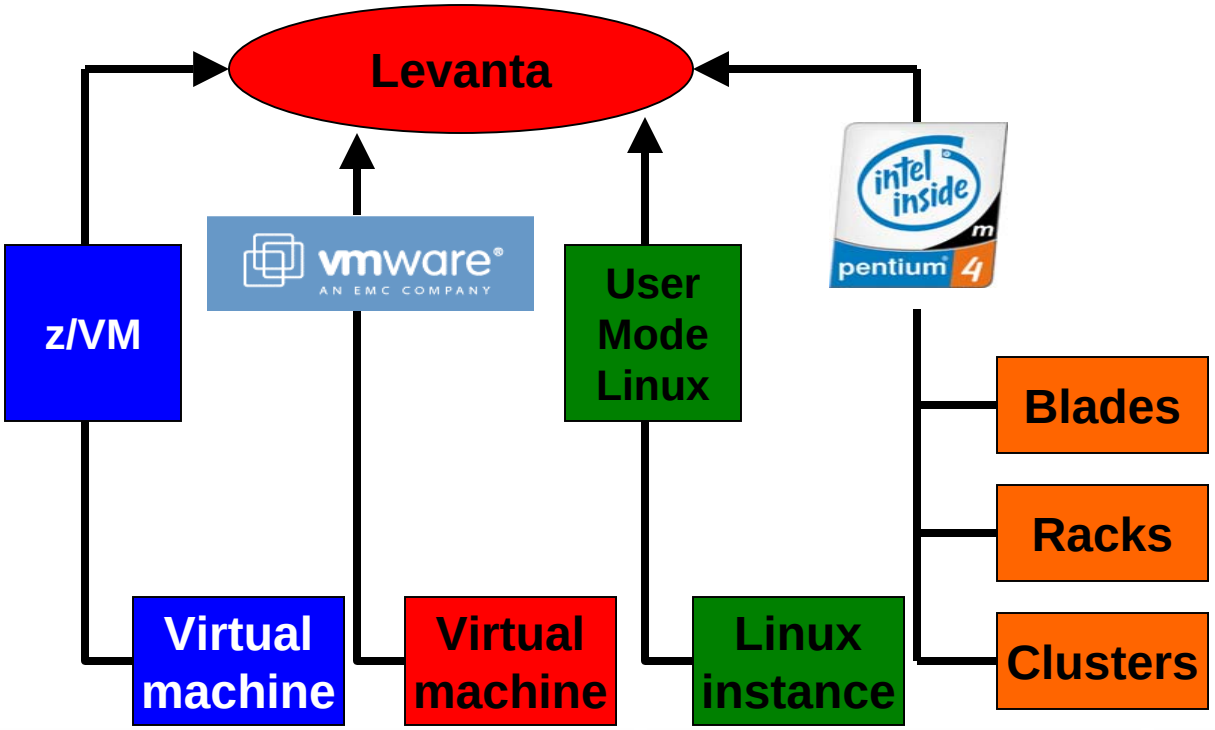


- Heterogeneous platforms
- Virtual/physical
- Best-fit provisioning from multiple pools
- Specification through templates
- Operational control: start/stop/reboot
- Develop on one hardware and deploy on another

