Ensemble Enabling z/VM and Linux for System z

Richard Young
IBM STG Lab Services

August 10th, 2011
Session Number 09879

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

- Quick review what an Ensemble is composed of
- z/VM Ensemble Components
- Virtual Switch Controllers
- DIRMAINT authorizations
- Enable SMAPI Servers
- Validating the Enablement
- Linux Ensemble Considerations

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

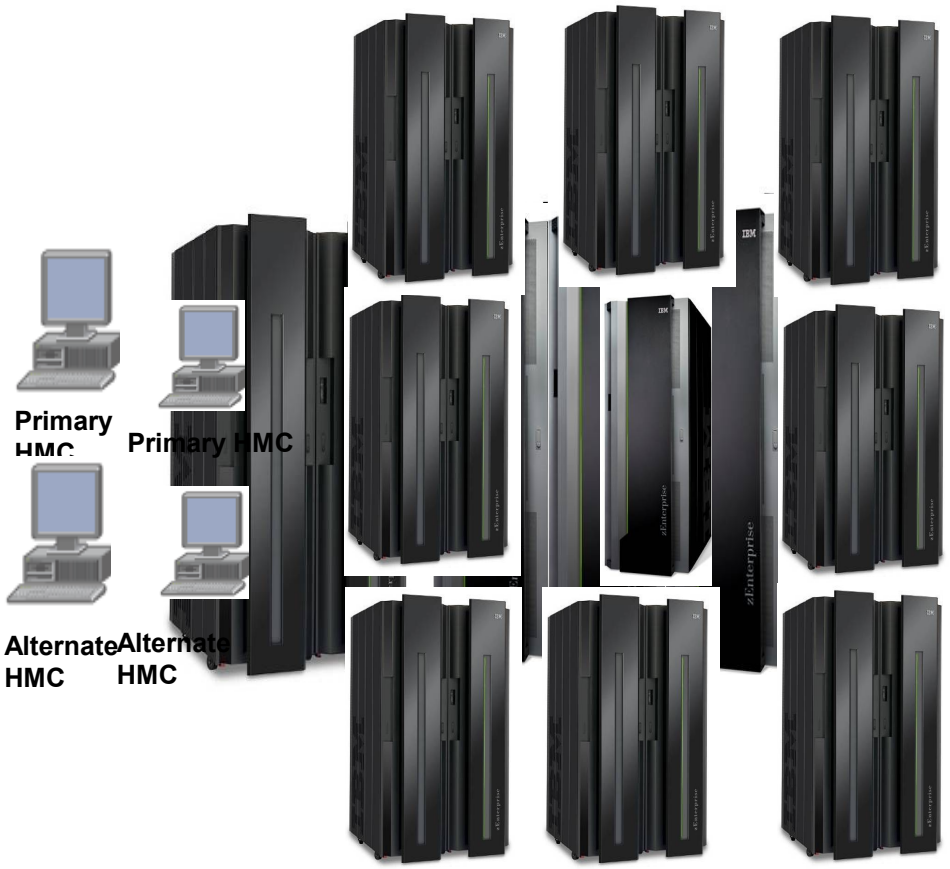

Enhancements for the Unified Resource Manager

- Software
  - Supported SLES and RHEL distributions
    - Optional - Guest Platform Management Provider
    - IEDN/INMN (OSX/OSM) NIC support
    - Legacy NIC connection to IEDN or INMN via virtual switch (OSDSIM support)
  - z/VM 6.1
    - z/VM Management Guest – ZVMLXAPP
    - 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 the Unified Resource Manager

- z/VM 6.1 Continued…
  - INMN and IEDN access provided via new z/VM virtual switch types
    - Up-link can be virtual machine NIC (for Management Guest)
    - Automatic connection to INMN
    - Ensemble membership enforce Ensemble MAC for each IEDN NIC
  - SMAPI validates and updates SYSTEM CONFIG
  - z/VM is authoritative source of virtual machine state
    - State automatically reflected in Unified Resource Manager

z/VM System Management APIs

- As part of the support for the IBM zEnterprise Unified Resource Manager, new SMAPI servers were also created:
  - AF_MGMT request server – Used to communicate between the SE and SMAPI
  - INET6 request severs – Use IPv6 to connect with clients
  - VSMGUARD worker server – Guard server to provide resiliency and error recovery
  - Management Guest (ZVMLXAPP) – Automatic instantiation by the Unified Resource Manager
- New Systems Management APIs added

z/VM SMAPI Servers

- **VSMGUARD** The VSMGUARD sever is a new worker server that provides better resiliency and error recovery. You start this server and it automatically start the remaining SMAPI servers and management guest. Unlike the worker servers, VSMGUARD does not process any request.

