

User Experience with DataMart on Linux for System z at FMS

Jim Moling
US Treasury, Bureau of the Public Debt

Wednesday, October 19, 2011

Hillgang Fall Meeting

Disclaimers

- **The opinions & ideas expressed herein are those of the author alone and do not necessarily reflect those of Financial Management Service, furthermore, Financial Management Service is hereby absolved of any and all responsibility or liability for the information contained herein.**
- **Copyrights & Trademarks:**
 - Any and all copyrights & trademarks are hereby acknowledged to be owned by their respective parties
 - All other brands, logos and products are trademarks or registered trademarks of their respective companies
 - All rights reserved
- **Disclaimer of Endorsement:**

Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government, and shall not be used for advertising or product endorsement purposes.
- **Disclaimer of Liability:**

With respect to this presentation, neither the United States Government nor any of their employees, makes any warranty, express or implied, including the warranties of merchantability and fitness for a particular purpose, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.

Overview

- Goals Of the DataMart Project
- VM Environment & Infrastructure for Linux applications
- DataMart Architecture
- Lab Environment a.k.a. “sandbox”
- Establishing the Baseline Application Server
- Server Security Considerations
- Finalizing the Baseline Server & Cloning
- DataMart Server Configuration
- DataMart Developement in the Lab
- Deploying DataMart across 4 Environments
- How is the DataMart doing? Were the goals met?

Goals of the DataMart Project

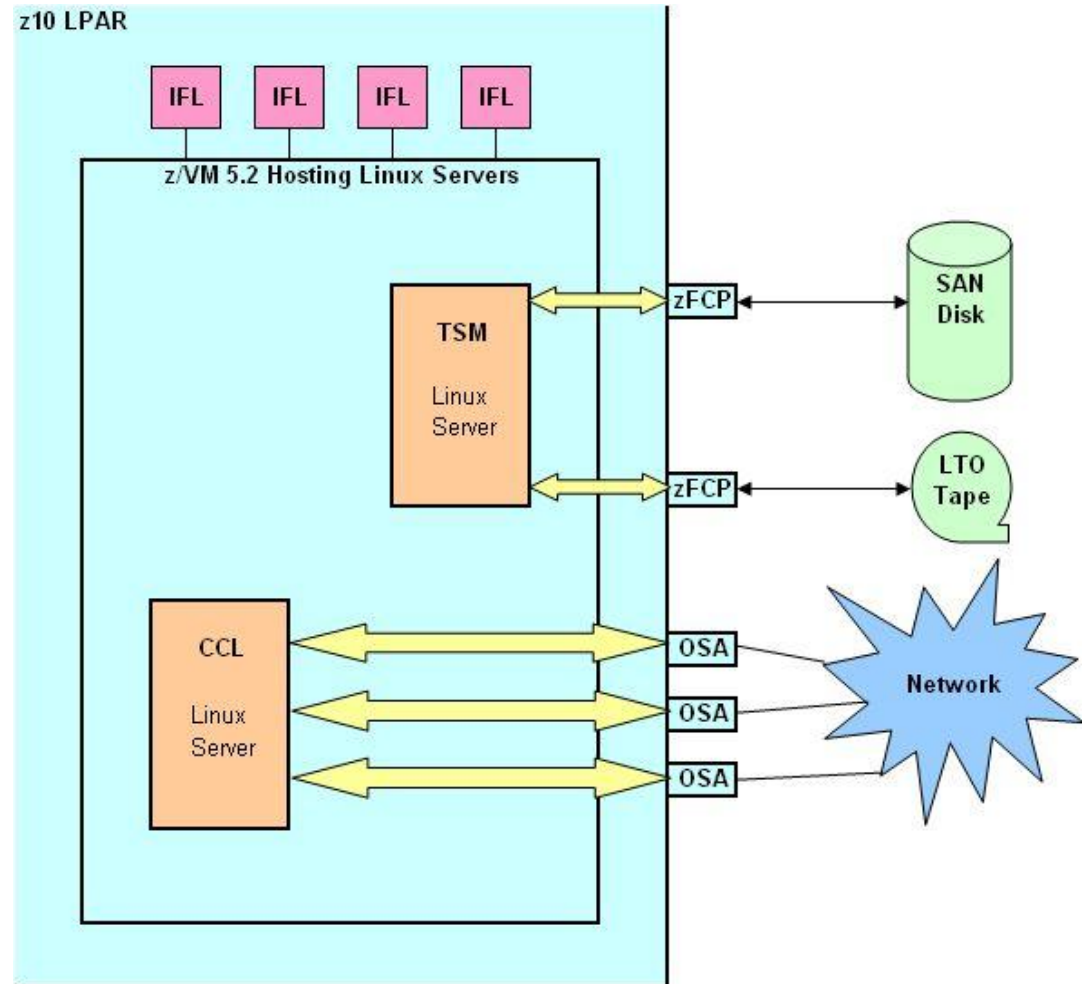
- Proof of concept – Linux on z Series Application Servers
- Establish the z/VM infrastructure for running SLES-based Virtual Linux Application Servers
- Develop a secured SLES server image for cloning
- Establish Roles & Responsibilities for running Linux on z Series
- Deploy a DataMart for executive review

