



Ensemble Enabling z/VM V6.2 and Linux for System z

Alan Altmark

Senior Managing z/VM and Linux Consultant

Alan Altmark@us.ibm.com



Trademarks & Disclaimer



The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DB2, e-business logo, ESCON, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/390, System Storage, System z9, VM/ESA, VSE/ESA, WebSphere, xSeries, z/OS, zSeries, z/VM.

The following are trademarks or registered trademarks of other companies

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries. LINUX is a registered trademark of Linux Torvalds in the United States and other countries. UNIX is a registered trademark of The Open Group in the United States and other countries. Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation. SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC. Intel is a registered trademark of Intel Corporation. * All other products may be trademarks or registered trademarks of their respective companies.

NOTES: Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply. All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions. This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography. References in this document to IBM products or services do not imply that IBM intends to make them available in every country. Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use. The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice. Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.



Agenda

- Hardware components of an ensemble
- z/VM Ensemble Components
- Virtual Switch Controllers
- DIRMAINT authorizations
- Enable SMAPI Servers
- Validating the Enablement
- Linux Ensemble Considerations





What is a zEnterprise Ensemble

- A zEnterprise node is a single zCEC with 0 to 4 zBX frames and up to two blade centers per frame
- A zEnterprise Ensemble is a collection of 1 to 8 zEnterprise Nodes managed as a single virtualized pool of server resources



- A zEnterprise node can be a member of a single ensemble
- An ensemble is the management scope for the Unified Resource Manager
- A primary / alternate pair of HMCs provide the management console for the ensemble





Enhancements for zEnterprise Ensembles

- Supported SLES and RHEL distributions
 - Optional Guest Platform Management Provider (GPMP)
 - IEDN/INMN (OSX/OSM) NIC support
 - Legacy NIC connection to IEDN or INMN via virtual switch ok via OSDSIM support in VSWITCH
- z/VM V6
 - z/VM Management Guest to forward Linux GPMP data to Unified Resource Manager
 - z/VM SMAPI enhancements
 - z/VM Directory Maintenance server (or equivalent)
 - INMN and IEDN virtual switch controllers
 - Control point for MAC assignment and VLAN access





Enhancements for zEnterprise Ensembles

- INMN and IEDN access provided via new z/VM virtual switch types
 - Uplink is the z/VM Management Guest
 - Automatic connection to INMN
 - Ensemble membership sets ensemble-defined MAC for each IEDN NIC
- SMAPI updates SYSTEM CONFIG
- z/VM is authoritative source of virtual machine state
 - State automatically reflected in Unified Resource Manager



z/VM System Management API Infrastructure Changes



- New SMAPI servers:
 - Support for IPv6 (INMN is IPv6)
 - Resiliency and error recovery (aka "guard" functions)
 - Management Guest, instantiated by the Unified Resource Manager
- New Systems Management APIs





z/VM SMAPI Family

VSMGUARD Is responsible for starting and monitoring remaining

SMAPI servers and management guest.