# Populating Servers

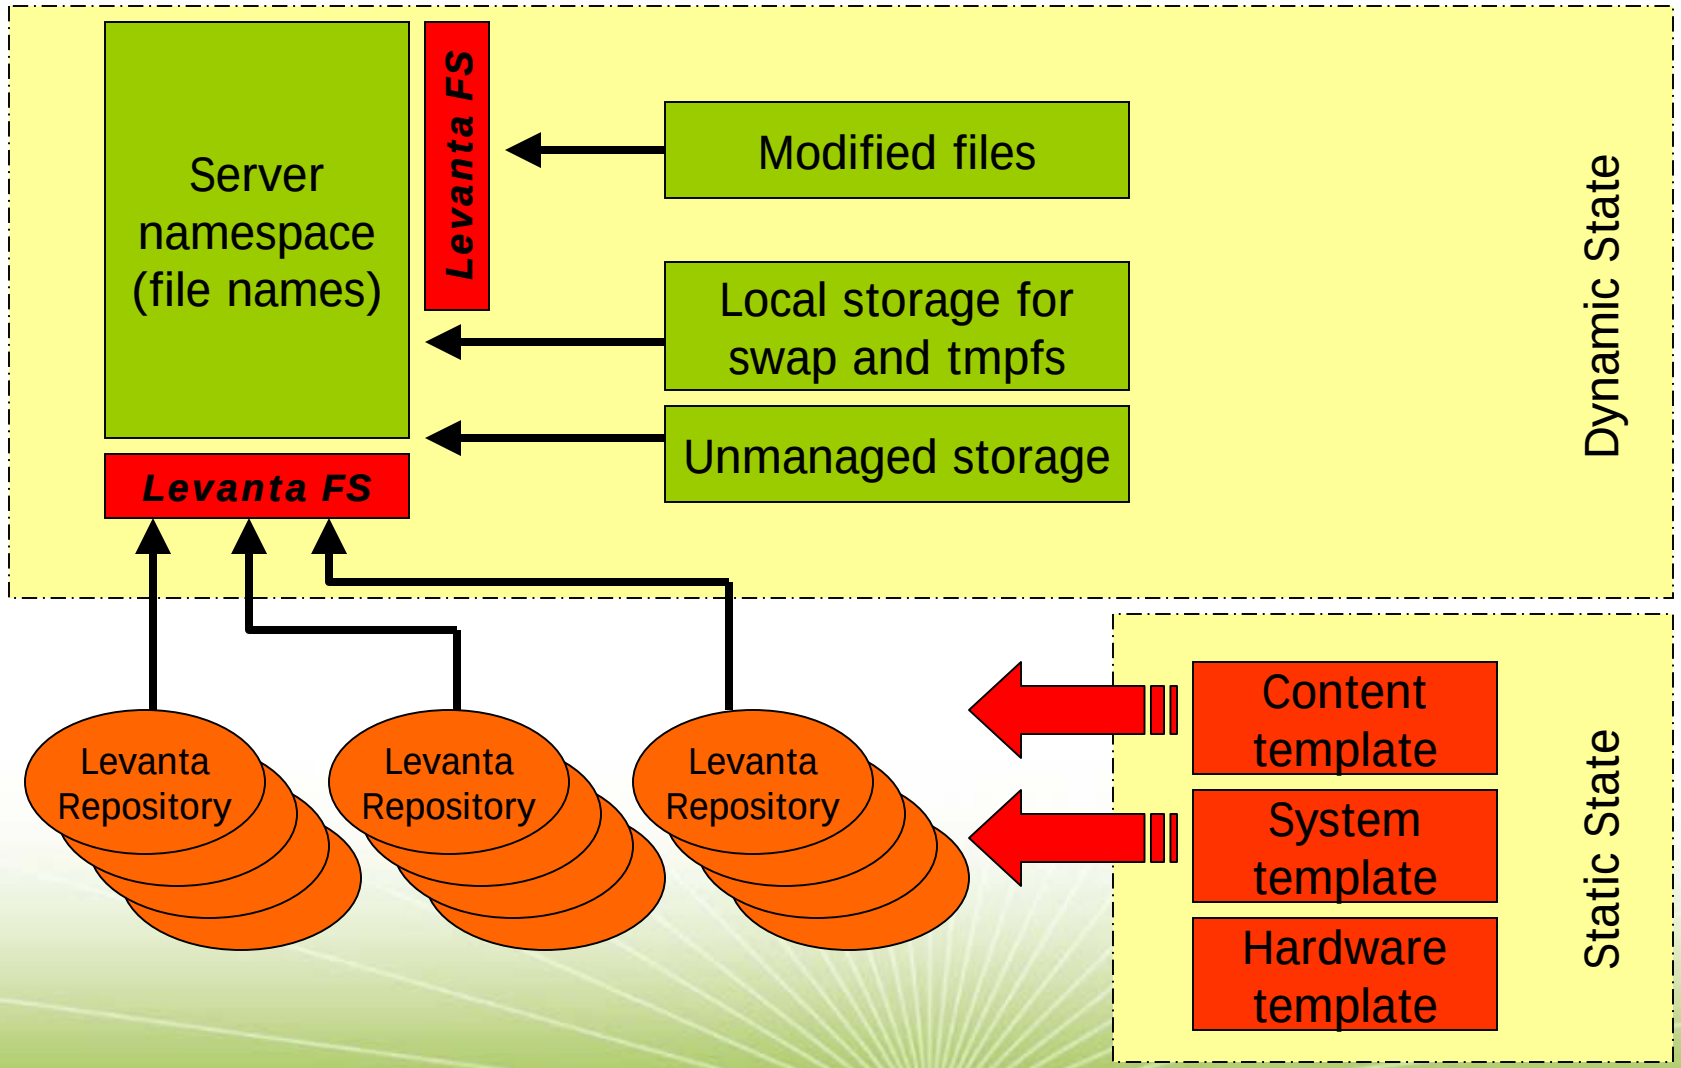

# Modifying Servers

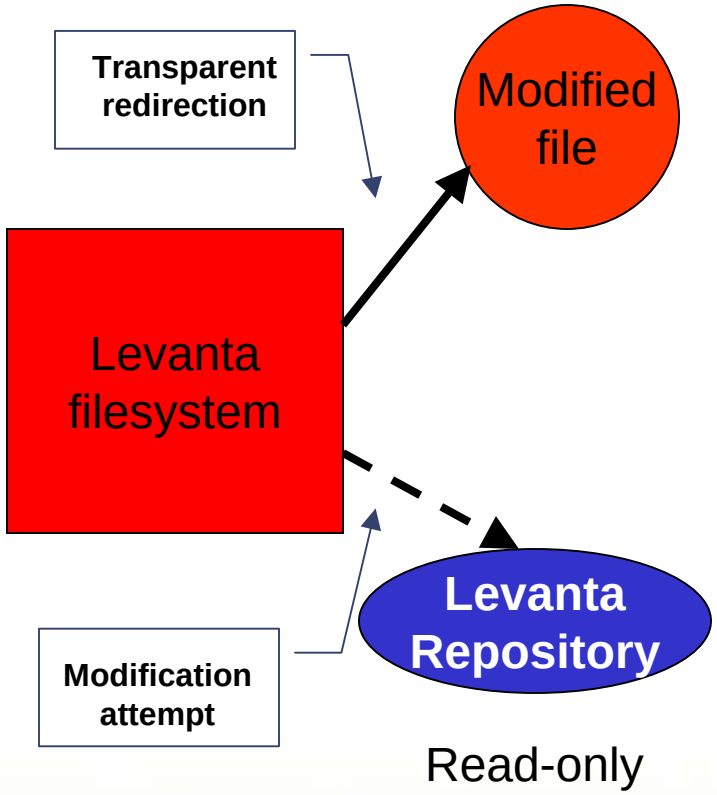


- Levanta filesystem is stacked atop existing filesystems
- Levanta filesystem enables lookup redirection
- All modifications are intercepted
- Copy-on-write creates private copy of file for server
- Subsequent references are redirected to modified file
- Rapid modified-file lookup

# Checkpoints and Rollback

- Checkpoints are incremental deltas of server content between two points in time
  - Include file and Repository changes