z/VM Environment for Linux Apps

- IBM z10 w/ 4 IFL's
- Used z/VM 5.2 (currently upgrading to 6.1)
- Virtual switches (trunk & access) used for network access
- zFCP adapters used to access SAN Disk & LTO Tape
- 3390 DASD used for server infrastructure
- SAN Disk used for server data

z/VM Environment for Linux Apps

- Already had a z/VM LPAR established for Linux
- TSM Server used for server backup & restore (both virtual & real)
- CCL Server used for z/OS SNA traffic (replaced 3745 controllers)




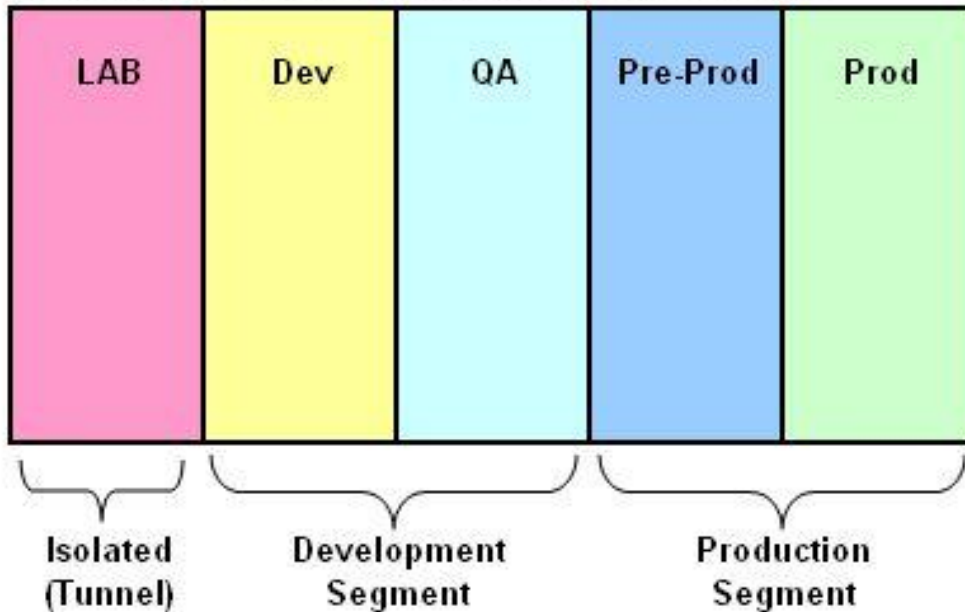
z/VM Infrastructure for Linux Apps

- 5 Environments: Lab, Dev, QA, Pre-Prod & Prod
- TSM Linux Server used for SAN Disk backup/restore
- MVS-based backup/restore used for 3390 DASD volumes
- 2-Step Server Recovery:
 1. Restore 3390's via MVS & IPL server
 2. Restore SAN Disk via TSM from running server
- Future: use XRC-based DASD mirroring for 3390 DASD volumes (after VM 6.1 upgrade)

z/VM Infrastructure for Linux Apps

- Servers are associated to one of the 5 environments by connection & name

Each server has an instance in each environment:



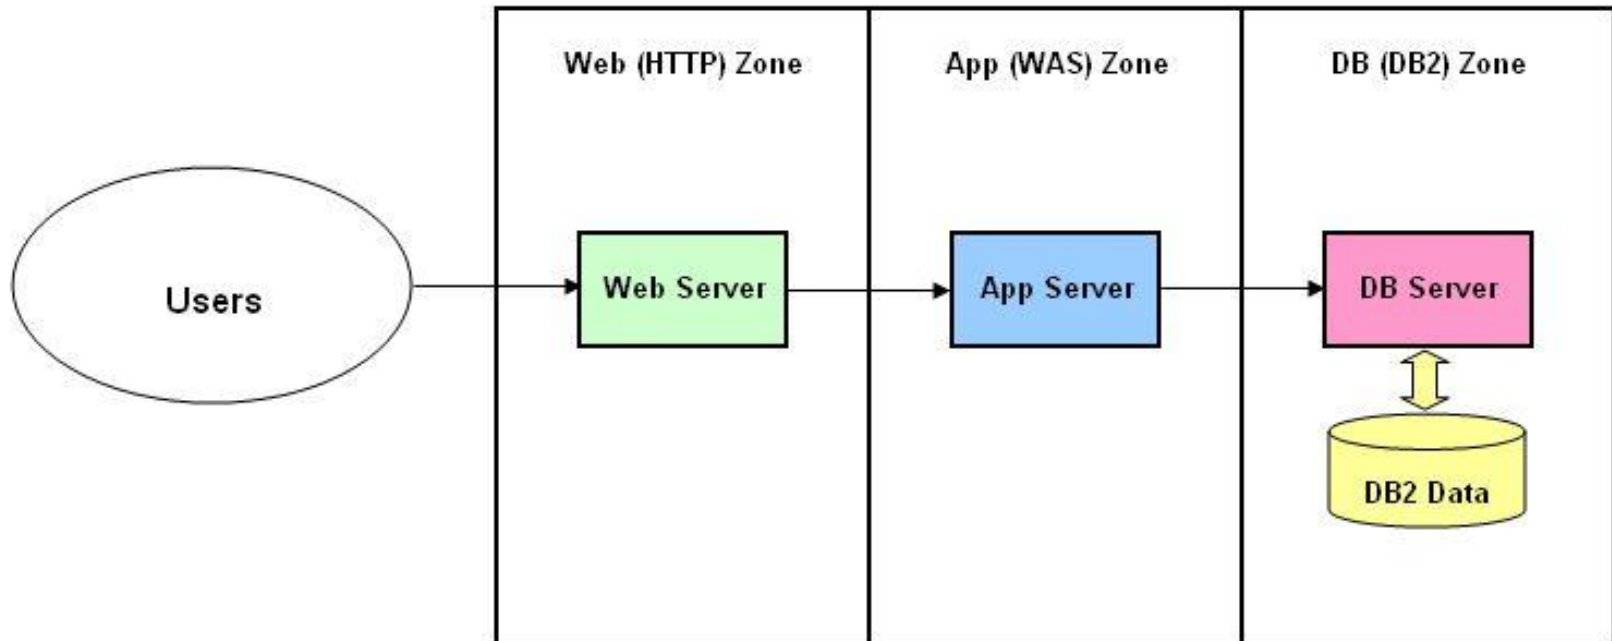
3390-9 (10016 Cyls)		
1b0	/boot	60 Cyls
1b1	/	500 Cyls
1b2	Swap	VDisk
1b3	Swap	VDisk
1b4	Swap	VDisk
1b5	/local	100 Cyls
1b6	/var	1019 Cyls
1b7	/usr	2290 Cyls
1b8	/opt	489 Cyls
1b9	/tmp	5557 Cyls

DataMart Architecture

- Adopted a 3-tiered Zone Architecture
 - Web (Front-end), App (Middle) and DB (Back-end) zones
 - Determined DataMart servers should reside in back-end DB zone (Info Server and DB Server)
- DataMart Process
 - DataMart Info Server Extracts, Transforms & Loads (ETL) records from DB Servers into the DataMart DB2 Server
 - (AIX) WAS-based App middle server accesses DataMart DB2 server to build reports
 - Executives access reports from HTTP Web server which accesses the WAS-based App server