VSMPROXY Speaks to the Support Element

VSMREQIN Requests from IPv4 clients

VSMREQIU Requests from other guests using AF_IUCV sockets

VSMREQIM Requests from Unified Resource Manager

VSMREQI6 Requests from IPv6 clients



VSMEVSRV Gathers data from *VMEVENT and *LOGREC

system services



ZVMLXAPP INMN gateway. Automatically started by

Unified Resource Manager







z/VM SMAPI Family

VSMWORK1 Request server for short transactions

VSMWORK2 Request server for long-running transactions

VSMWORK3 Request server for long-running transactions

LOHCOST Caching server for Query-type operations



DTCSMAPI Private TCP/IP stack for SMAPI components

that require IP connectivity



PERSMAPI Performance monitor used if managing z/VM

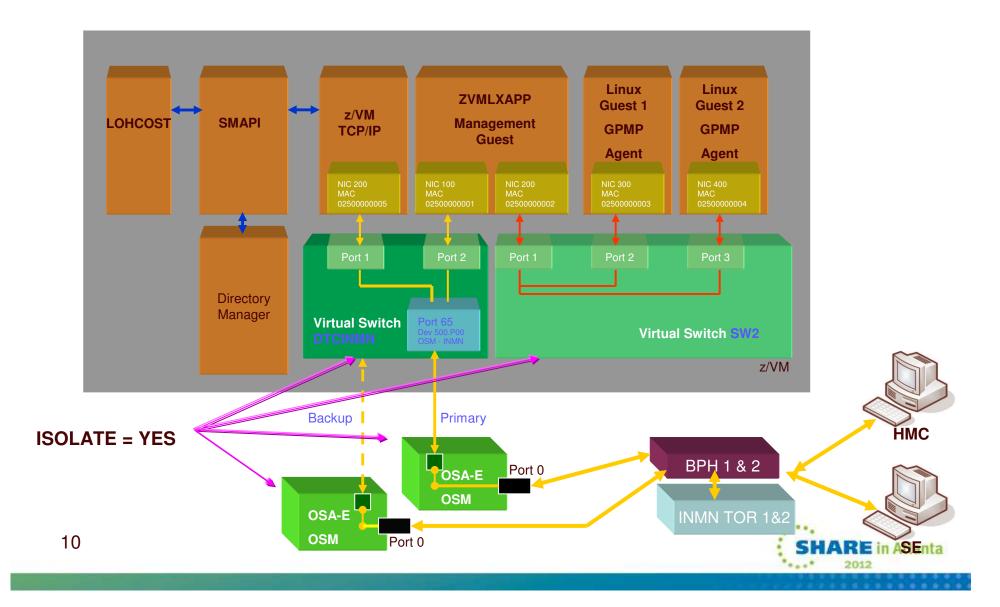
exclusively by Unified Resource Manager





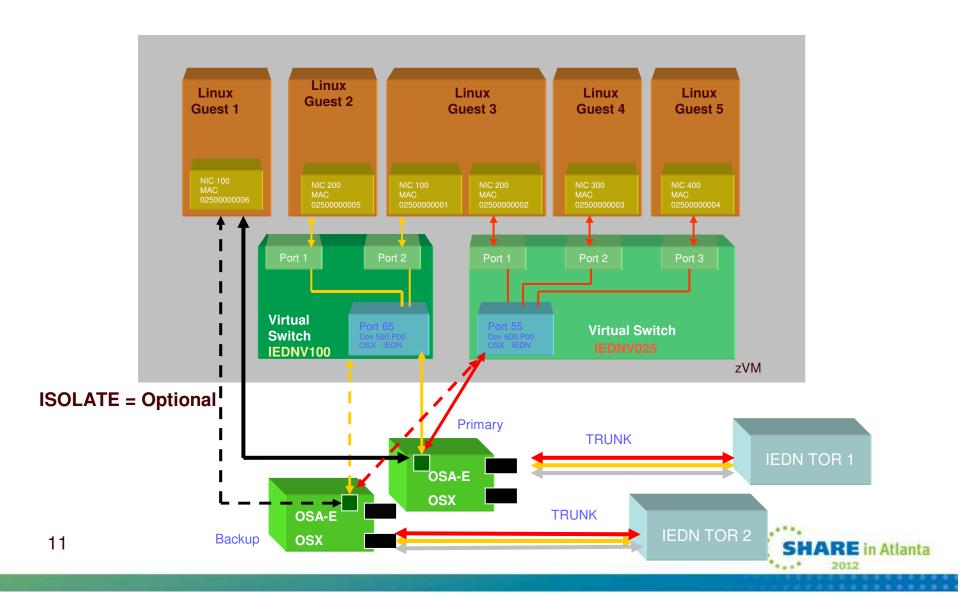


z/VM Ensemble INMN Infrastructure





z/VM Ensemble IEDN Infrastructure







- Software
 - z/VM 6.2 at Service Level 1101 or higher
 - PTFs UM33623 (VM65083), UM33547 (VM65055), and PTF for VM65151
 - Get the latest APAR information from
 - http://www.vm.ibm.com/service/vmrequrm.html

Hardware

- System z firmware bundle 41z or higher
- OSX and OSM CHPIDs genned and cabled
 - If not cabled, you will see error codes E080







- Books
 - CP Planning and Administration Guide, Chapter 15
 - z/VM Systems Management Application Programming, Chapter 4





Validate OSX/OSM devices are available

Q OSA TYPE ENSEMBLE

OSA	2300	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX
OSA	2301	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX
OSA	2302	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX
OSA	2303	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX
OSA	2304	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX
OSA	2305	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX
OSA	2306	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX
OSA	2307	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX
OSA	2308	OFFLINE	DEVTYPE	IEDN	CHPID	18	OSX





Validate OSX/OSM devices are available

Continued ...

OSA	2340	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM
OSA	2341	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM
OSA	2342	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM
OSA	2343	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM
OSA	2344	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM
OSA	2345	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM
OSA	2346	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM
OSA	2347	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM
OSA	2348	OFFLINE	DEVTYPE	INMN	CHPID	0A	OSM





Configure DIRMAINT Authorizations

- Command permission
- Surrogate permission





DIRMAINT Authorization

- Give SMAPI worker virtual machines permission to issue privileged DIRMAINT commands
- Update AUTHFOR CONTROL file:

```
ALL VSMWORK1 * 140A ADGHMOPS
ALL VSMWORK1 * 150A ADGHMOPS
ALL VSMWORK2 * 140A ADGHMOPS
ALL VSMWORK2 * 150A ADGHMOPS
ALL VSMWORK3 * 140A ADGHMOPS
ALL VSMWORK3 * 150A ADGHMOPS
ALL VSMGUARD * 140A ADGHMOPS
ALL VSMGUARD * 150A ADGHMOPS
```





DIRMAINT CONFIGXX DATADVH Additions

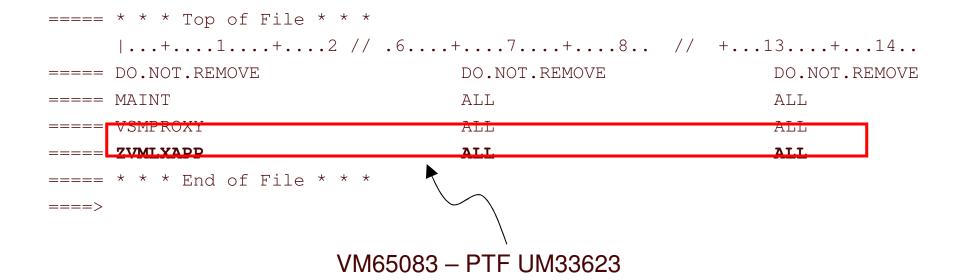
- Allow SMAPI worker virtual machines to issue requests on behalf of already-authenticated SMAPI clients
- Update CONFIGxx DATADVH file:

```
ALLOW_ASUSER_NOPASS_FROM= VSMWORK1 *
ALLOW_ASUSER_NOPASS_FROM= VSMWORK2 *
ALLOW_ASUSER_NOPASS_FROM= VSMWORK3 *
ALLOW ASUSER NOPASS FROM= VSMGUARD *
```





Authorize the Management Guest



VSMWORK1 AUTHLIST in VMSYS:VSMWORK1.

Note column numbers: 1, 66, 130





VSMGUARD ADMIN authority in DMSPARMS

On VMSERVS 191 minidisk

```
00000 * * * Top of File * * *

00001 ADMIN MAINT 6VMTCP10 VSMGUARD

00002 NOBACKUP

00003 SAVESEGID CMSFILES

00004 FILEPOOLID VMSYS

00005 USERS 100

00006 * * * End of File * * *
```





How to Operate This New Infrastructure?

