Building a strong z/VM and Linux on the mainframe architecture

International zSeries Oracle SIG

April 30, 2008

David Kreuter

dkreuter@vm-resources.com

Table Of Contents

Client context

- IT service provider for many government offices (125)
  - Already a mainframe shop
  - 5 z890 + 2 z800 + 1 G5 on the floor on 3 sites
  - 1 z9 EC dedicated to Linux on z/VM
  - 450+ physical servers (750+ logical) (HP, SUN, pSeries, …)

- Orientations :
  - Promote the mainframe environment
  - z/VM is the prime choice for future projects
  - Server consolidation is a priority
  - This project is in line with the new « online government » policy

Client context
Project origin

- Initial needs :
  - Must solve many issues with the intermediate platform
    - Many operation systems
    - Many versions
    - Unsupported software
    - Unsatisfactory DR
    - Fast growing (unprecedented growth)
  - Understaffed
  - Need a flexible solution with rapid deployment
- Mainframe is a stable and mature environment
  - Staff is available and at early stages of their careers
  - Solid and well controlled DR process (MVS-like)
- The conclusion : GO with z/VM

Client context
Timeline of the project origin

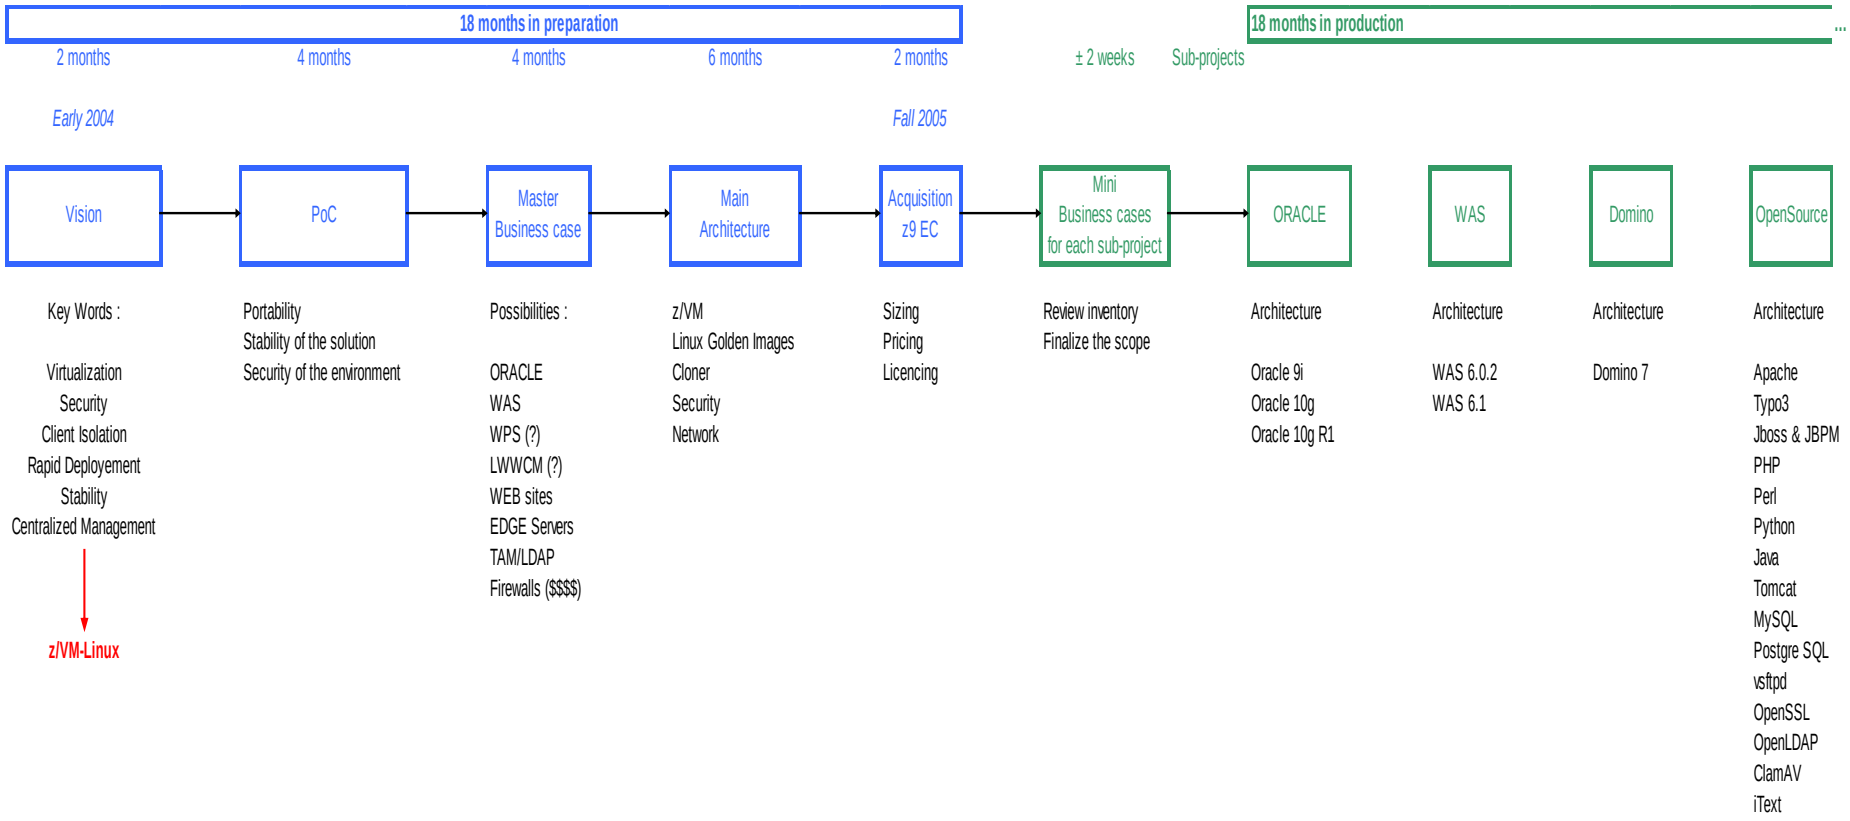

Client context
z/VM Linux Environment

- 1 z9 EC mainframe with 5 IFLs (~ 2750 mips) + 96GB
- 10 LPARs
  - Oracle/DB
  - WAS (2)
  - WAS Portal (2)
  - OpenSource
  - Domino (2)
  - Service Zone
  - Lab Zone
- 40+ internal networks
- Software
  - Novell SLES (versions 9 & 10)
  - z/VM v.5.3
  - Oracle/DB (versions 9i & 10g & 10gR1)
  - Velocity Software Performance Tools
  - CA products (Automation, Scheduler)

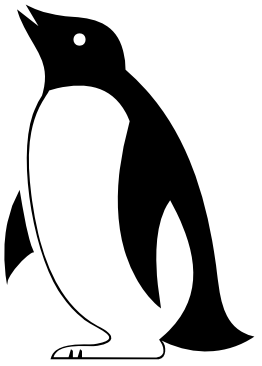

Architecture
Guaranteed isolation of multiple clients

- The environment serves the needs of over 125 clients, some large, some small.
- Clients must have their applications and data separated from each other.
- The challenge is how to do this in a centralized shared environment.
- System z hardware with z/VM leads the way!