- Rapid checkpoint mechanism enables frequent invocations to track server file modifications
- “Undo” function for server state
- Discriminatory rollback versus
  - Install/uninstall
  - Image rollback
- Can rapidly rollback file, server, or group of servers to any checkpoint

# Archiving

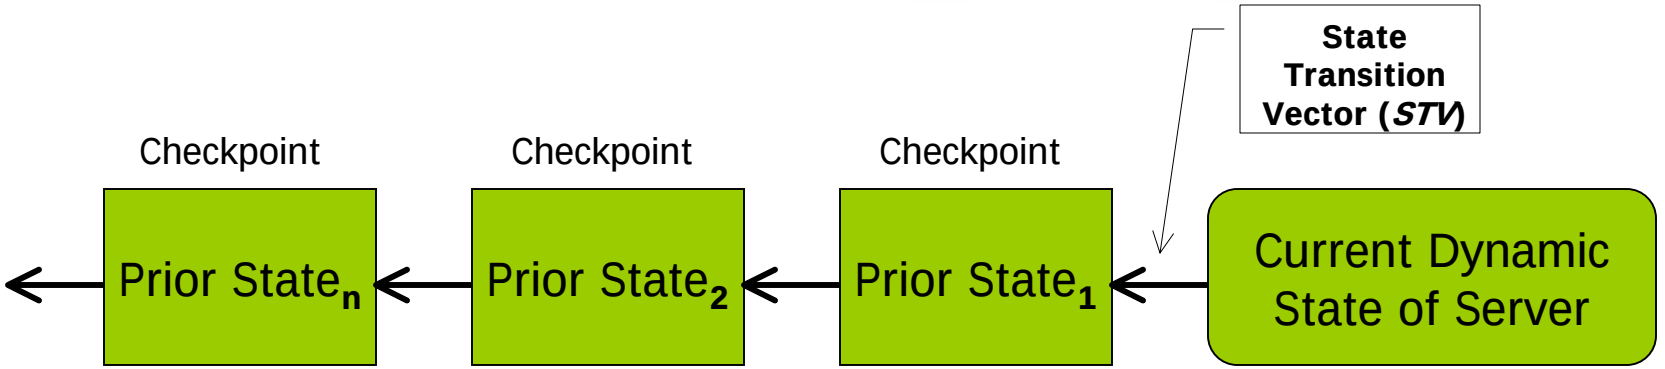


- ***State Transition Vector*:** the delta between two states
- STVs require minimal storage
  - Megabytes of changes from petabytes of data
- Static state does not change (template change)
- Rapid election of modified files — no **`diff`** required
- Automatic mirror of modified files