- To start the SMAPI servers, XAUTOLOG VSMGUARD
 - Add it to the PROFILE EXEC of AUTOLOG1 or AUTOLOG2
- VSMGUARD will start the SMAPI servers and the Management Guest will start automatically.
- ZVMLXAPP can be restarted via the Unified Resource Manager as a task of the z/VM Hypervisor



Validating the configuration



```
q vmlan
VMLAN maintenance level:
 Latest Service: VM64780
VMLAN MAC address assignment:
 System MAC Protection: OFF
 MACADDR Prefix: 020000 USER Prefix: 020000
 MACIDRANGE SYSTEM: 000001-FFFFFF
            USER:
                    000000-000000
VMLAN Unified Resource Manager status:
 Hupervisor Access: YES Status: MANAGED
  ID: 52BD737254BF11E0B85A0010184CB262
 MAC Prefix: 023C90
VMLAN default accounting status:
 SYSTEM Accounting: OFF USER Accounting: OFF
VMLAN general activity:
 PERSISTENT Limit: INFINITE Current: 5
 TRANSIENT Limit: INFINITE
                              Current: 0
Ready; T=0.01/0.01 16:19:45
```





Validating the INMN configuration

```
g vswitch dtcinmn
VSWITCH SYSTEM DTCINMN Type: INMN
                                     Connected: 2
                                                      Maxconn: INFINITE
  PERSISTENT RESTRICTED
                            ETHERNET
                                                      Accounting: OFF
 VLAN Unaware
  MAC address: 02-3C-90-00-00-01 MAC Protection: Unspecified
  State: Ready
  IPTimeout: 5
                       QueueStorage: 8
  Isolation Status: ON
Uplink Port:
  RDEV: 236D.P00 VDEV: 236D Controller: DTCENS1
  RDEV: 234D.P00 VDEV: 234D Controller: DTCENS1
                                                BACKUP
```





Validating the SW2 configuration

```
vswitch sw2
VSWITCH SYSTEM SW2
                       Type: INMN
                                     Connected: 1
                                                     Maxconn: INFINITE
 PERSISTENT RESTRICTED
                           ETHERNET
                                                     Accounting: OFF
 VLAN Unaware
 MAC address: 02-3C-90-00-00-03 MAC Protection: Unspecified
 State: Ready
 IPTimeout: 5
                      QueueStorage: 8
 Isolation Status: ON
Uplink Port:
 NIC: ZVMLXAPP VDEV: 0200
```





Validating the DTCENSx controllers

```
q controller
Controller DTCVSW2 Available: YES VDEV Range: * Level 610
 Capability: IP ETHERNET VLAN_ARP GVRP LINKAGG ISOLATION
           NO_ENSEMBLE NO_INMN
  SYSTEM VSWITCH1 Primary Controller: VDEV: 2100
Controller DTCVSW1 Available: YES VDEV Range: * Level 610
 Capability: IP ETHERNET VLAN_ARP GVRP LINKAGG ISOLATION
           NO_ENSEMBLE NO_INMN
  SYSTEM VSWITCH1 Backup Controller: VDEV: 2120
Controller DTCENS1 Available: YES VDEV Range: * Level 610
 Capability: IP ETHERNET VLAN_ARP GVRP LINKAGG ISOLATION
           ENSEMBLE INMN
  SYSTEM DTCINMN Primary Controller: DTCENS1 VDEV: 236D
  SYSTEM DTCINMN Backup Controller: DTCENS1 VDEV: 234D
Controller DTCENS2 Available: YES VDEV Range: * Level 610
 Capability: IP ETHERNET VLAN_ARP GVRP LINKAGG ISOLATION
           ENSEMBLE
                     NO INMN
```





Validating your SFS configuration

```
q auth vmsys:vsmwork1.
 Directory = VMSYS: VSMWORK1.
 Grantee R W
              NR NW
 MAINT X X
              X
                 X
 VSMWORK1 X X X X
 VSMGUARD X X X X
 VSMPROXY X - X -
 VSMREQIM X - X -
 VSMREQIN X - X -
 VSMREQIU X - X -
 VSMREQI6 X - X -
 VSMWORK2 X - X -
 VSMWORK3 X - X -
```





Validating your SFS configuration

```
q auth vmsys:vsmwork1.data.
Directory = VMSYS:VSMWORK1.DATA
```

```
Grantee
      R
          M
             NR NW
MATNT
       X
         X
             X
                X
VSMWORK1 X X
            X = X
VSMGUARD X X X X
VSMPROXY X X X X
VSMREQIM X X X X
VSMREQIN X X X X
VSMREOIU X X X X
VSMREQI6 X X X X
VSMWORK2 X X X X
VSMWORK3 X X X X
```





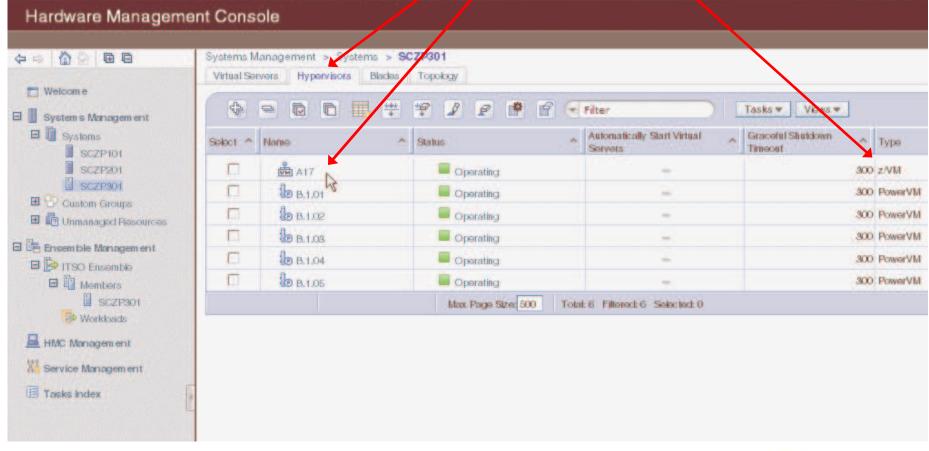
Validating your SMAP configuration

```
netstat
VM TCP/IP Netstat Level 610 TCP/IP Server Name: TCPIP
Active IPv4 Transmission Blocks:
User Id Conn Local Socket
                                 Foreign Socket
                                                         State
INTCLIEN 1006 *..TELNET
                                                        Listen
INTCLIEN 1007 9.12.4.189..TELNET 9.76.158.39..50358 Established
VSMREQIN 1002 *..44444
                                                        Listen
VSMPROXY 1003 *..55555 ________
                                                        Listen
Active IPv6 Transmission Blocks:
User Id Conn State
VSMREQI6 1001 Listen
 Local Socket: *..44445
 Foreign Socket: *..*
```