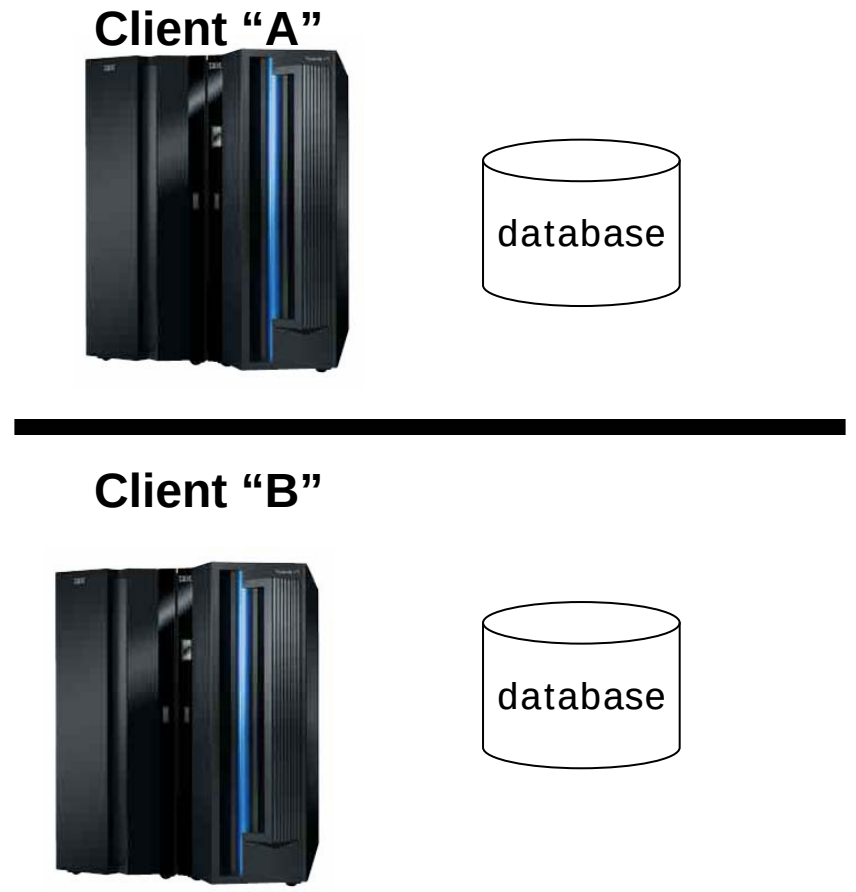

Architecture
RACF Serves and Protects

- RACF ensures isolation as it provides security for these protected resources and events in z/VM:
  - Logon
  - Link
  - VSWITCH
  - VLAN
  - Shared File system
  - VM FTP

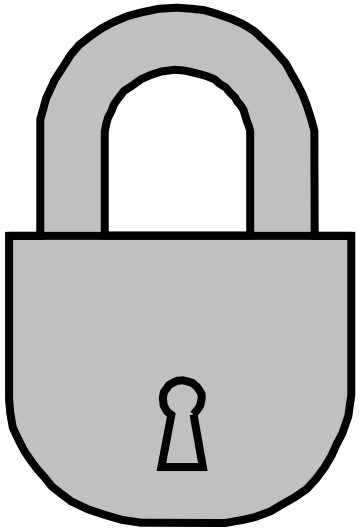

Architecture
We harden our Linux on the mainframe servers

- The Linux golden images are hardened, tested and certified by an independent team before allowing the image to be cloned.

Hardening tasks:

- Removing unneeded login accounts
- Removing many supplied services such as FTP, Telnet, and NFS.
- Sifting through the startup /etc/rc.d tasks and removing unneeded tasks.
- Using PAM authentication with strong password practices.
- Using Tripwire to inventory software and for file anomaly detection.
- Ethical hacking done on a regular basis for penetration testing and cracking.
    - Certified by an independent team.

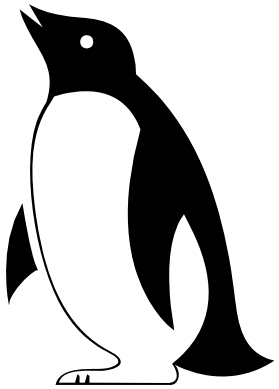

Architecture
Networks within the box

- A variety of choices for networks that keep the clients isolated
  - **OSA devices**
    - *Traditional connectivity from mainframe to physical switches*
  - **HiperSockets**
    - *Inter and Intra LPAR connectivity*
  - **Guest LANs**
    - *Connect virtual machines on virtual networks within an LPAR*
  - **VSWITCHes**
    - *Connect **guest LANs** to physical switches using **OSA devices***
    - *40+* VSWITCHes on 8 LPARs

Architecture
Networking with the real world

- We use a lot of VSWITCH networks.
  - 40 …
- VSWITCH connects to OSA port as conforms to the physical network topology.
- Redundancy provided only for production networks.
  - Handled within VSWITCH connecting to multiple unique OSA ports.
  - Does not require VIPA
- Some OSA ports shared across zones in multiple LPARs.
- Firewalling done downstream from the mainframe.

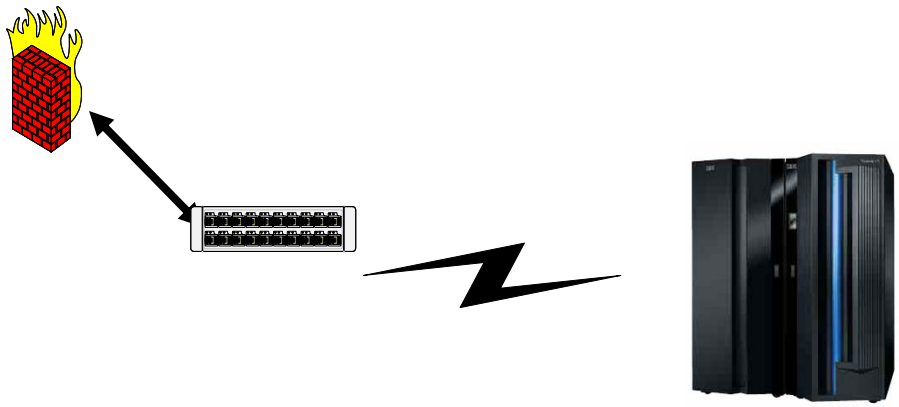

Architecture
Production Network Redundancy

- 7 of the networks (production) have redundancy with dual OSA ports.
  - *All others networks (30+) do not have redundant networking*
- Managed by Vswitch.
- Connect to different physical switches.
- Switches are bridged.