- **VSMREQIM** The VSMREQIM is a AF_MGMT request server. The AF_MGMT request server is used to communicate between the support element and the z/VM SMAPI server environment, only when z/VM is managed by the Unified Resource Manager. There can be one and only one AF_MGMT request server

- **VSMREQI6** VSMREQI6 is the AF_INET6 request server. This server handles requests over the IPV6 sockets

- **VSMREQIN** VSMREQIN is the AF_INET request server. This server handles request over the IPV4 sockets

- **VSMPROXY** VSMPROXY is the AF_SCLP request server. This server is used for communication between the support element and the z/VM SMAPI server environment. There can be one and only one AF_SCLP server

z/VM SMAPI Servers & Management Guest

- **VSMREQIU** VSMREQIU is the AF_IUCV request server. There can be one or more AF_IUCV request servers.

- **VSMWORK1** VSMWORK1 is the short call request server. It is one of the three default worker servers. There must always be at least one short call worker server. The default SFS directories are owned by the “short call” request server VSMWORK1.

- **VSMWORK2** VSMWORK2 is a long call request server. This is the one of two long call request servers. If all the request servers are busy, the request will be queued until on becomes available.

- **VSMWORK3** VSMWORK3 is a long call request server. This is the one of two long call request servers. If all the request servers are busy, the request will be queued until on becomes available.

- **ZVMLXAPP** ZVMLXAPP is a the new Management Guest. The Management Guest is automatically instantiated by the Unified Resource Manager. It is also part of the INMN network communication past to Linux guest.