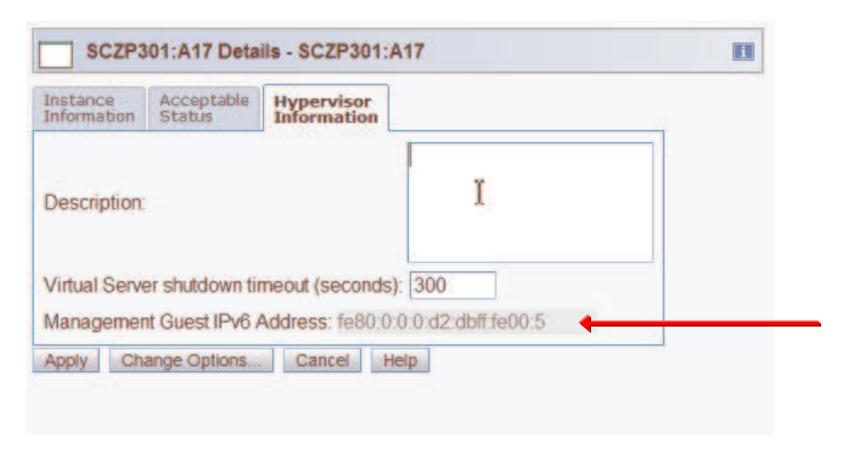
Validating your configuration







Validating your configuration



 The IPV6 IP address will display when the management guest is activated and z/VM is part of the Ensemble



Implementation Tips

- Access to the IEDN and OSX connections must be configured using Unified Resource Manager
 - Ensemble will reject "out-of-band" connection attempts
- VMSYS filepool needs to be backed up with the rest of your system!
 - This is where access rights and status are kept
 - If you lose it, you will start over
- No miracles. Unified Resource Manager doesn't solve connectivity problems.
 - E.g. FCP devices must be able to access LUNs without zManager if they are going to be able to do it with zManager)
- Console output from VSMGUARD, VSMWORK1, and VSMREQIU can hold clues if you have trouble.
 - But don't believe everything you see





Implementation Tips

- If you vary all devices offline in the SYSTEM CONFIG and then vary on only the ones you know about, zManager-defined FCP devices could be a problem.
 - You may want to have a predefined range of devices for this
- Unified Resource Manager is not a RACF security administration application.
 - Enable DIRMAINT-RACF Connector (USE_RACF=YES)
 - VSWITCH and RDEV authorizations must be handled separately
- If ZVMLXAPP does not start, the other SMAPI service machines will not be started.
- Depending on the size and volume of the virtual server directories to be managed, you may find that the SMAPI servers will run out of memory.
 - The default is 128MB. You can increase up to 512MB.





Next Steps: Use Unified Resource Manager

- Create IEDN Virtual Switches and give guests access to the IEDN
- Define disk storage resources
 - System and user
- Define virtual server containers for Linux guests or migrate existing guest
- Manage guest resources





Managing guest priorities from zManager

- "Too many cooks spoil the broth."
 - Only one resource manager at a time
 - If you are managing a guest with VMRM don't add it to a managed workload in zManager
 - Enable one or the other, not both





Ensemble Enabling Linux on System z Guests





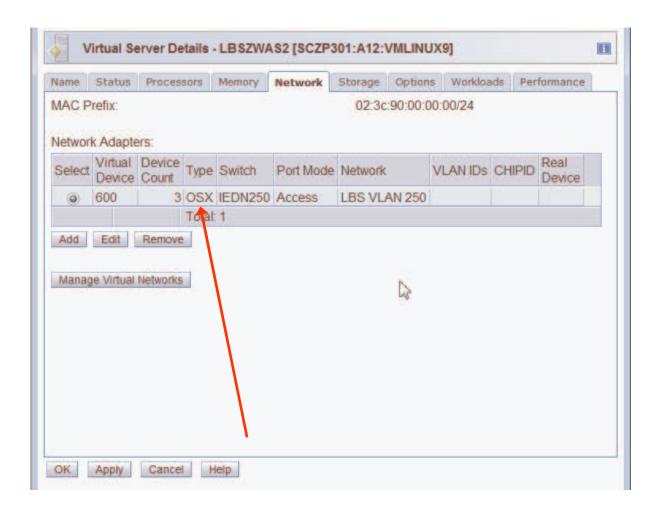
Linux considerations for Ensembles

- Native OSX interfaces are supported by the more recent kernel levels
 - RHEL 6.1, SLES 10 SP3, SLES 11 SP1
 - No Linux access to virtual or real OSX without kernel support
- VSWITCH can simulate OSD for older systems
 - OSDSIM support when you can't get to the latest kernel level
- Do not attempt to set a MAC address in the guest
 - z/VM will not allow (MACPROTECT)
 - The ensemble could assign a different MAC next time













Installer Boot With OSX Interface

- The SLES 11 SP1
 Installer System will not find an OSX interface
- Install with OSDSIM interface and add or convert after kernel upgrade

```
15:27:00 3) NFS
15:27:00 4) SMB / CIFS (Windows Share)
15:27:00 5) TFTP
15:27:00
15:27:00 > 15:27:02 1
15:27:02 Detecting and loading network drivers
15:27:02 netiucv.8db02b: driver initialized
15:27:03
15:27:03 Choose the network device.
15:27:03 1) IBM IUCV
15:27:03 2) IBM IUCV
15:27:03 3) IBM IUCV
15:27:03 3) IBM IUCV
```





Installer boot with OSD simulation

If the Network
 Adapter is
 redefined as OSD
 instead of OSX,
 the OSA devices
 are discovered

```
15:29:47 1
15:29:47 Detecting and loading network drivers
15:29:47 netiucv.8db02b: driver initialized
15:29:48
15:29:48 Choose the network device.
15:29:48
15:29:48 1) IBM OSA Express Network card (0.0.0600)
15:29:48 2) IBM OSA Express Network card (0.0.0601)
15:29:48 3) IBM OSA Express Network card (0.0.0602)
15:29:48 4) IBM IUCV
15:29:48 5) IBM IUCV
15:29:48 6) IBM IUCV
```