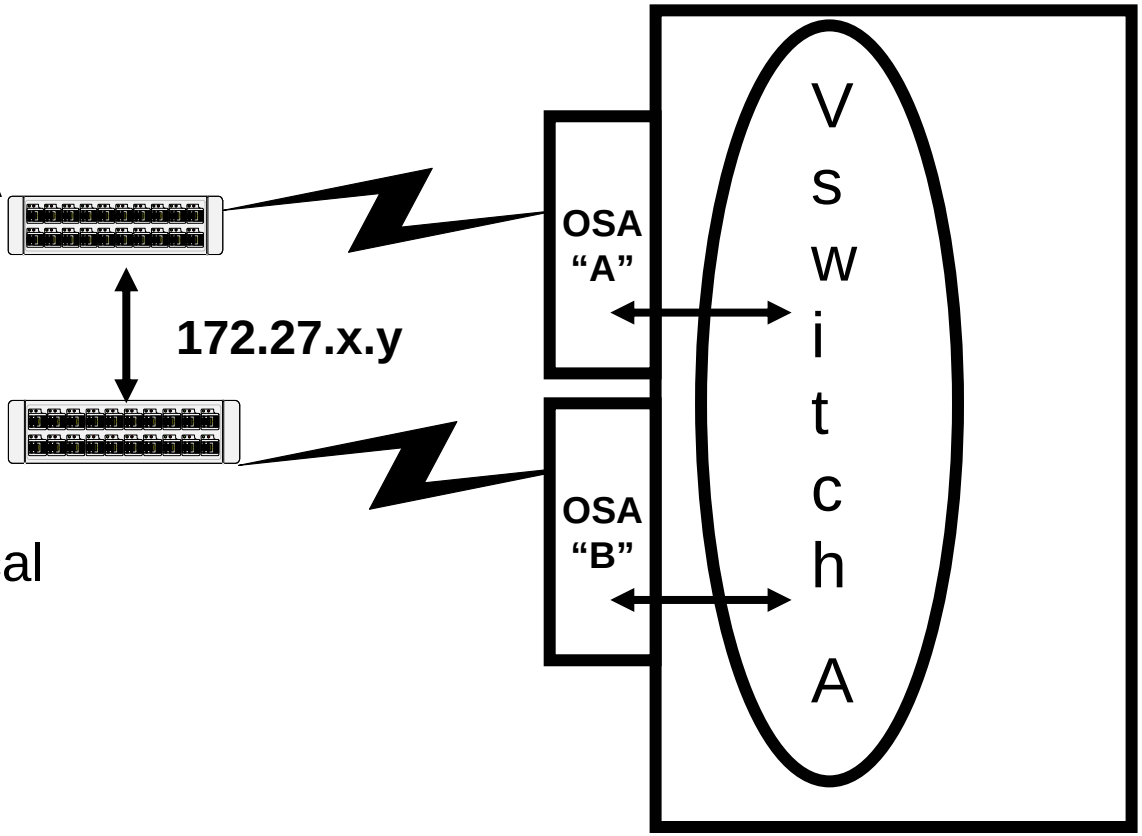

Best Practices

Best practices

- In our project we planned to utilize best practices for systems and network management.
- Examples of in use best practices:
  - Networking:
    - Performance data collection using private VSWITCHes
    - Manage multiple networks from a single TCPMAINT
  - Systems:
    - Golden images (z/VM & Linux)
    - Cloning engine
    - Sharing resources the client way
    - Sharing resources the IBM way

Best Practices
Performance data collection using private VSWITCHes

- A TCPIP stack with multiple guest LANs and VLANs collects data for the Velocity SNMP data collection.
- The VSWITCHes are defined without real devices.
- Membership in the VSWITCH and VLAN is RACF protected.