3-Tier Zone Architecture

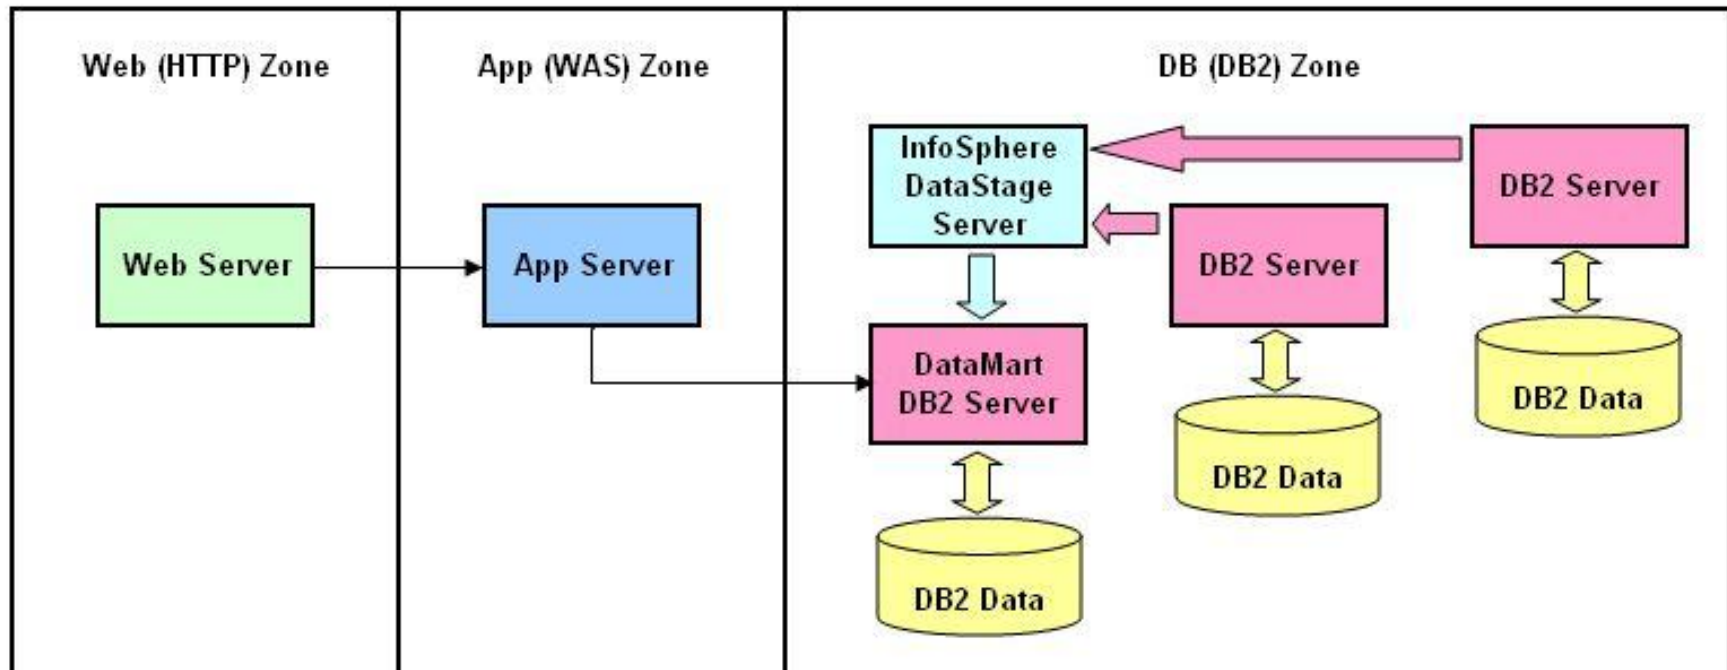
- At first, thought that the IBM InfoSphere DataStage server would go in the App Zone...



Zones Are Separated By Firewalls

3-Tier DataMart Architecture

- IBM InfoSphere DataStage Server was relocated in DB Zone for better performance