Layer 2 MAC Address

- The layer 2 MAC address can be automatically recorded in the /etc/sysconfig/network scripts when the interface is configured
- The virtual MAC assigned to the guest by the Unified Resource Manager may change
- Remove the LLADDR entry from your IEDN interfaces
- An update to Linux should be available to correct this behavior





Configuration with MAC address specified

```
16:00:40 cat ifcfg-eth0
16:00:41 B00TPR0T0='static'
16:00:41 IPADDR='172.27.250.7/24'
16:00:41 BR0ADCAST='172.27.250.255'
16:00:41 STARTMODE='onboot'
16:00:41 LLADDR='02:3c:90:00:00:0e'
16:00:41 NAME='OSA Express Network card (0.0.0600)'
16:00:41 lbxzwas1:/etc/sysconfig/network #
```





Remove LLADDR

```
cd /etc/sysconfig/network
```

cp ifcfg-eth0 backup-ifcfg-eth0

sed '/LLADDR/d' backup-ifcfg-eth0 > ifcfg-eth0







```
16:05:04 cat ifcfg-eth0
16:05:05 B00TPROTO='static'
16:05:05 IPADDR='172.27.250.7/24'
16:05:05 BROADCAST='172.27.250.255'
16:05:05 STARTMODE='onboot'
16:05:05 NAME='OSA Express Network card (0.0.0600)'
16:05:05 lbxzwas1:/etc/sysconfig/network #
```



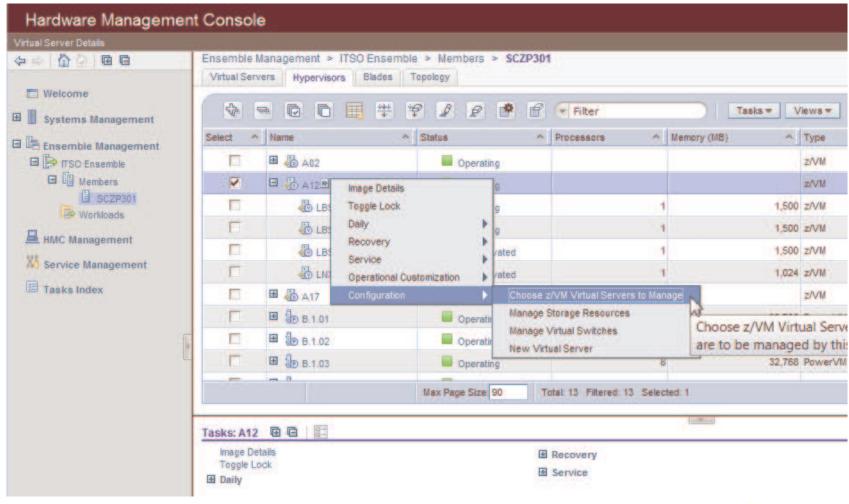


- You could either create new guest containers and copy or point them at existing disk storage
- Or you can migrate them directly to be "Managed" by the Unified Resource Manager
- You do not have to manage all guests. You may chose to just manage the ones you want.



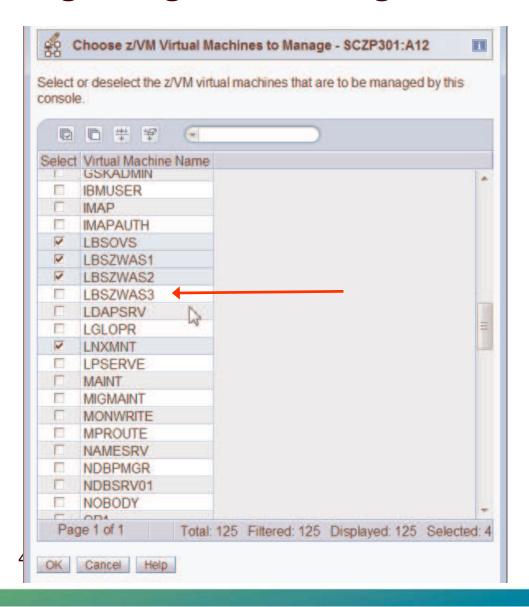














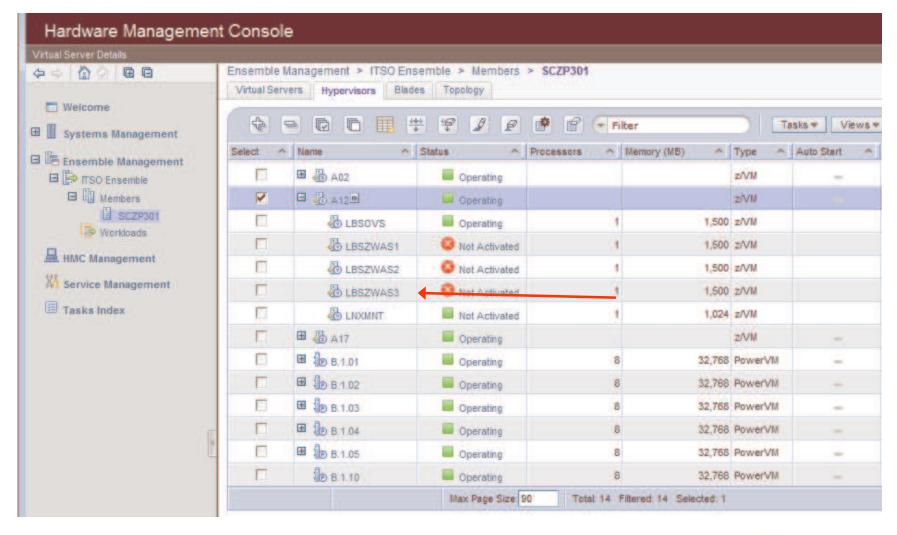


console	nsole.					
Select	Virtual Machine Name					
	GSKADMIN		1			
	IBMUSER					
	IMAP					
	IMAPAUTH					
V	LBSOVS					
V	LBSZWAS1					
V	LBSZWAS2					
V	LBSZWAS3		h			
	LDAPSRV					
	LGLOPR					
V	LNXMNT					
	LPSERVE					
	MAINT					
	MIGMAINT					
	MONWRITE					
	MPROUTE					
Г	NAMESRV					
	NDBPMGR					
	NDBSRV01					
	NOBODY		,			