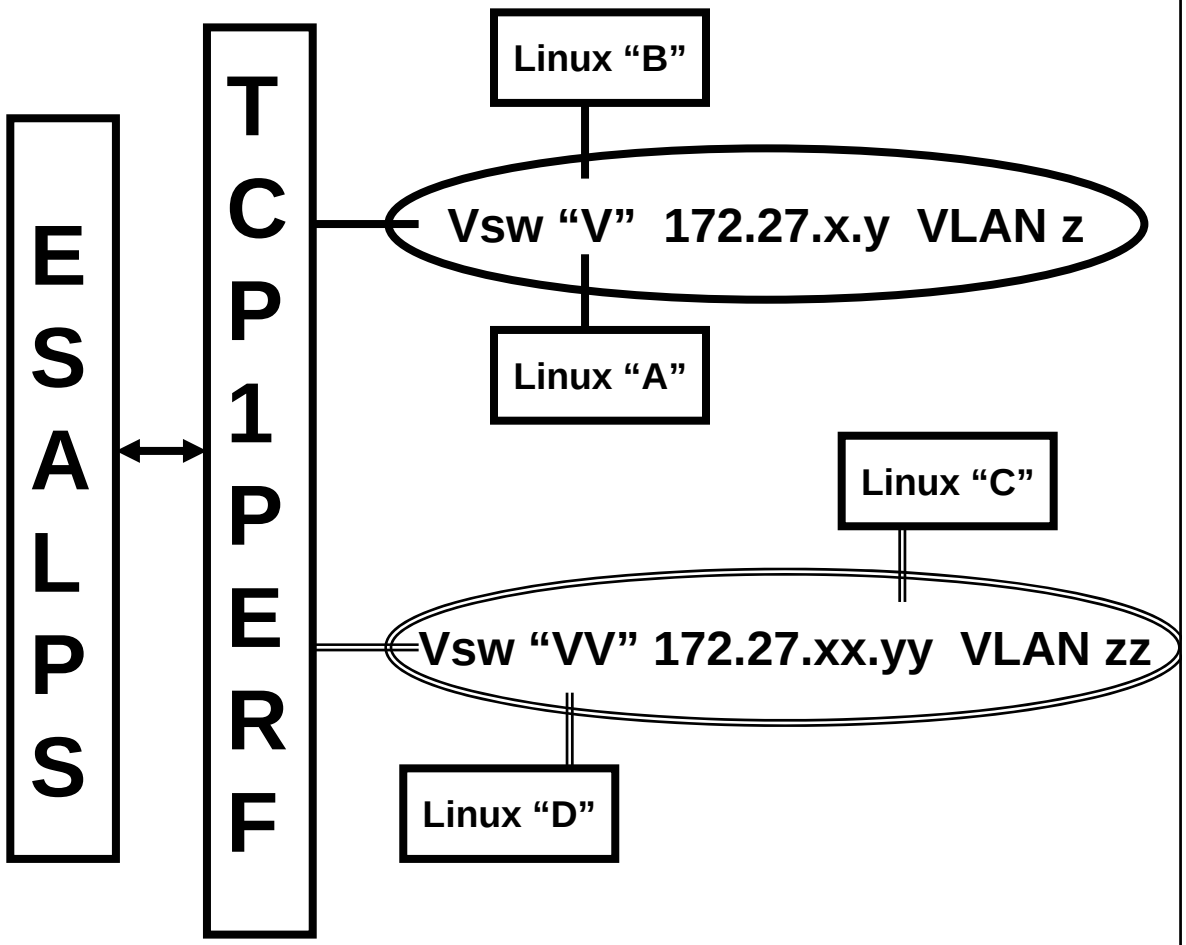

Best Practices
Administering multiple z/VM TCPIP machines from a single TCPMAINT

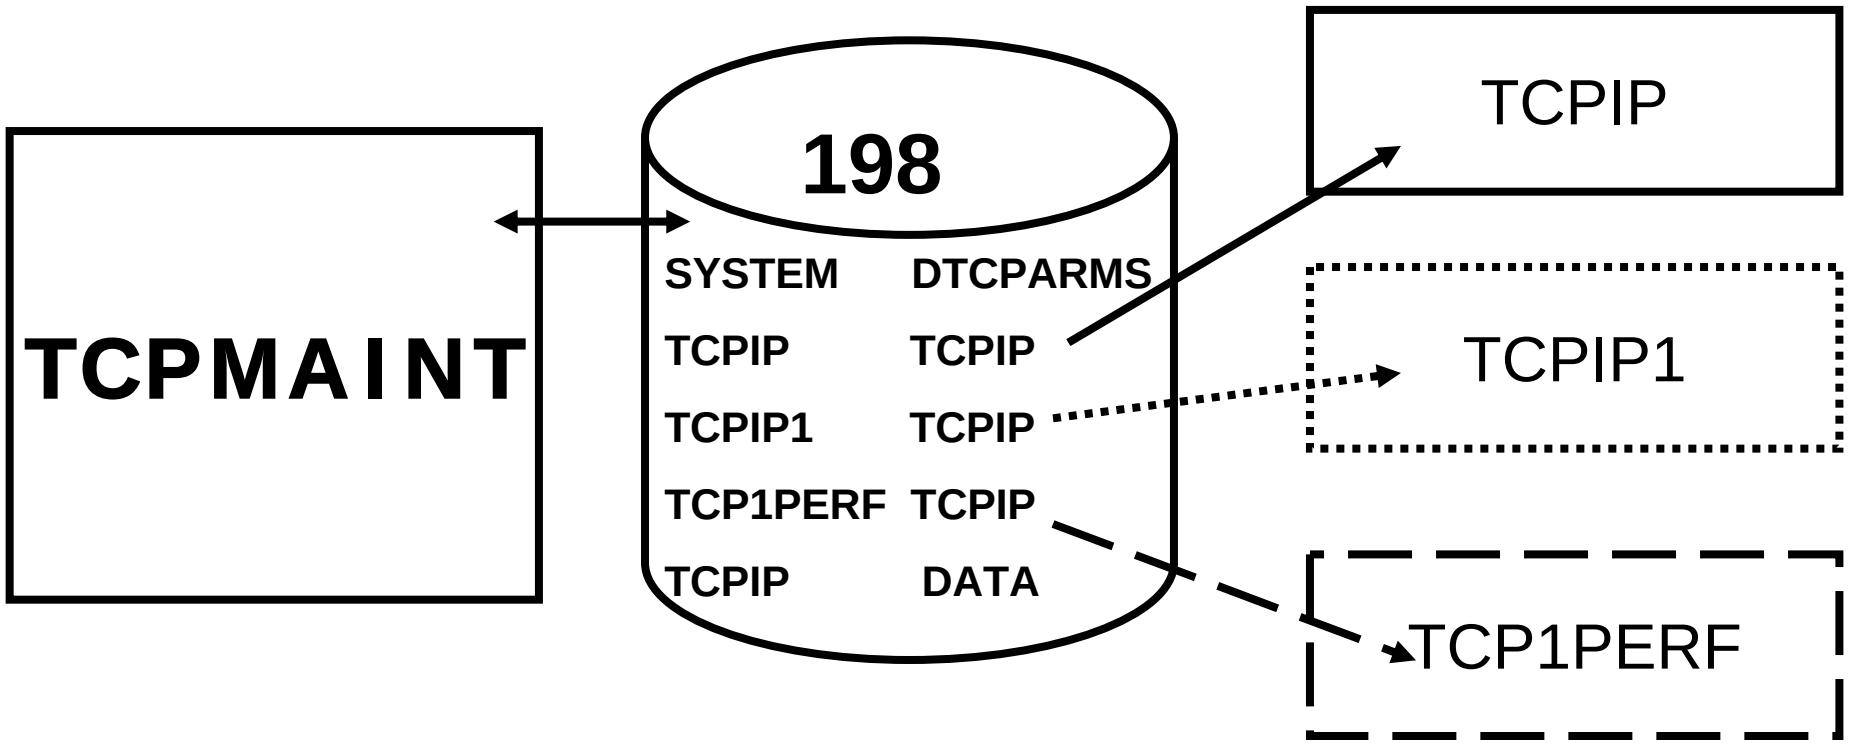

Resource sharing
HiperSockets network on the z9 EC

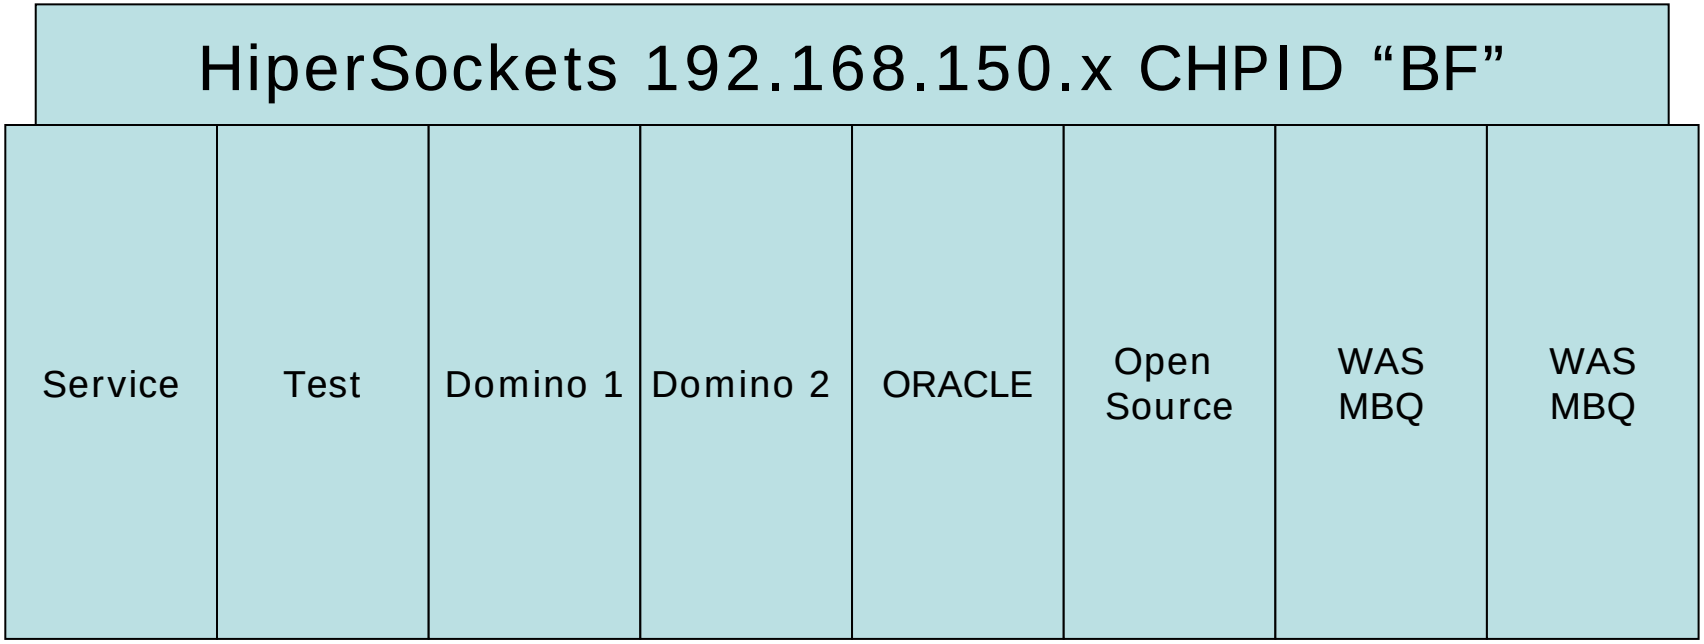

- **Internal network only.**
- **Used for administrative purposes.**
- **Applications include the cloner, telnet, RSCS (file transfer and message queues).**
- **Secure memory-to-memory transfer.**

Best Practices
Golden images : z/VM & Linux

- Our z/VM golden image:
  - z/VM 5.3.0
  - All production MDISKs on one volume per system
  - Serviced from one system
  - One flavor
- Our Linux on the mainframe golden images:
  - Novell SLES 9 or 10
  - Hardened
  - One application flavor (Oracle, WAS, Domino, MBQ …)
  - Input to the cloner
- Both are rigorously tested and certified

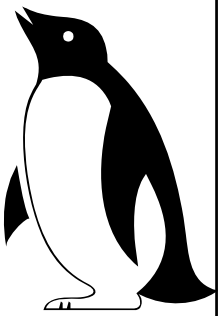

Best Practices
The Linux Golden Image
"*install once and clone often*"

- The golden image is really black and white and waddles on ice but not until:
  - Installed
  - Serviced
  - Hardened
  - Tested by various groups
  - Passes security penetration tests and certification
- There are a few masters and many many many clones!