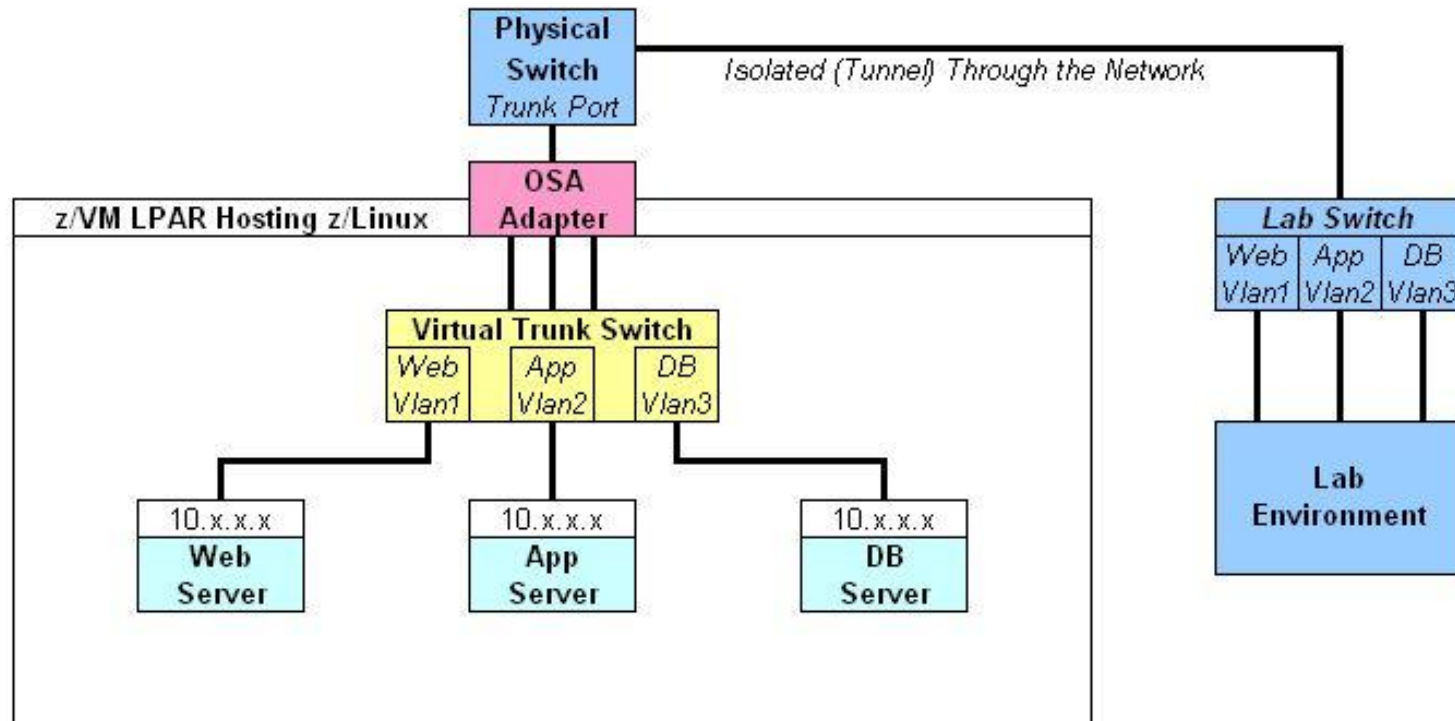
3-Tier Zone Structure For DataMart

Lab Environment

- Lab environment is completely isolation from rest of network via a tunnel
- Used isolated lab environment for two major tasks:
 1. Establish server baseline for application servers
 2. Build prototype DataMart servers (w/o baseline)
- Once baseline image is ready, use it for cloning
- Used baseline image to clone new DataMart Servers

Lab Environment (Tunnel)

- Tunnel connects Lab servers to the Lab environment
- Using VLANs to maintain separation between 3 zones.




Establishing the Baseline App Server

- Initially developed SLES 9 prototype baseline servers in Lab
 - SLES 9 was based on a monolithic R/W root file system
 - Discovered IBM InfoServer required SLES 10 (+ bug fixes)
- Built new servers using SLES 10 SP2
 - Server image based on a 3390-9 DASD volume
 - SLES 10 servers were also configured to allow for a read-only root file system

3390-3 (3339 Cyls)

100	/	3038 Cyls
101	Swap	VDisk
102	Swap	300 Cyls



3390-9 (10016 Cyls)

1b0	/boot	60 Cyls
1b1	/	500 Cyls
1b2	Swap	VDisk
1b3	Swap	VDisk
1b4	Swap	VDisk
1b5	/local	100 Cyls
1b6	/var	1019 Cyls
1b7	/usr	2290 Cyls
1b8	/opt	489 Cyls
1b9	/tmp	5557 Cyls