# Packaging Any Application

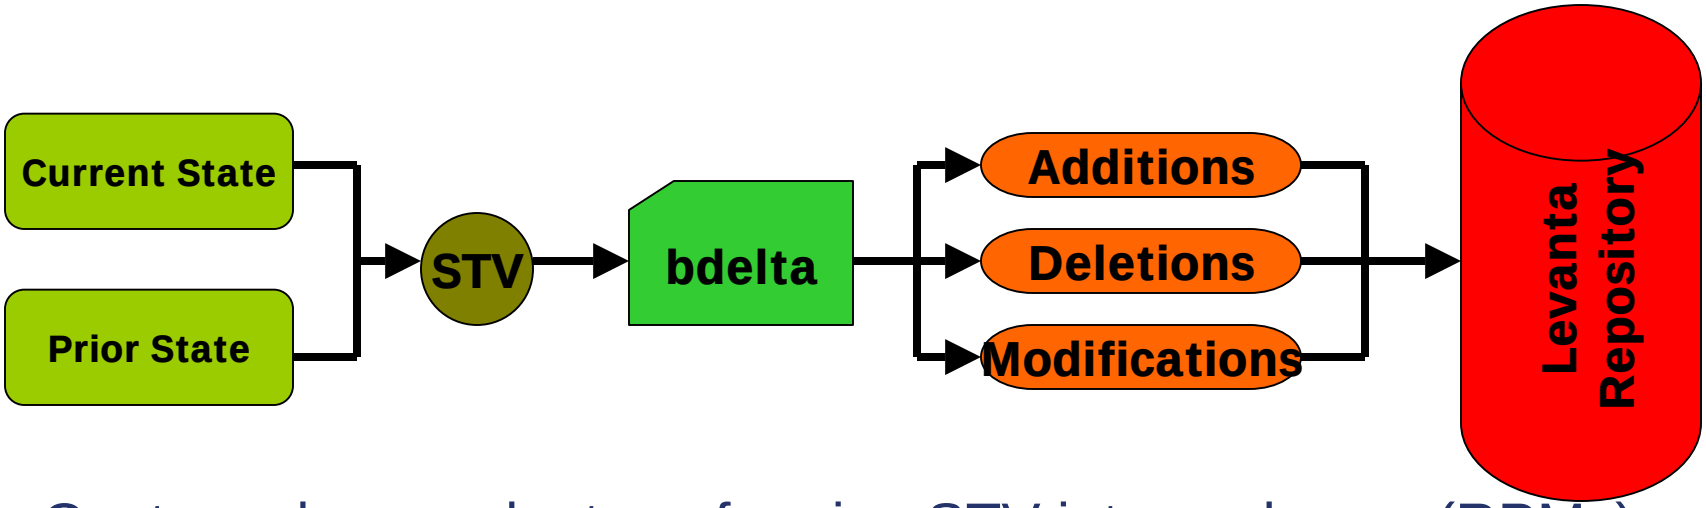


- Capture changes by transforming STV into packages (RPMs)
- Redeploy changes to one server or to hundreds of similar servers

# Depopulating Servers


- Very rapid depopulation by removal of installed Repositories
- Can remove single Repository or entire content of server
- Rollback to a previous Repository topology
- Apply to running or stopped server

# Relinquishing Hardware

- Levanta decouples hardware from software (virtualization)
- Hardware is released back to pool when either:
  - All software is depopulated from server
  - Server is destroyed voluntarily
- Hardware is then available for reuse for same or different service