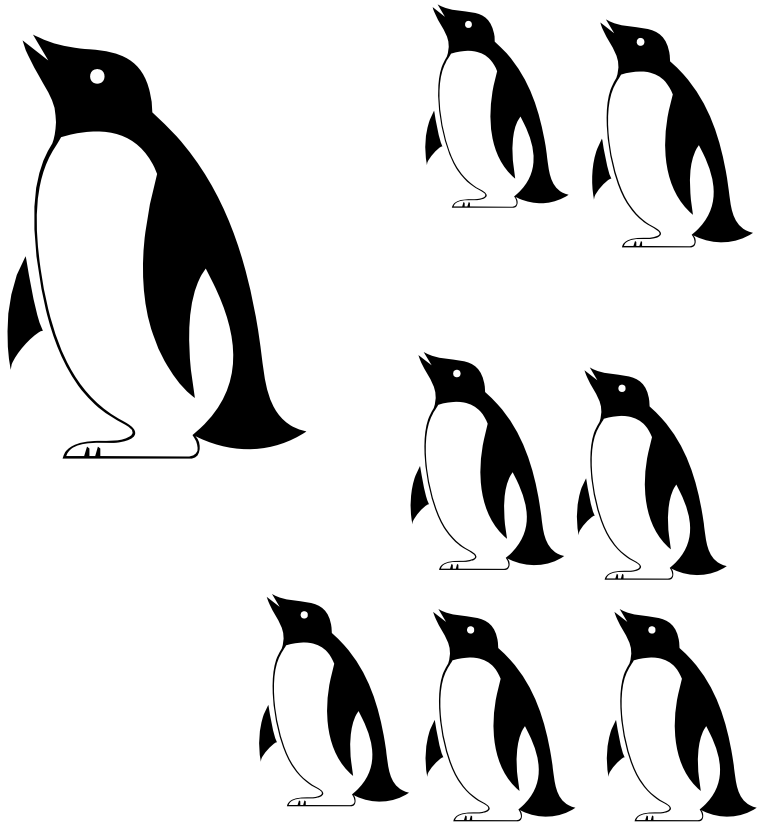

Best Practice
Our cloner : Overview

- Hand crafted
- Pride of ownership
- Not a disk copier
- Intelligent decisions:
  - Choice of Linux
  - Choice of application
  - System and application position
  - VSWITCH membership
  - VLAN membership
  - IP address
  - Data replicated
  - Strong passwords

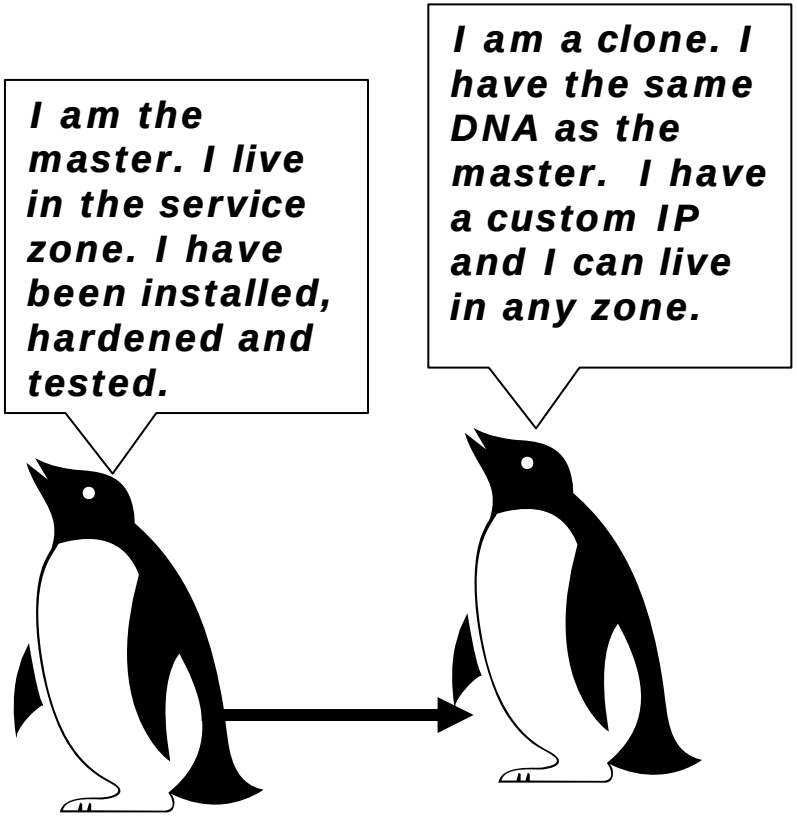

Best Practices
The cloner

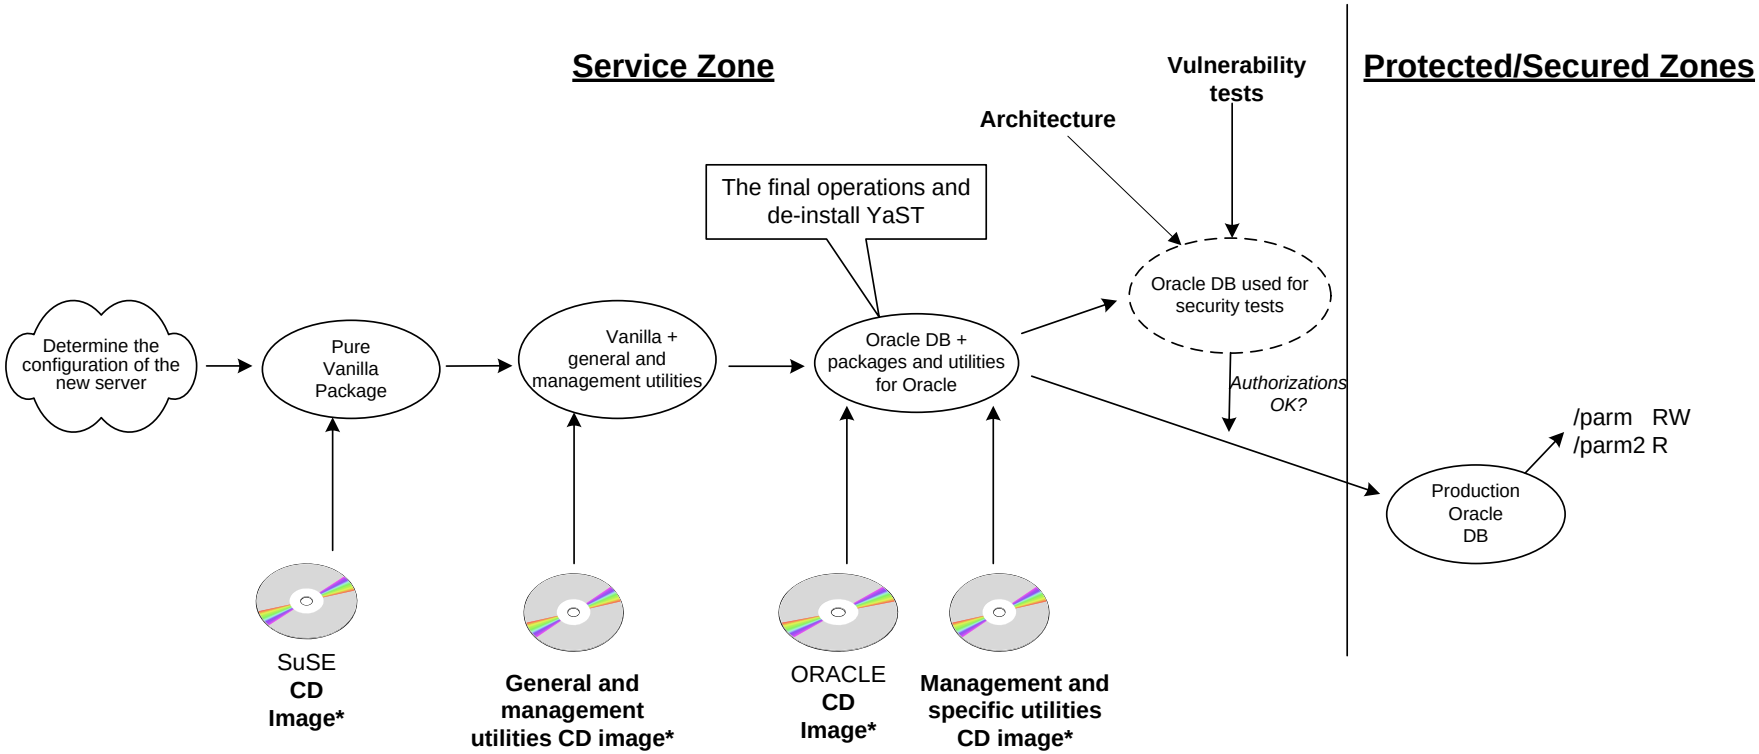

* : These images reside on a virtual Linux server in the service zone for access via FTP. This server contains software libraries.