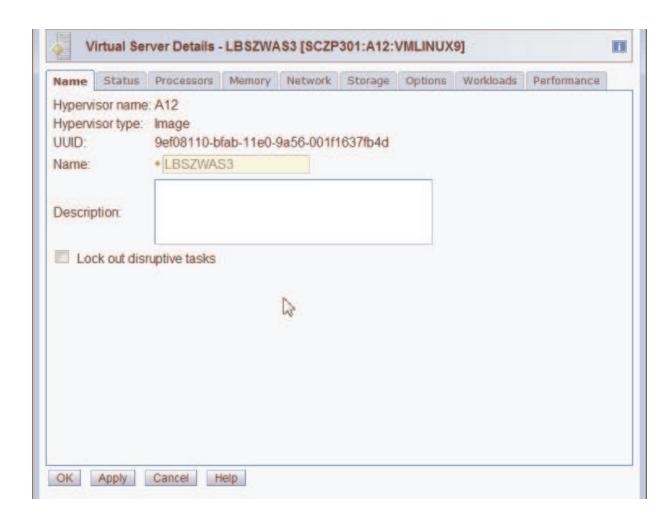








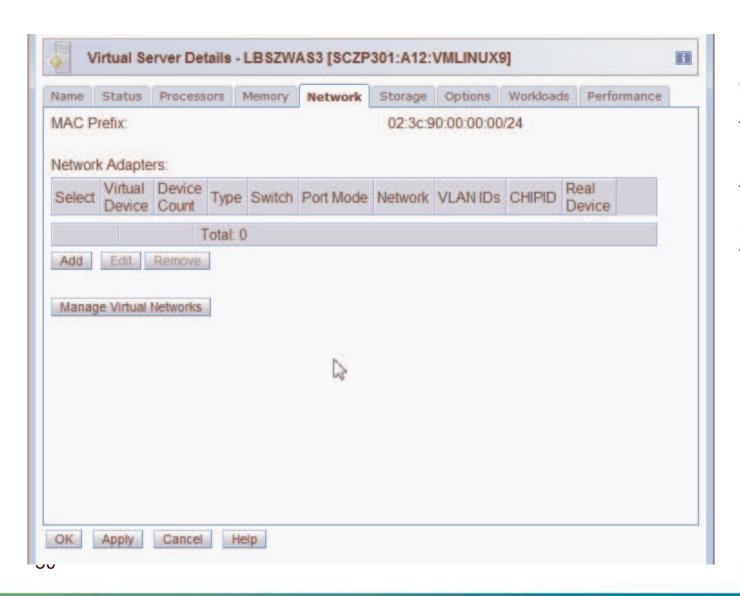




Once a guest is known to the Unified Resource Manager as a Virtual Server, you can change its configuration from the Unified Resource Manager



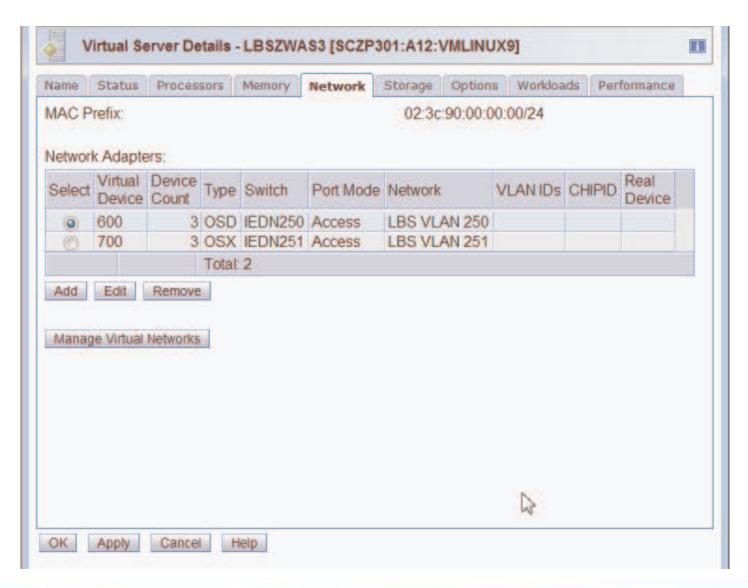




When migrated this guest had no NICs defined to it, so we will add two of them via the Unified Resource Manager











The Guest Platform Management Provider

- Provides detailed performance data from the guest operating system to the Unified Resource Manager
 - Is provided via the zManager code stream
- Data flows across the INMN
- Optional, but is required if you want to feed Application Response Measurement data (ARM) from middleware to the zManager
- Lets take a look at the steps to enable it







Password:	NOPASS		Options Workloads Perfo	MINISTERNIA PROPERTY AND ADDRESS OF THE PROPERTY AND ADDRE
Privilege class				
PL device:	Add RMC Device			
PL load parar PL parameter Enable Gf GPMP version	resource monitoring this virtual server. Device	e Guest Platform Man and control network d * 1000 3 RMC	agement Provider, a evice will be added to	







```
Port Name

Options

[ ] Enable IPA Takeover

[x] Enable Layer 2 Support
Layer2 MAC Address

Read Channel

Write Channel

0.0.1000

0.0.1002
```

Remember we do not want to code a MAC address in the configuration files. The Ensemble could assign a different virtual MAC the next time this guest is started. Communications would fail if we tried to assign a different MAC







```
Network Card Setup

General—Address—Hardware—
Device Type
QETH—
() No IP Address (for Bonding Devices)
(x) Dynamic Address

Zeroconf——â DHCP version 6 only——â

() Statically assigned IP Address

IF Address

Alias Name IP Address Netmask
```

- Utilize "Zeroconf" to dynamically assign the IP address
- The assignment is made by the Ensemble





Validating INMN on Linux for System z

```
Link encap:Ethernet HWaddr 02:D2:DB:00:00:37
inet addr:169.254.250.130 Bcast:169.254.255.255 Mask:255.255.0.0
inet6 addr: fe80::d2:dbff:fe00:37/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1492 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:11 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:0 (0.0 b) TX bytes:678 (678.0 b)
```