# Server State Cycle

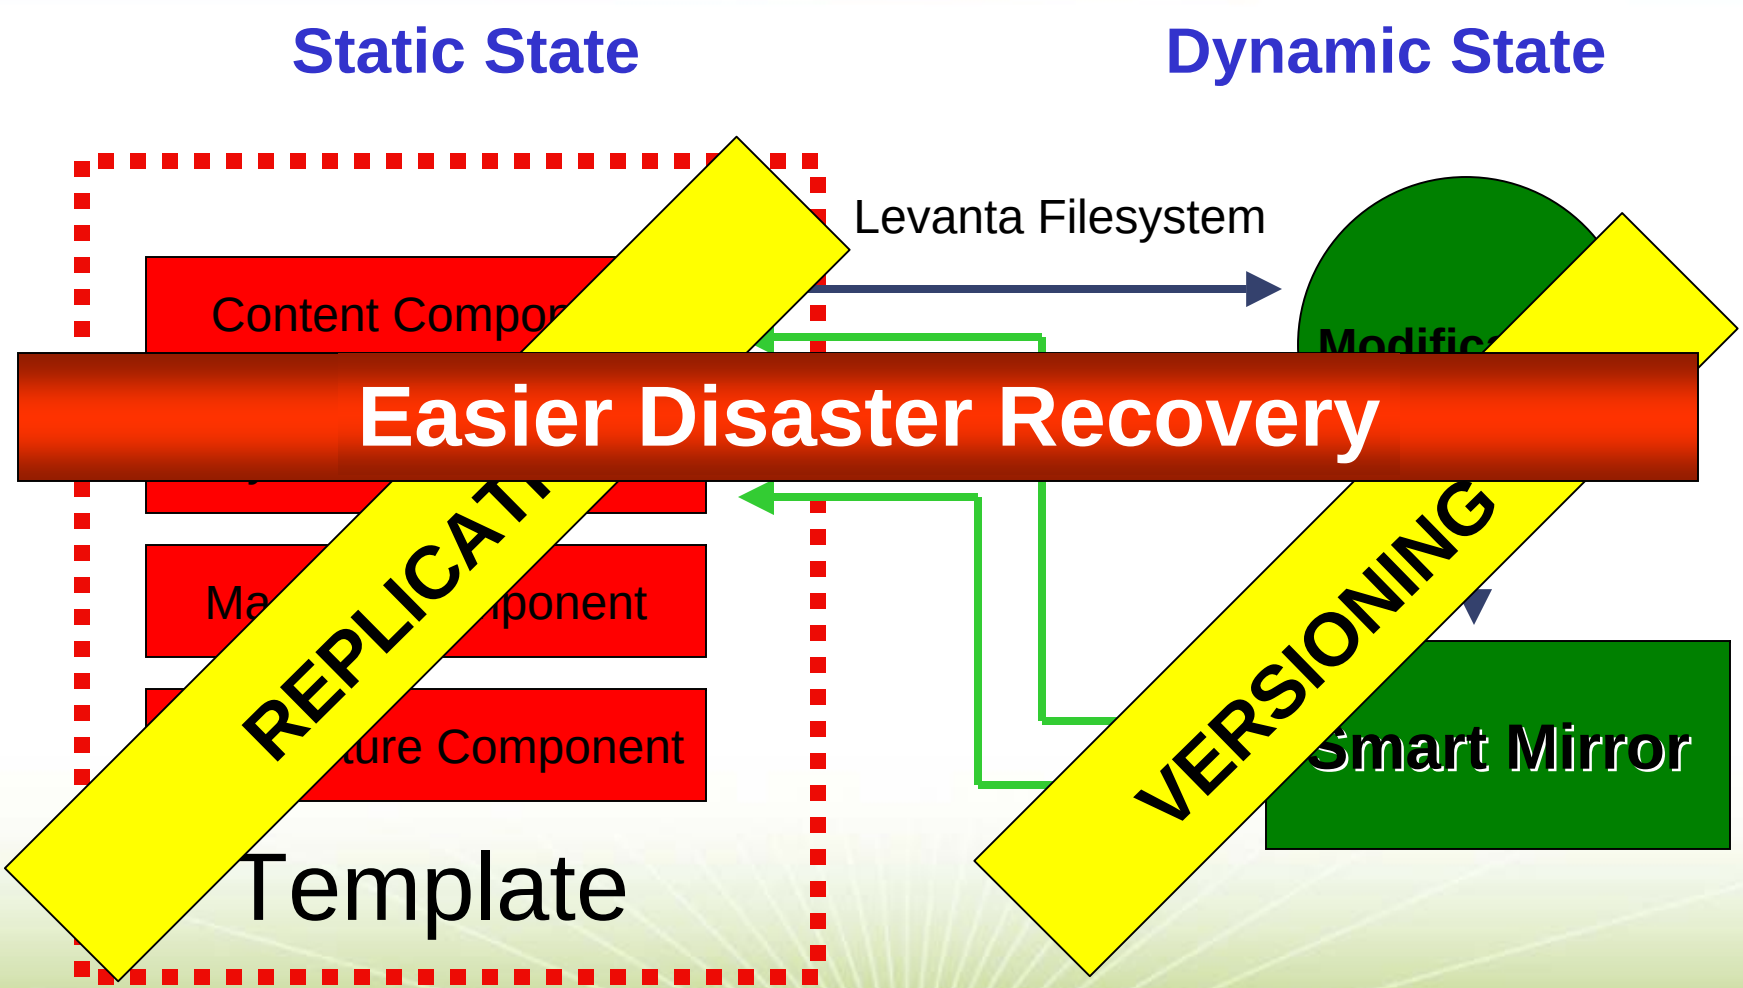

# Levanta Architecture