Best Practices
The big picture of the cloning

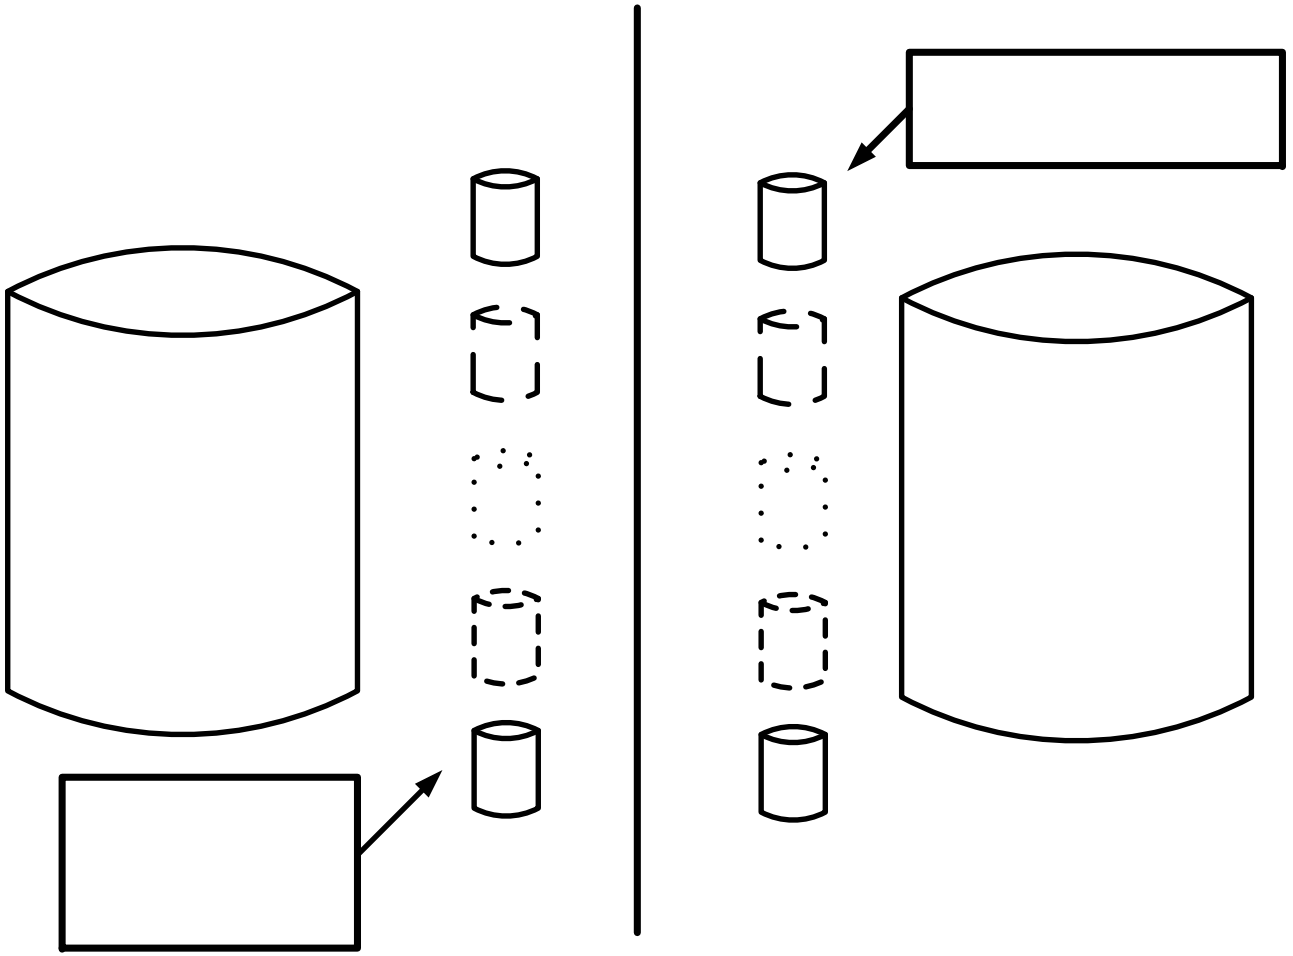

Best Practices
Our cloner: Coding and interfaces

- Coded in REXX and PIPELINEs.
- Interfaces to DIRMAINT and RACF.
- Inputs include which system, application, storage size, etc.
- Interfaces with 3270.
- Can clone only from service zone to any other zone.

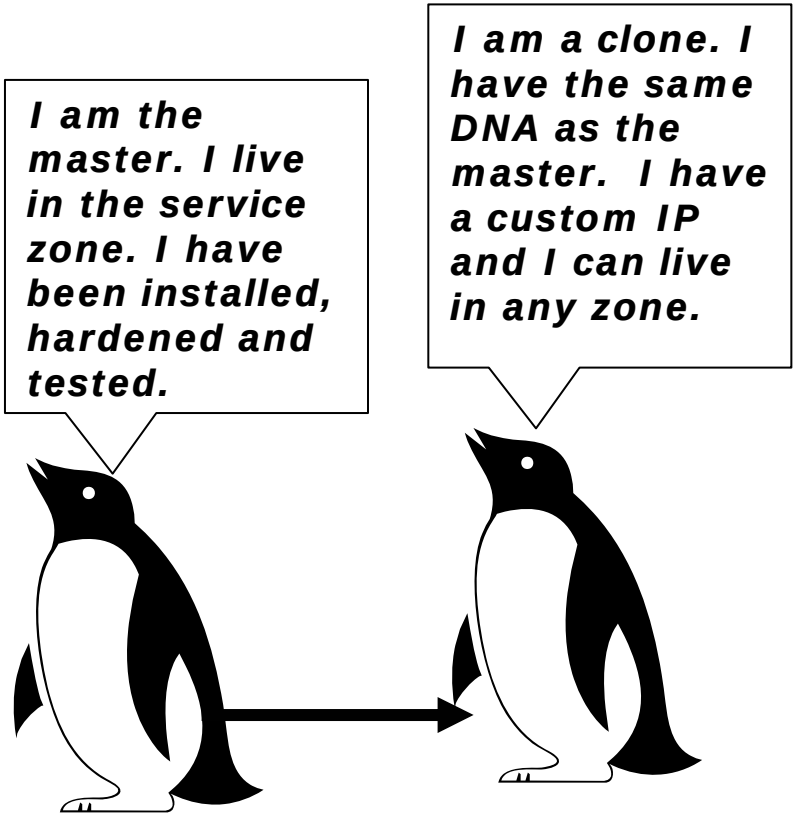

Best Practices
Resource sharing

- IBM way (old school – 35+ years)
  - CPU
  - Memory
  - Minidisk i/o
  - Spooling
- Also sharing:
  - Linux file systems (/usr)
  - Heavy usage of VSWITCH
    - Lots of guest LANs
    - Many OSA ports