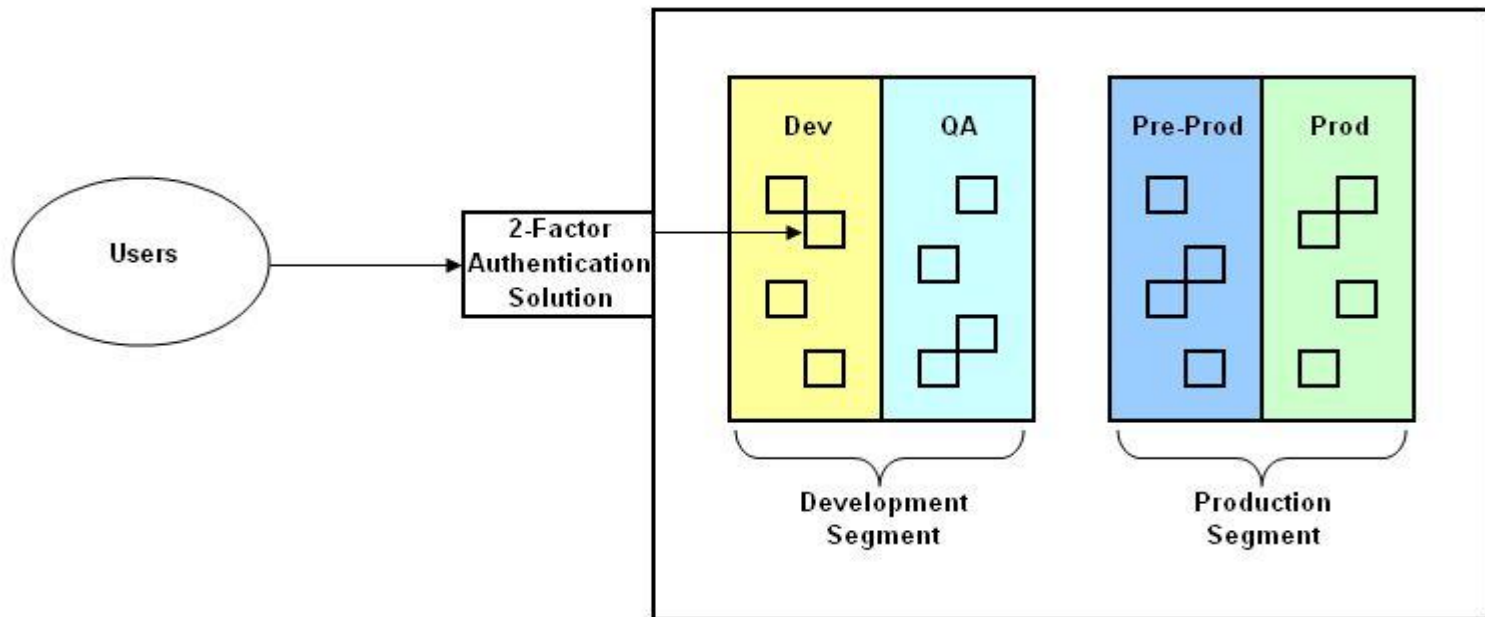
Server Security Considerations

- Established a new security hardened SLES baseline
- Based on security standards used for other (Unix) platforms, including:
 - HIDS (Host Intrusion Detection Software) requirement
 - Failed – working on solution with chosen vendor (Now Solved!)
 - Two-factor authentication, i.e. SecurID requirement
 - Failed – working on solution with vendor (Now Solved!)
 - Single sign-on management software requirement - OK
 - Configuration Management (CM) software requirement – OK
 - Meanwhile, failures are mitigated through external solution



External Security Solution

- 2-Factor authentication must be performed before the servers can be accessed
- Auditing is enabled and Syslog facilities monitored for intrusion activity



Finalizing the Baseline Server & Cloning

- Security Baseline Validation
 - Subjected to a vulnerability scan in the Lab
 - Issues were identified and resolved (patches applied, etc.)
- Deployed baseline server outside of Lab environment
 - Placed on development network segment
 - Confirmed connectivity with infrastructure servers
 - Finalized server baseline image for cloning

Finalizing the Baseline Server

- Final baseline server image takes a whole 3390-9 DASD volume
- Extra space is needed for some of the baseline software, i.e. HIDS, CM & auditing

R/W Server Configuration


1b0	/boot	R/O	60 Cyls
1b1	/	R/O	500 Cyls
1b2	Swap	VDisk	64 MB
1b3	Swap	VDisk	64 MB
1b4	Swap	VDisk	512 MB
1b5	/local	RAW	100 Cyls
1b6	/var	RAW	1019 Cyls
1b7	/usr	R/O	2290 Cyls
1b8	/opt	R/O	489 Cyls
1b9	/tmp	RAW	5557 Cyls

Total Cylinders: 10015 100%

3390-9 (10016 Cyls)		
1b0	/boot	60 Cyls
1b1	/	500 Cyls
1b2	Swap	VDisk
1b3	Swap	VDisk
1b4	Swap	VDisk
1b5	/local	100 Cyls
1b6	/var	1019 Cyls
1b7	/usr	2290 Cyls
1b8	/opt	489 Cyls
1b9	/tmp	5557 Cyls