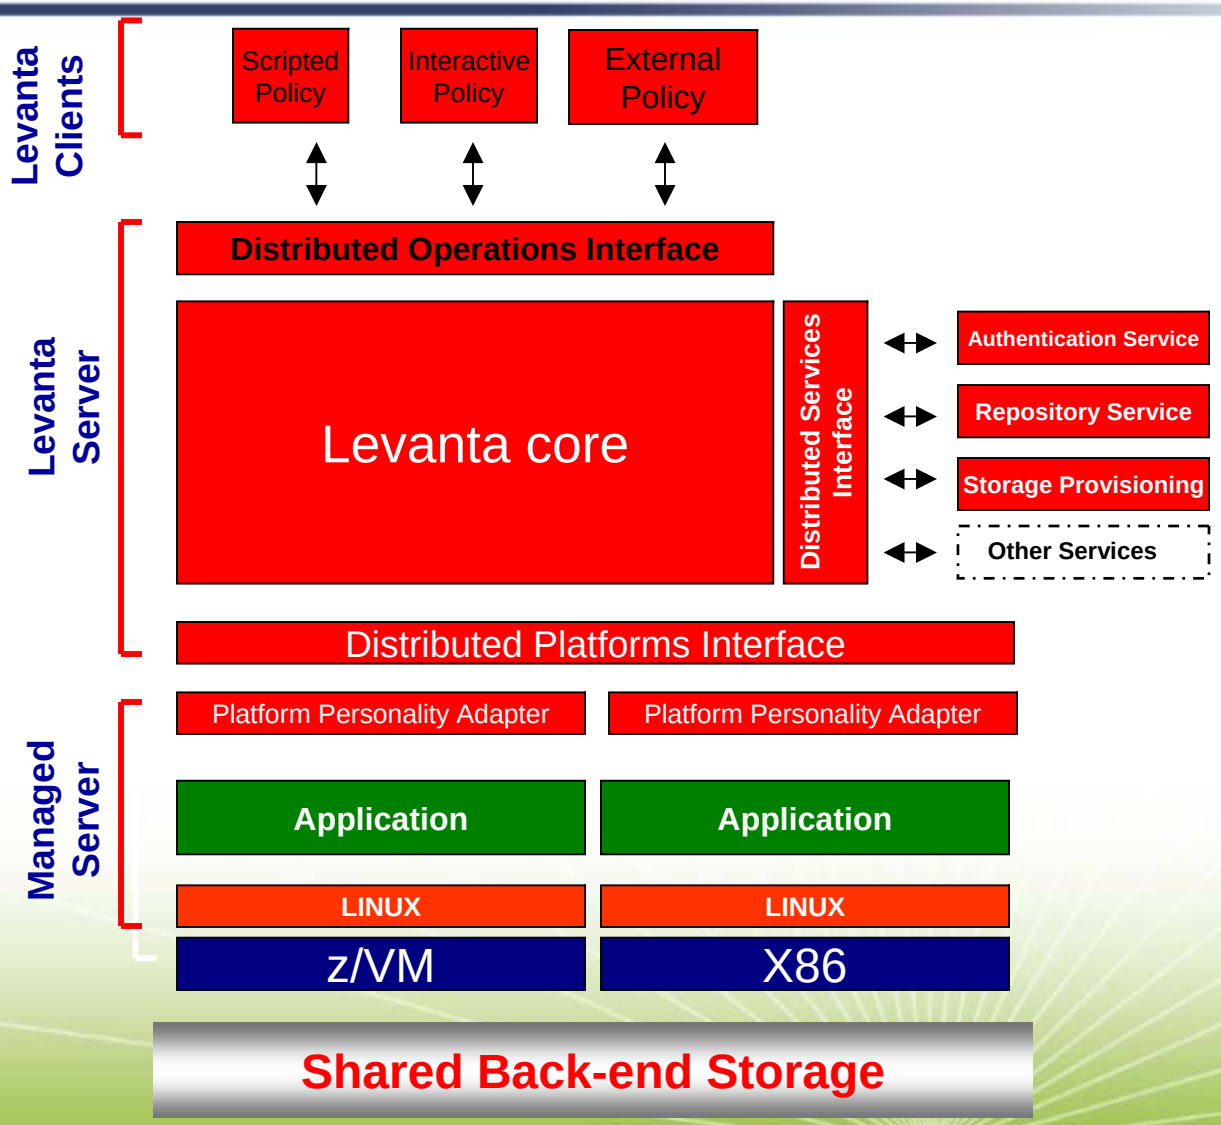


Levanta Policy Layer

System Calls (API)