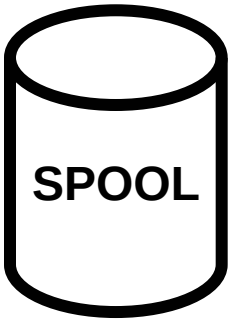

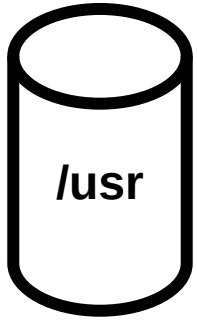

Base Cloner links and minidisks

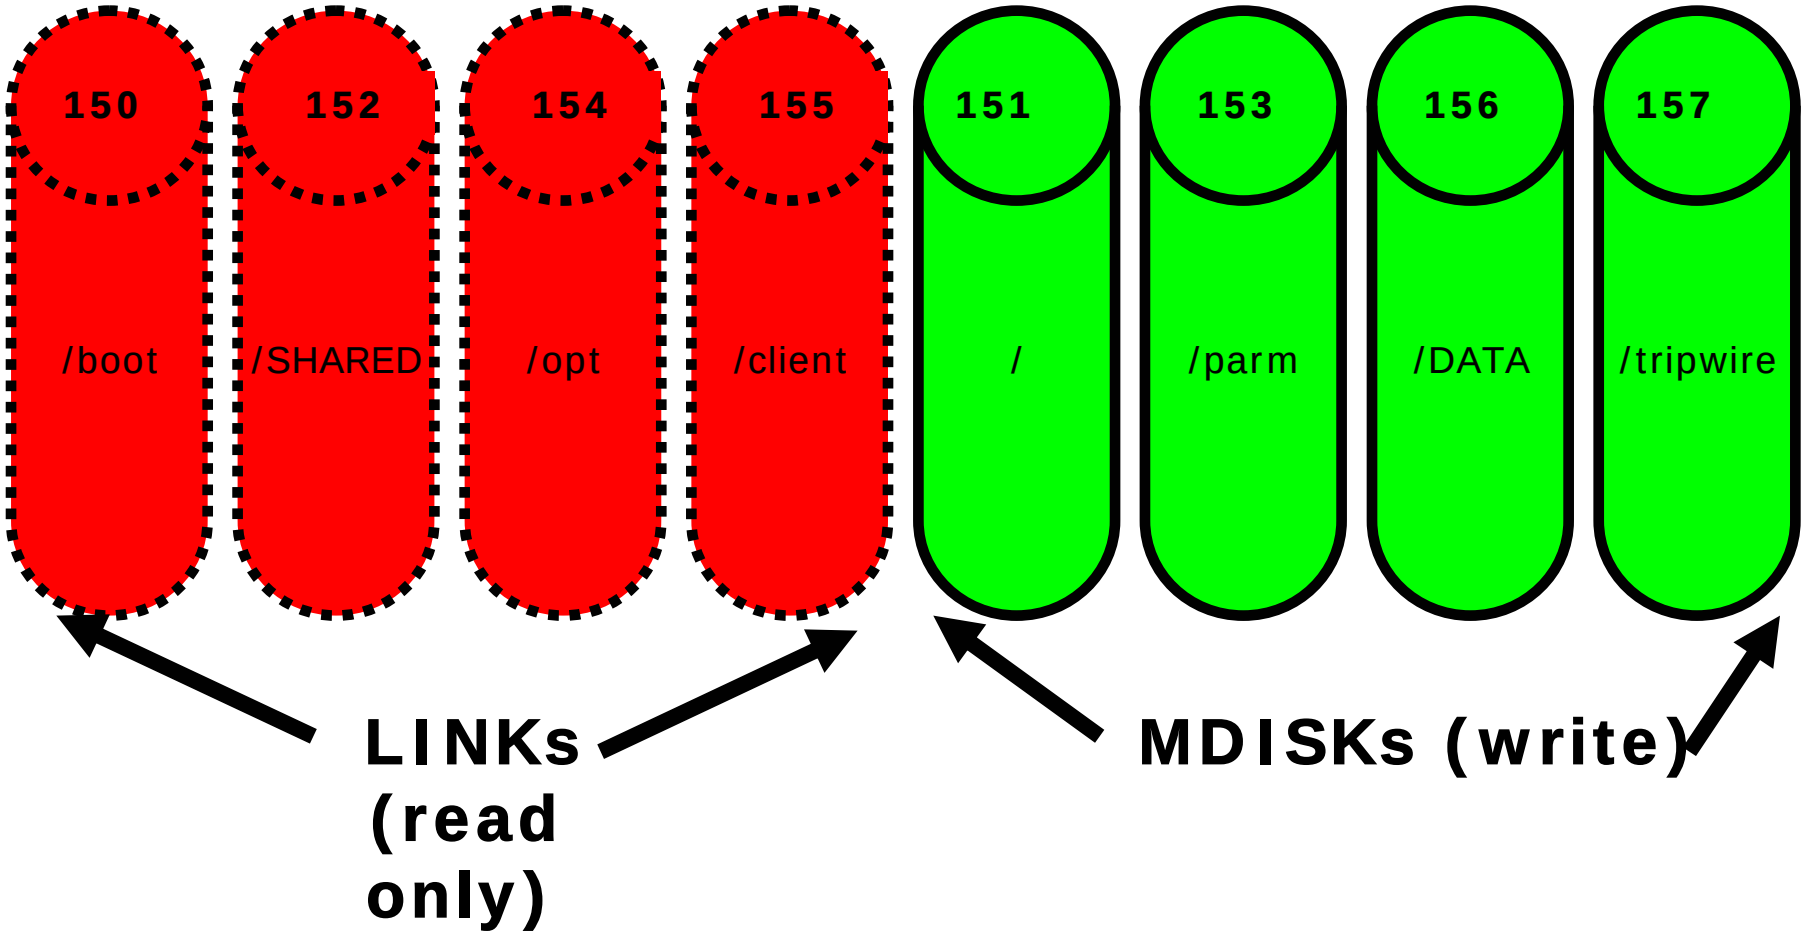

Best Practice
OpenSource initiatives

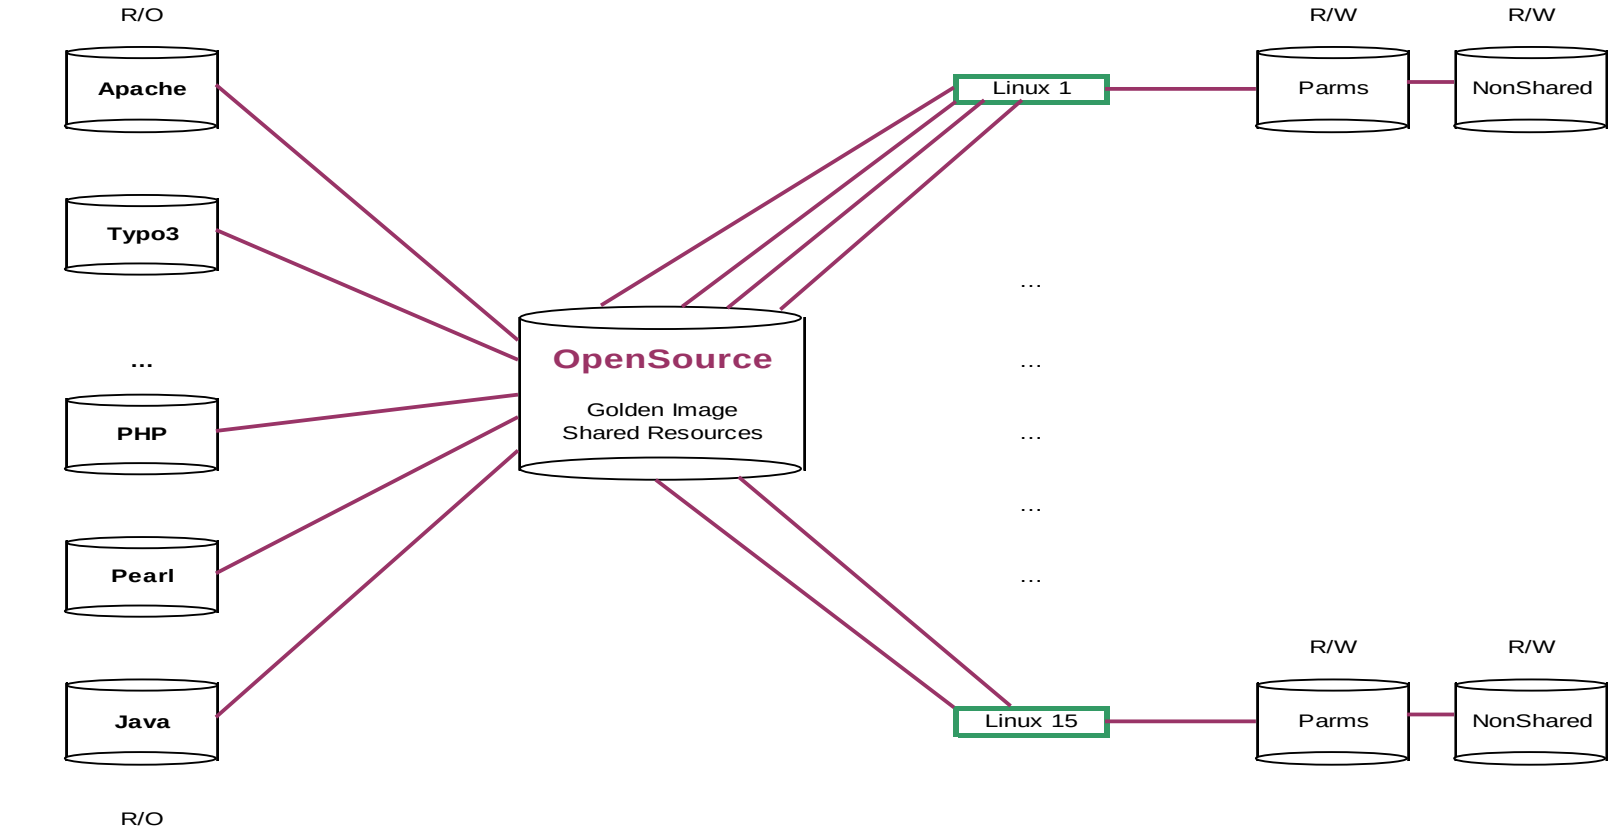

Project Status

- Oracle Databases
- WAS Migration
- Domino Pilot
- OpenSource initiatives

Project Status
Oracle Databases

- Original migration almost complete (95% completed)
  - Close to 200 Oracle instances
  - Running within 150 virtual machines (Linux)
- Current projects
  1) Oracle upgrade from v.9i to v.10gR1
  2) Oracle upgrade from v.10g to v.10gR1
  3) Linux upgrade (Novell SLES 10 + maintenance)
  4) Second generation of the clones (Golden Image v.2)
  5) Data migration from Lightning to USP from HDS
- Because of the cloning engine and sharing strategy
  - Realization time = 4 months (Sept. To Dec.) vs. 12+ months using the old fashioned way (mid-range hardware without cloning)

Project Status
WAS Migration

- Migration project WAS, WMB & MQ to System z
  - Close to 90 instances of WAS
  - 12 WMB & MQ servers
- Current projects
  1) WAS upgrade from v.5 to v.6.0.2
  2) WMB upgrade from v.5 to v.6
  3) MQ upgrade from v.6 to v.6
  4) Second generation of the clones (Golden Image v.2)
  5) Platform change
- Because of the cloning engine and sharing strategy
  - Realization time = 4 months (Sept. To Dec.) vs. 12+ months using the old fashioned way (mid-range hardware without cloning)

- ***Note*** : We can do both projects (Oracle & WAS) in parallel !!!

Project Status
Domino Pilot

- Current project
  - Pilot for dozens of users
  - Domino Server v.7
  - Notes clients v.5, v.6.5 & v.7
  - Outlook clients (with DAMO)
- Go/No-Go decision in Sept. 2007
- Total migration of one Government department
  - About 1800 users
- If migration successful (by the end of 2007)
  - Development of a Government Offering
    - Reduction of 900 mail servers

Project Status
OpenSource initiatives

- Current projects
  1) Government Intranet
  2) WEB sites (15)
  3) JBoss project with a Government department
- The current offering (16)
  - HTTP server : Apache
  - File Transfer : VsFTPD
  - Languages: JAVA & PHP & Pearl & Python
  - Databases : MySQL & PostGreSQL
  - Content manager : Typo3
  - WAS-like : JBoss & JBPM & TomCat
  - Security : OpenSSL & OpenLDAP
  - Anti-Virus : ClamAV
  - PDF generator : iText

- More packages can be added on-demand within the OpenSource Golden Image (between 2 days and 2 weeks of effort including adjustments of the cloner)
- Customers waiting avidly for this offering

Lessons learned Volume 1

- Acceptance of virtual servers quicker than expected.
  - Grew to 200 Oracle servers ahead of plan.
- Fully tasked personnel (big shoulders):
  - Confirmed our expectation that 2 Linux administrators can support all virtual Linux servers.