Cloning the Baseline Server Image

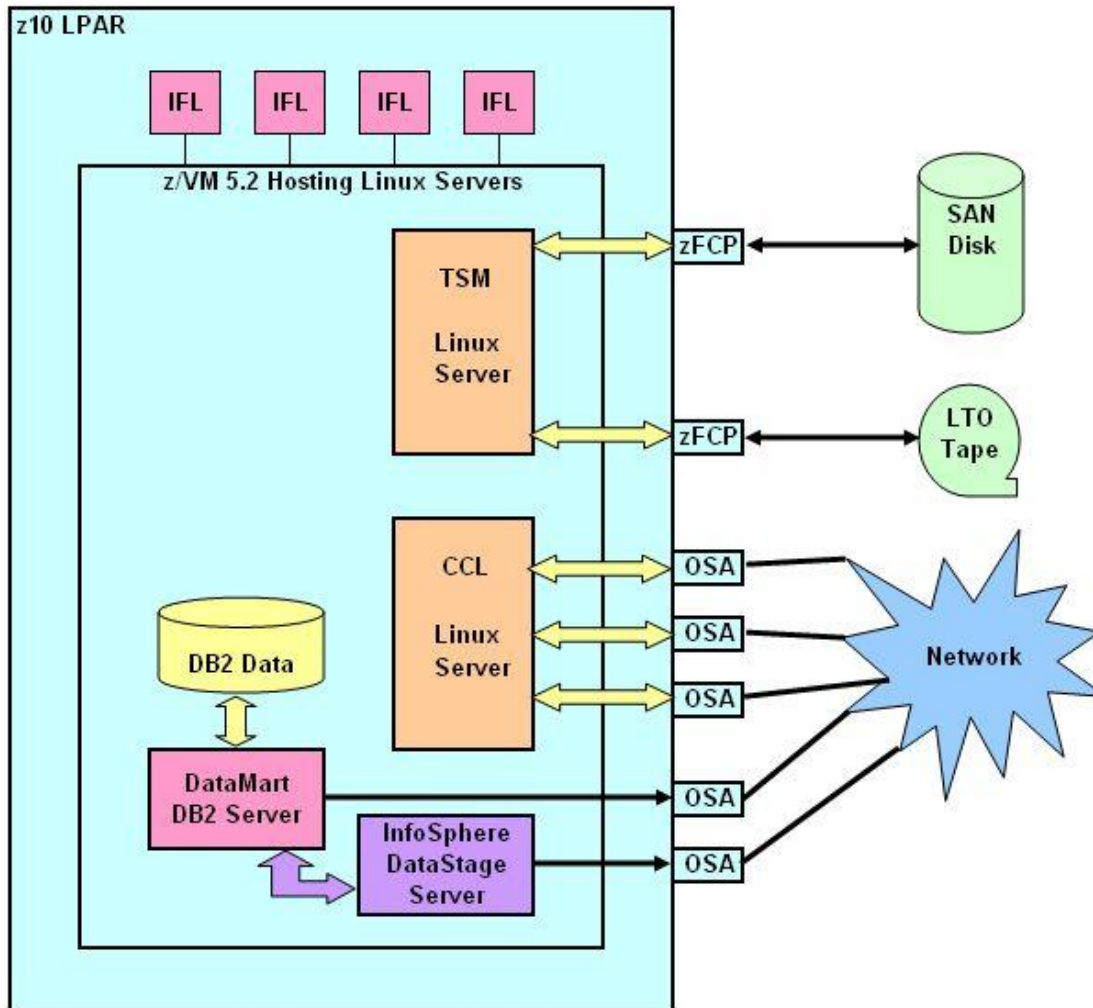
- Cloning uses the Flashcopy Service (when available)
- After cloning, server is IPLed & network reconfigured
- Initially, we are working with a R/W root file system
- Once the baseline image with a R/O root file system is finalized, then will start cloning R/O baseline servers



DataMart Server Configuration

- Two servers were cloned for DataMart from the baseline image and customized:
- Cloned the baseline image for the IBM InfoServer
 - Installed IBM Information Server Version 8.1
 - Server has 2 CPs, 5 GB Storage, 3390 DASD, no SAN space
 - Also established directory structure for (4) DataMart projects
- Cloned the baseline image for the DataMart DB Server
 - Installed DB2 Version 9.5 on DB Server
 - Server has 2 CP's, 2 GB Storage, 3390 DASD + SAN space

z/VM Environment With DataMart



- DataMart Servers were added to z/VM Hosted LPAR with existing infrastructure Linux servers
- Employ z/VM controls as needed – QUICKDSP, SET SHARE

Deployed DataMart in 4 Environments

Deployed 1st pair of servers in Development environment

- Began ETL testing with a few thousand records
- Initial testing of DataMart report code

Deployed 2nd pair of servers in QA environment

- Increased ETL testing to millions of records
- Further testing of DataMart report code

Deployed 3rd pair of servers in Pre-Prod environment

- Increased ETL testing to 10's of millions of records
- Pre-tested using full set of records
- Final testing/confirmation of DataMart report code