Levanta Kernel and Services

Levanta DDI/DKI Platform Driver

Data center Applications

Platform Operating System

Geographically Distributed Platforms

# Levanta VM Implementation

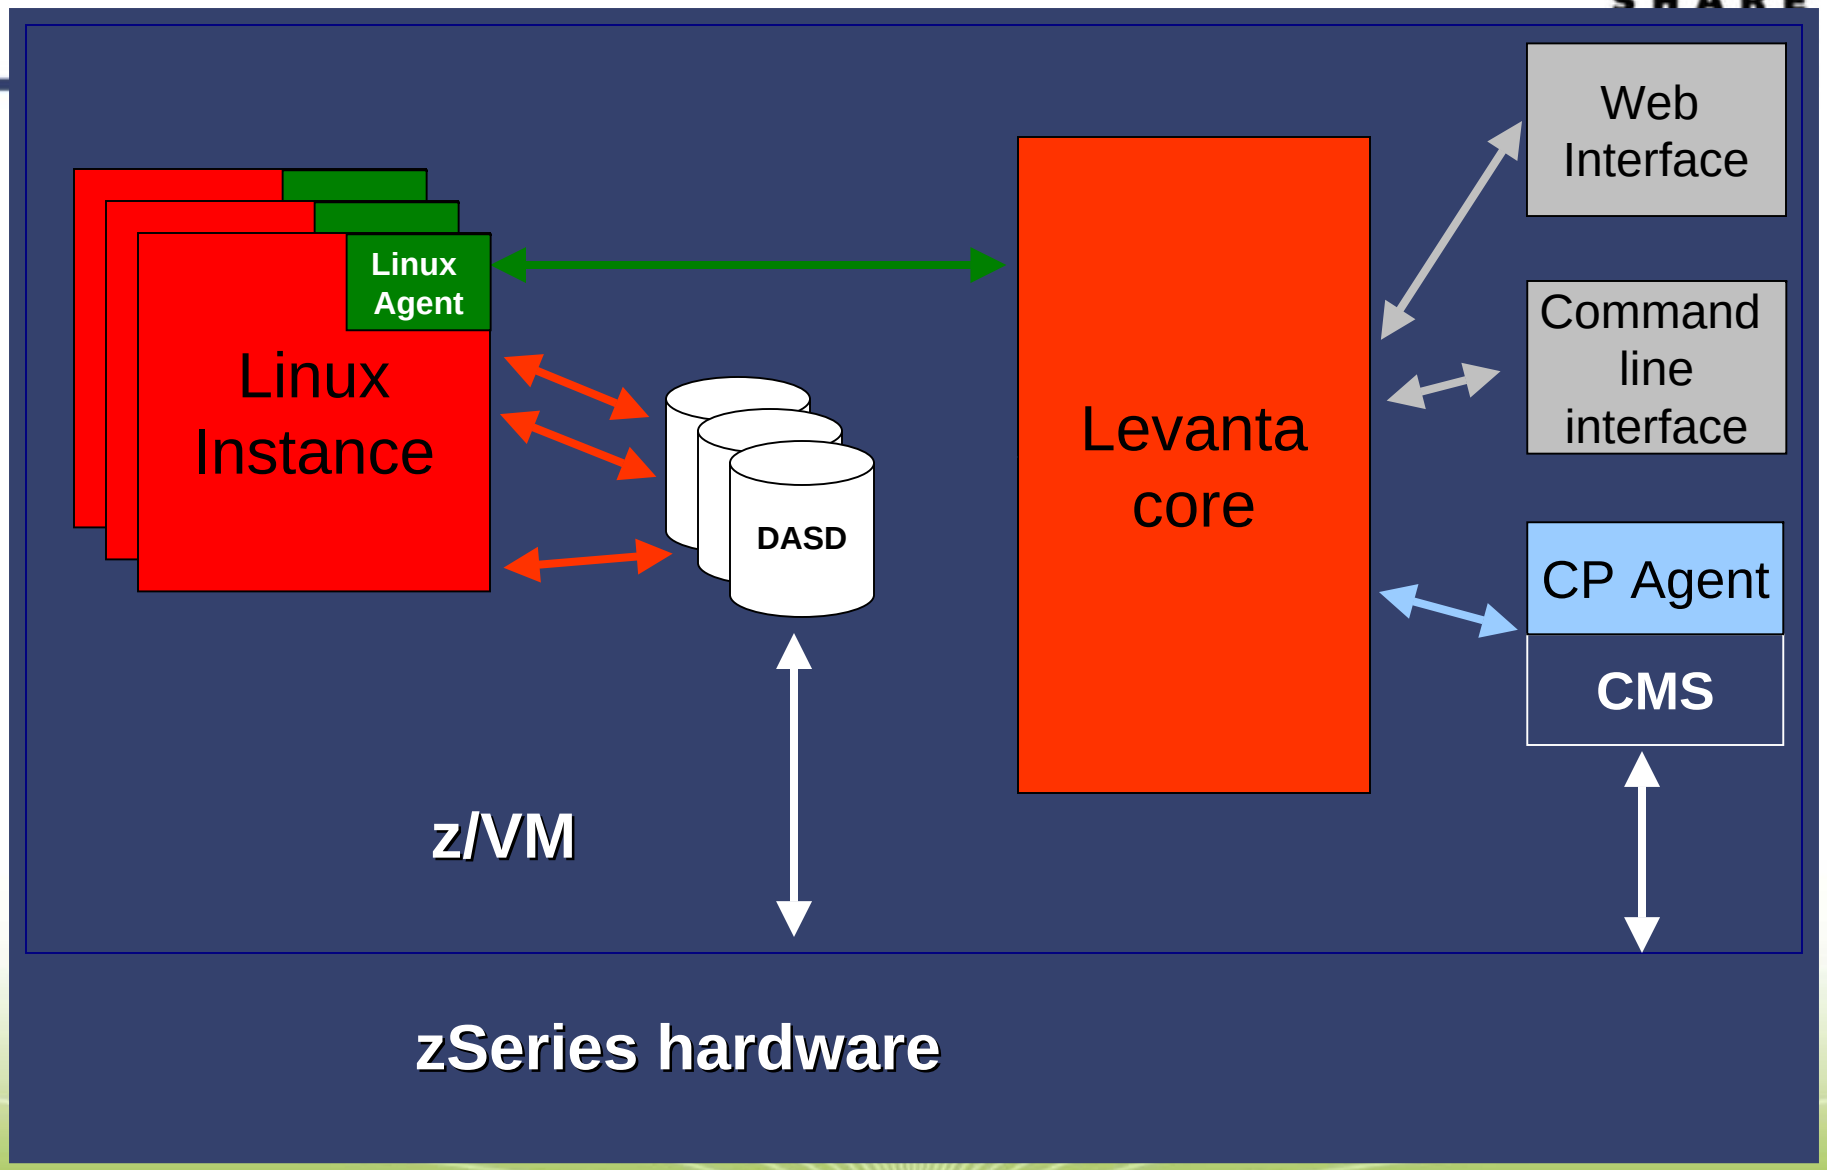

# Levanta Benefits: Staff Efficiency

- 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

# 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


USER EVENTS
SHARE
Technology · Connections · Results
SHARE.ORG

# The Future…


Enterprise IT management
SYMPOSIUM


Application Development & integration
SYMPOSIUM


SHARE
User-Driven Training
Event & Expo


Enterprise Systems management
SYMPOSIUM


USERBLUE

# Beyond Levanta 3.0


- Other hardware (e.g., IBM pSeries)
- Other operating systems (e.g., Sun SOLARIS™)
- Dynamic server provisioning for load
- Dynamic server failover
- Server provisioning without disk sharing
- Synchronizing replicated servers without disk sharing