    - *100:1 ratio of Linux virtual machines to administrator*
  - 2 z/VM systems programmers supporting 10 LPARs: (could support many more)
    - *New to z/VM*
    - *Mentored by consultant*
    - *z/VM support integrated into MVS group (year end 2007)*
- Less than fully tasked personnel (arms and legs):
  - *Security administrator*
  - *Network programming*
  - *Storage*
  - *Automation*
  - *Performance*

Lessons learned
Volume 2

- Big win early win with successful disaster recovery.
- Administration and reporting on centralized servers is excellent.
- Lots of new documentation and procedures integral part of project.
- Lots of training required.

Lessons learned
Volume 3

- Critical mass of servers required – use more than 1 Linux virtual machine for benchmark, POCs, and business case!
- Initially, project was done for the $ savings, now the important gains:
  1) The flexibility of the solution
  2) Disaster recovery
  3) $ savings
- You must have a sponsor. Our sponsor was the operations directorate for the mainframe business interested in solving DR issues.

Lessons learned Oracle Project

- Mostly business as usual for the DBAs:
  - Use SSH client or “X” windows (no 3270 usage)
  - DBAs comment on rapid performance of I/O
  - DB loading faster than in other platforms.
- Benign ignorance of the virtual machine
  - Linux administration performed by Linux sysadmin.
  - z/VM administration performed by VM sysprogs.

- Rapid creation of new databases in virtual machines for testing, acceptance, and production.
- Initial install was difficult but once incorporated into cloning methods subsequent installs quick and easy.
- Almost all client needs satisfied with ORACLE cloned image (they don’t know).
  - ~ 2% require some sort of customizing.

Conclusion

- z/VM and Linux on the mainframe: a powerful combination for applications hosting.
- Supported open source software on the mainframe provides the stability of z/VM with the ability to run modern applications.
- Service being offered to many government offices and agencies.
- The word is out that z/VM and Linux on the mainframe is a good place to host your applications:
  - Internal government emails and announcements from the project office promoting z/VM and Linux on the mainframe solution.
- Rapid growth is forecasted.

Conclusion

- We’re providing infrastructure to many offices and agencies.
- Building and nurturing business case critical to success of the project.
- The training was a vital part of the client acceptance of the concept.
- Architecture was developed and polished for over one year (on going activity).
- z/VM and Linux on the System z natural fit for the vertical and horizontal growth.
- Oracle project won the Share 2007 Award for Excellence in technology
- Project success will continue into the future!