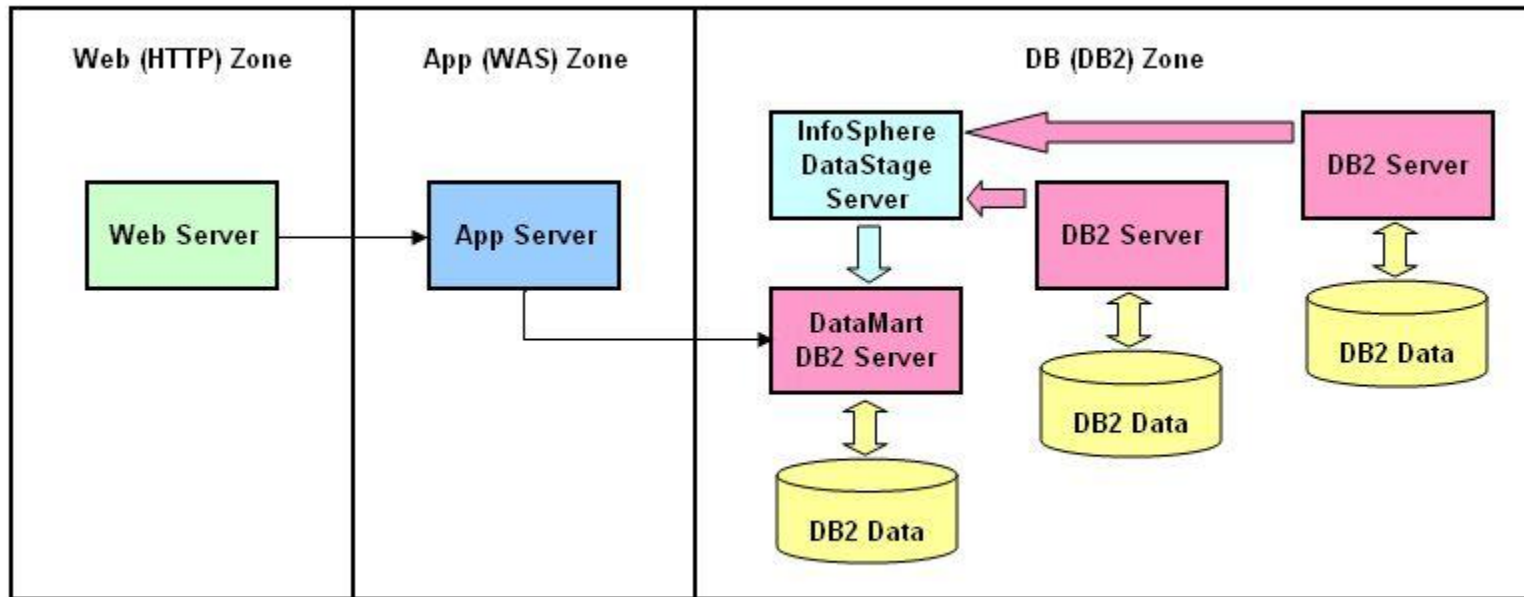
Deploying DataMart in 4 Environments

Deployed 4th pair of servers in Production environment

- Using full set of records
- Records coming from multiple DB sources
- Over 70 million records currently loaded
- Growth rate estimated at over 10 million records annually
- Performing Daily/Weekly/Monthly DB ETLs & Updates

Deployed Configuration

- In this DataMart implementation, DB zone is on z Series
- Web & App zones are on mid-tier (AIX) servers



3-Tier Zone Structure For DataMart

How is the DataMart Doing?

- Executives are very excited about the DataMart results
- Performance
 - CPU peaks during ETL processing, up to 60-70% (4 IFLs)
 - Storage over commit ratio is 2:1 – so far, so good
- Additional projects are being added to original project
 - Planning to make the DataMart available to a wider audience
- Growing pains
 - As DataMart projects are added, demand for resources increases
 - Already see the need for more CPU, storage and SAN space!

Were the goals met?

- We proved the concept – we now have Linux Application Servers running on z Series
- We establish the z/VM infrastructure for running SLES-based Virtual Linux Application Servers
- We developed a secured SLES server image for cloning
 - However, still resolving a couple of issues (Resolved now!)
- We establish roles & responsibilities for running & maintaining Linux App Servers on z Series
- We deployed a DataMart for executive review

Next Steps

- Planning to upgrade Info Server & DB2 versions
 - Developing a server upgrade & roll-out strategy
- Upgrading z/VM to Version 6.1
 - Planning to employ XRC-based DASD mirroring
 - Employing new SSL Server for secure access to z/VM
 - Employing Terminal Server for centralized access to servers
- Creating another Linux LPAR for Production Linux servers
 - Positioning for a VM Cluster configuration in the future (6.2)
- Preparing to add another project that completed a proof of concept in the Lab environment

Thank You For Attending!

Questions?

Hillgang Fall Meeting