Validating INMN on Linux for System z

nstance nformation	Acceptable Status	Hypervisor Information	
Description:			
/irtual Serv	er shutdown t	imeout (seconds): 300	





Validating INMN on Linux for System z

```
wass1:~ # ping6 -I eth3 fe80:0:0:0:d2:dbff:fe00:17
PING fe80:0:0:0:d2:dbff:fe00:17(fe80::d2:dbff:fe00:17) from fe80::d2:dbff:fe00:37 eth3:
64 bytes from fe80::d2:dbff:fe00:17: icmp_seq=1 ttl=64 time=8.28 ms
64 bytes from fe80::d2:dbff:fe00:17: icmp_seq=2 ttl=64 time=0.084 ms
64 bytes from fe80::d2:dbff:fe00:17: icmp_seq=3 ttl=64 time=0.096 ms
64 bytes from fe80::d2:dbff:fe00:17: icmp_seq=4 ttl=64 time=0.092 ms
64 bytes from fe80::d2:dbff:fe00:17: icmp_seq=5 ttl=64 time=0.113 ms
64 bytes from fe80::d2:dbff:fe00:17: icmp_seq=6 ttl=64 time=0.094 ms
^C
--- fe80:0:0:0:d2:dbff:fe00:17 ping statistics ---
6 packets transmitted, 6 received, 0% packet loss, time 4998ms
rtt min/avg/max/mdev = 0.084/1.460/8.282/3.050 ms
```

- Here we ping the IPV6 address of the management guest (ZVMLXAPP) from the previous slide
- The INMN interface is also coded on this ping







 The installation of the GPMP rpm is performed by retrieving the code from the Management Guest







```
chown ibmgpmp:ibmgpmp /opt/ibm/gpmp/CollectFFDC.sh /opt/ibm/gpmp/armsad /opt/ibm/gpm
k /opt/ibm/qpmp/qpmpmain /opt/ibm/qpmp/qpmpsad /opt/ibm/qpmp/qpmpshm /opt/ibm/qpmp/
tall-config /opt/ibm/gpmp/post-uninstall /opt/ibm/gpmp/shmdump
chmod 550 /opt/ibm/gpmp/CollectFFDC.sh /opt/ibm/gpmp/armsad /opt/ibm/gpmp/gpasetuid
pmp/qpmpmain /opt/ibm/qpmp/qpmpsad /opt/ibm/qpmp/qpmpshm /opt/ibm/qpmp/java /opt/ibr
/opt/ibm/gpmp/post-uninstall /opt/ibm/gpmp/shmdump
chmod 555 /opt/ibm/gpmp/java
chown ibmgpmp:ibmgpmp /var/opt/ibm/gpmp
chmod 770 /var/opt/ibm/gpmp
chown ibmqpmp:ibmqpmp /opt/ibm/qpmp/java/arm4.jar
chmod 444 /opt/ibm/gpmp/java/arm4.jar
chown ibmlarm /usr/sbin/lsarm
chgrp ibmlarm /usr/sbin/lsarm
chmod 550 /usr/sbin/lsarm
chown ibmlarm:ibmlarm /var/opt/ibm/arm/
chmod 770 /var/opt/ibm/arm/
chown ibmlarm /opt/ibm/qpmp/qpmpshm
chgrp ibmlarm /opt/ibm/gpmp/gpmpshm
chmod 6550 /opt/ibm/gpmp/gpmpshm
chown root /opt/ibm/qpmp/qpasetuid
chmod 4550 /opt/ibm/qpmp/qpasetuid
Adding the gpmpcheck process to crontab for user ibmgpmp.
Note: Run /opt/ibm/qpmp/post-install-config to grant permissions
to other users to access GPMP and ARM components. Currently, only the
ibmorph user has all the necessary access permissions.
```



Enabling the GPMP on Linux for System z



```
wass1:~ # su ibmgpmp -c "/opt/ibm/gpmp/gpmp start"
FEW6101I The guest platform management provider is starting.
wass1:~ # su ibmgpmp -c "/opt/ibm/gpmp/gpmp autostart on"
FEW6002I Setting guest platform management provider autostart on.
wass1:~ #
```

- The gpmp is started with the ibmgpmp user name
 - You cannot start the GPMP under root!

 The second command will cause it to start automatically on subsequent IPLs



Enabling the GPMP on Linux for System z



```
2200
                                        00:00:00 /usr/sbin/cupsd
                      0 11:51 2
root
                                        00:00:00 /usr/sbin/nscd
          2209
root
                                        00:00:00 /usr/lib/postfix/master
          2271
root
          2299
                                        00:00:00 /usr/sbin/gdm --no-console
                      0 11:51 ?
root
          2308
                                        00:00:00 /usr/sbin/cron
                      0 11:51 ?
root
                                        00:00:00 /usr/sbin/xinetd -pidfile /var/run/xinet
          2323
                      0 11:51 ?
root
          2325
                                        00:00:00 [flush-253:1]
                      0 11:51 ?
root
          2332
                      0 11:52 ttyS0
                                        00:00:00 /sbin/mingetty --noclear /dev/ttyS0 dumb
root
          2414
                                        00:00:00 sshd: root@pts/0
                2192
                      0 12:41 ?
root
                2414
                      0 12:41 pts/0
          2417
                                        00:00:00 -bash
root
                                        00:00:00 [flush-94:0]
          2449
                      0 12:41 ?
root
                                        00:00:00 /sbin/autoip -B eth3
          4072
                      0 12:43 ?
root
                                        00:00:0/ pickup -1 -t fifo -u
                2271
postfix
          4544
                                        00:00 00 amgr -1 -t fifo -u
          4545
                2271
postfix
                                        00:00:00 [flush-253:0]
          4659
                      0 12:45 ?
                                        00/00:00 gpmpmain daemon
ibmopmp
          4668
                      0 12:46 pts/0
                                         0:00:00 ps -4
          4681
root
                2417
```