# Beyond Levanta 3.0

- Distributed data center management

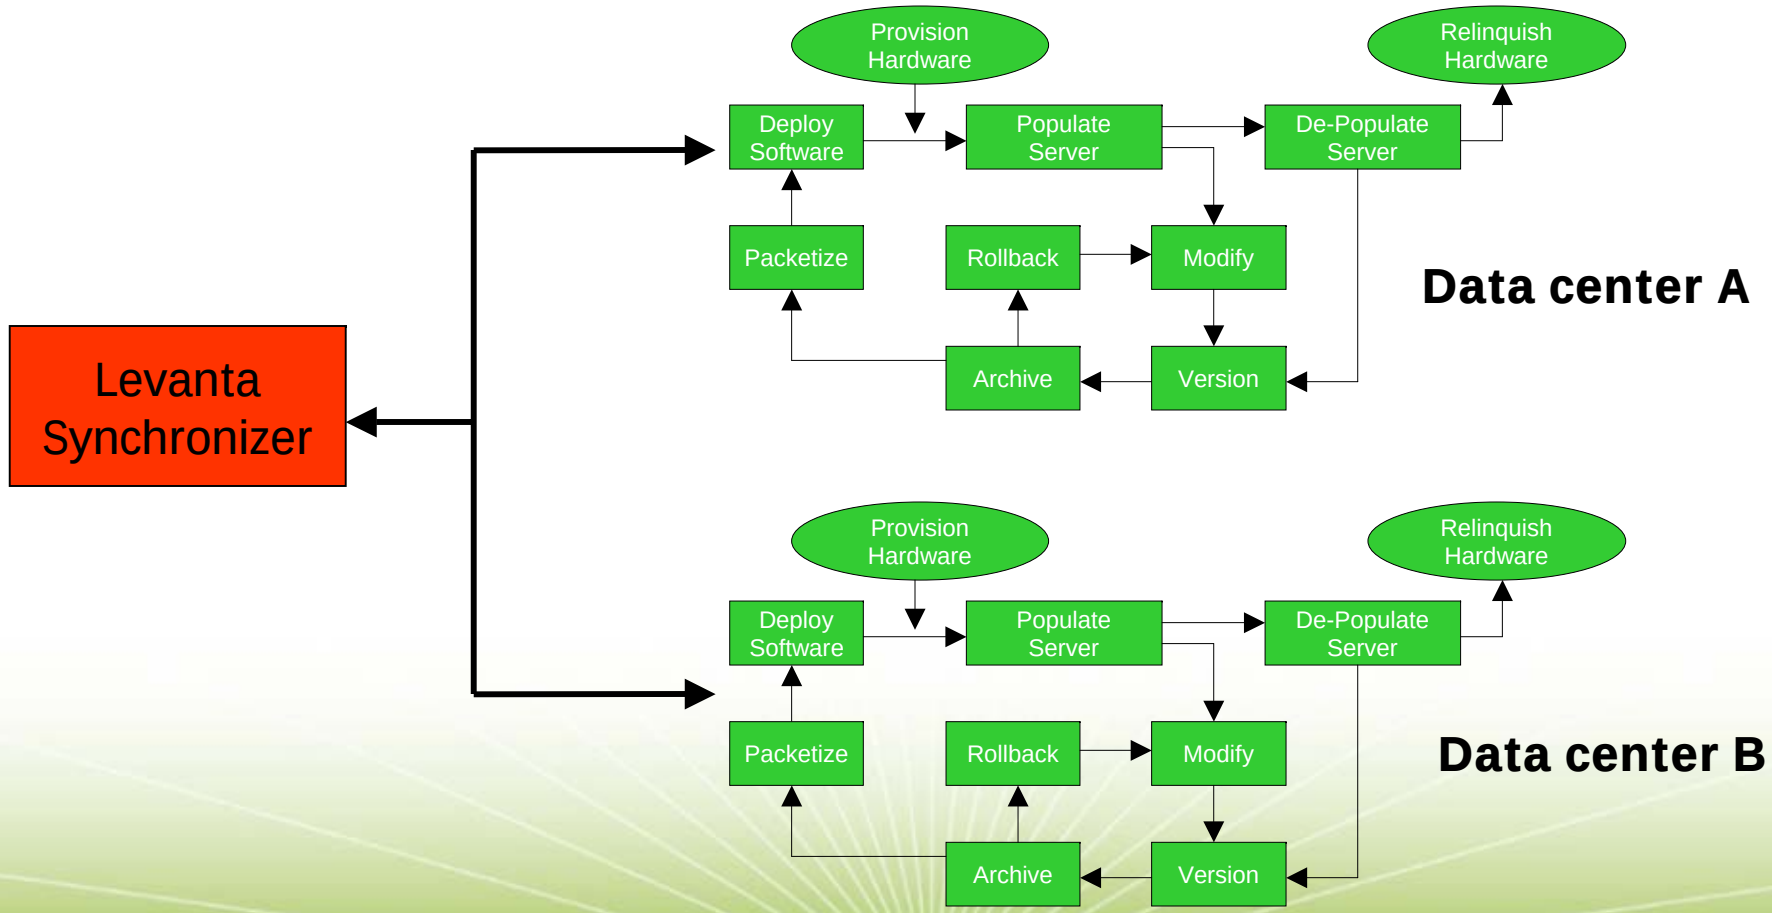

# Levanta: Proven Benefits


- Lowering system administration effort reduces costs
- “Cultural collaboration” capabilities ramp pilot-to-production process
- Self-service model for server deployment eases administrative management
- “Best practices” capabilities and captured intelligence via templates simplify operation
- End-to-end virtual server deployment process saves time
- Simplifying z/VM configuration tasks aids those with limited expertise
- DASD sharing saves costly resources

# Summary

# The Levanta Value Proposition


- Increased staff efficiency (“force multiplier”)
  - Current staff can handle workload without retraining/hiring more staff
- Enables “Best practices”, improved RAS
  - Extends data center change management processes/control to Linux
- Cross-team collaboration
  - Opens up IBM zSeries virtual platform to UNIX/NT teams

# 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 view is worth the climb”, 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
- Levanta can help!
  - “Operators are standing by…”

# Questions?

Phil Smith III

LEVANTA

703.476.4511

psmith@levanta.com

www.levanta.com