Can see that the autoip service and the gpmp daemon are running





```
wasq1:/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin # /opt/ibm/qpmp/post-install-config
GPMP configuration:
Default GPMP owner: ibmgpmp
Default GPMP group: ibmgpmp
Default ARM owner: ibmlarm
Default ARM group: ibmlarm
You can change these values by specifying different answers below.
In addition, you can add additional users to the above named groups.
Default user for GPMP processes? [ibmgpmp]
Default group for GPMP access? [ibmgpmp]
Users to be granted access to GPMP? (comma separated)
Default owner for Lightweight ARM shared memory? [ibmlarm]
Default group for access to Lightweight ARM data? [ibmlarm]
Users to be granted access to ARM data? (comma separated) [] root
About to change GPMP configuration:
Default GPMP owner: ibmgpmp
Default GPMP group: ibmgpmp
Additional users granted access to GPMP:
Default ARM owner: ibmlarm
Default ARM group: ibmlarm
Additional users granted access to ARM: root
Is this okay? [default is NO] yes
Proceeding with re-configuration.
Adding user root to group ibmlarm
Note: If any of the users you listed have active sessions,
they will need to log out and back in for the changes to take effect.
```



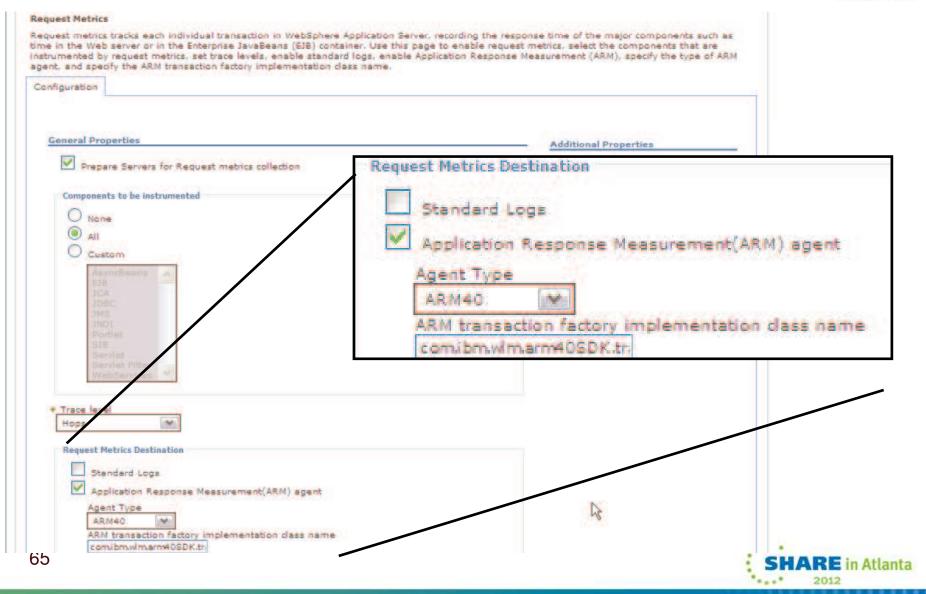




Add two custom properties to WAS JVM









```
case SPLATFORM in
 AIX)
   WAS LIBPATH="$WAS HOME"/bin
   NLSPATH=/usr/lib/nls/msg/%L/%N:/usr/lib/nls/msg/en US/%N:${}
 Linux)
   WAS LIBPATH="SWAS HOME"/bin:/usr/lib64
   NLSPATH=/usr/lib/locale/%L/LC MESSAGES/%N:S{NLSPATH:=}
   JAVA HIGH ZIPFDS=2
 SunOS)
```

Update WAS setupCmdLine.sh to add /usr/lib64 to WAS_LIBPATH





Start the WebSphere server

- Use Isarm to query the status of the ARM data
- You should now have ARM data flowing to the Unified Resource Manager





References

- z/VM CP Planning and Administration Guide
- z/VM CP Commands and Utilities Reference
- z/VM Directory Maintenance Facility Commands Reference

- IBM zEnterprise Ensemble Performance Management Guide
- IBM zEnterprise Ensemble Planning and Configuration Guide
- IBM zEnterprise Unified Resource Manager Redbook





Questions?



IBM zEnterprise System Technical Training Offerings



- Basic: IBM System z Technical Overview (2 days)
 - ES820/EZ820 IBM System z: Technical Overview of HW and SW Mainframe Evolution
- Basic: IBM zEnterprise System Technical Introduction (1 day)
 - ESA0/EZA0 Describes new terminology, functions and provides technical details for each
 of the main components that make up the zEnterprise System. It describes how the
 resources of the zEnterprise System provides the necessary infrastructure for hybrid
 computing that can be managed and virtualized as a single pool of resources.
- Advanced: IBM zEnterprise System: Using zManager to Provision Virtual Servers
 (4 days) NEW with hands on labs
 - ESA1 Lectures and hands-on labs provide the information and skills required to use the IBM zEnterprise Unified Resource Manager to provision virtual servers on blades and under z/VM. Students using hands-on labs will use Unified Resource Manager tasks to:
 - Audit an existing ensemble configuration, delete ensemble resources to both hypervisers and virtual servers
 - Define Virtual networks to the ensemble, add Storage resources manually and imported via an SAL, create virtual servers and associated resources.
 - Enable Guest Platform Management, create and start performance policies, run workload and review reports
- For complete technical training catalog, see www.ibm.com/training



Customized Services

IBM zEnterprise System Technical Training Offerings

Ensemble Acceptance Services

Customized Services

zEnterprise
Ensemble
Enablement
Jumpstart
Assistance for
zBX Blades

3-4 weeks

zEnterprise Ensemble Enablement for zBX Blades

(Starter Kit)

~9 weeks

Customized Services

zEnterprise Ensemble Enablement Jumpstart Assistance for

3-4 weeks

IBM Smart
Analytics
Optimizer
Enablement
Services

3-4 weeks

zEnterprise
Ensemble
Enablement
Jumpstart
Assistance
for
DataPower
XI50z Blades
1-2 weeks

Pre-Sales

Rapid Workload Optimization Assessment and TCO for IBM zEnterprise System

4 weeks



Session 10331

Alan C. Altmark

Senior Managing IT Consultant
IBM STG Lab Services

z/VM & Linux on System z

IBM

1701 North Street Endicott, NY 13760

Tel 607 429 3323 Fax 607 429 3323

Email: alan_altmark@us.ibm.com