z/VM Ensemble Management (INMN) Infrastructure

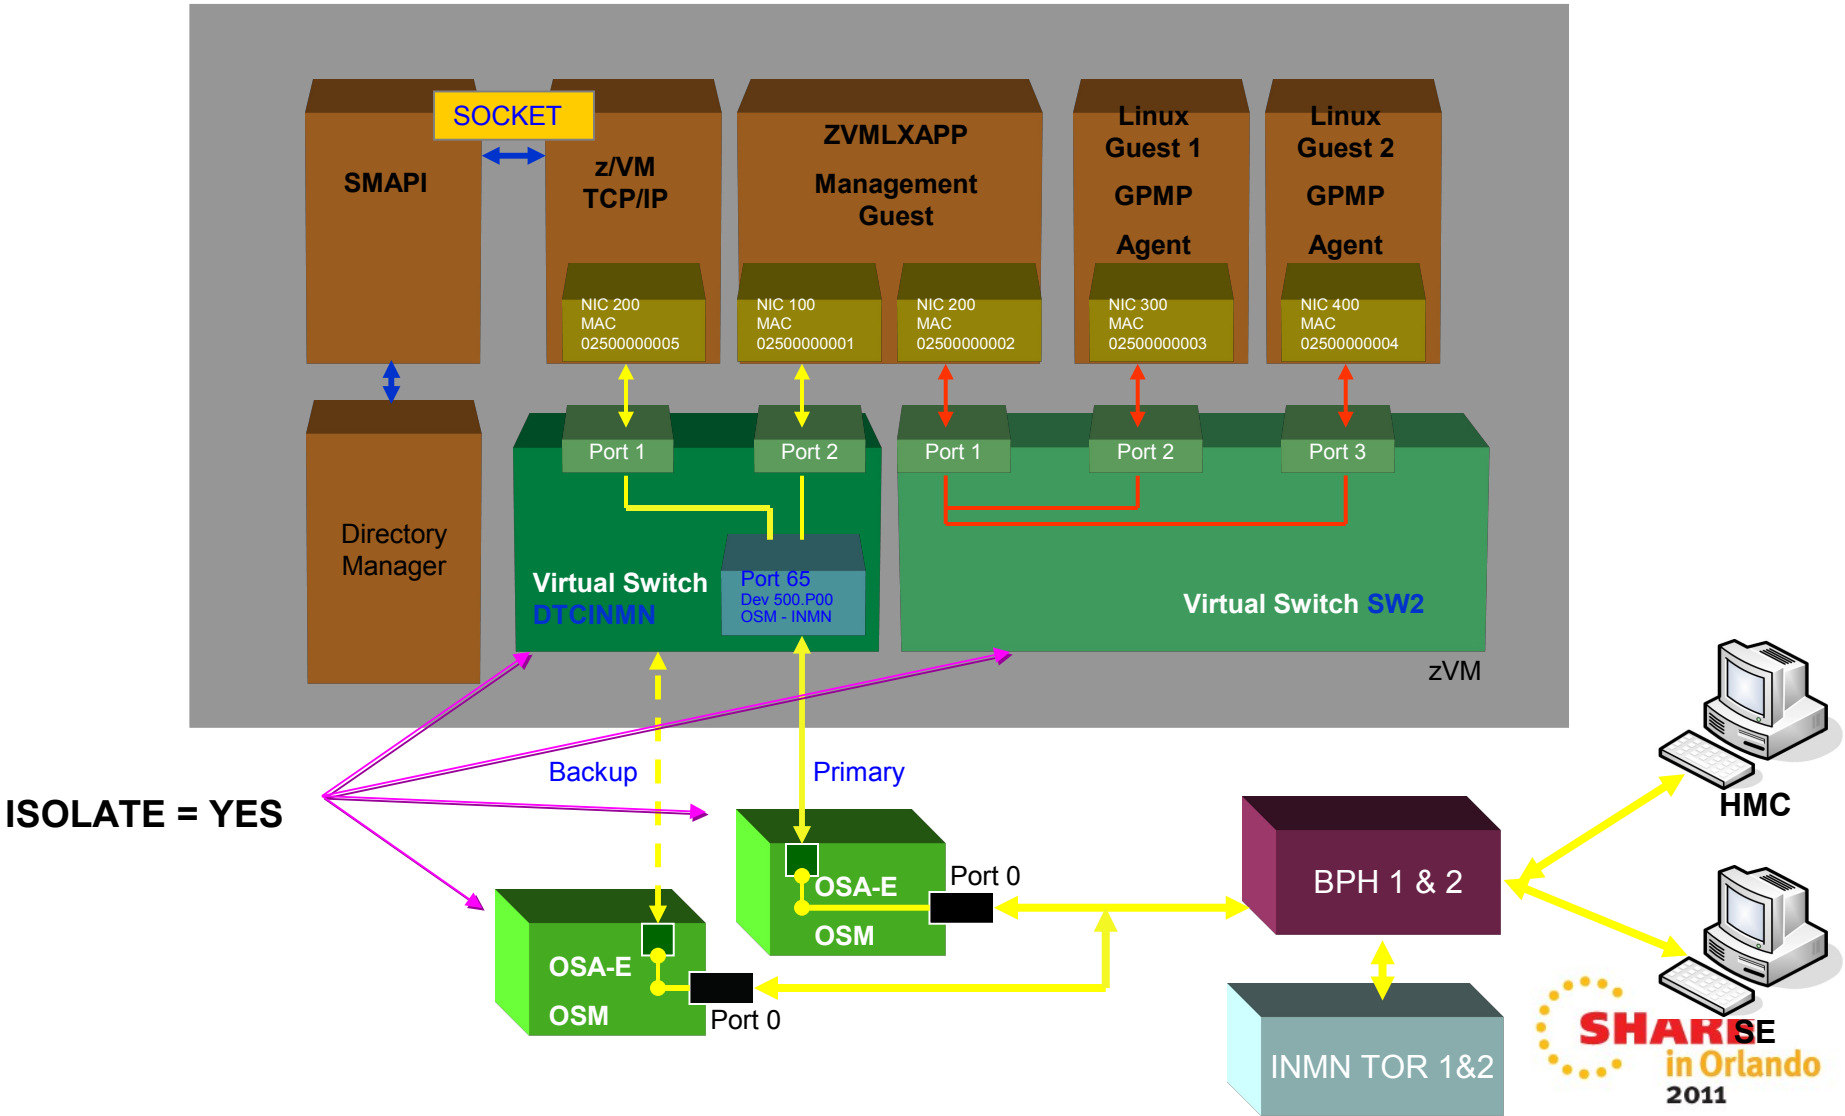

SHARE in Orlando 2011

z/VM Ensemble IEDN Infrastructure Options

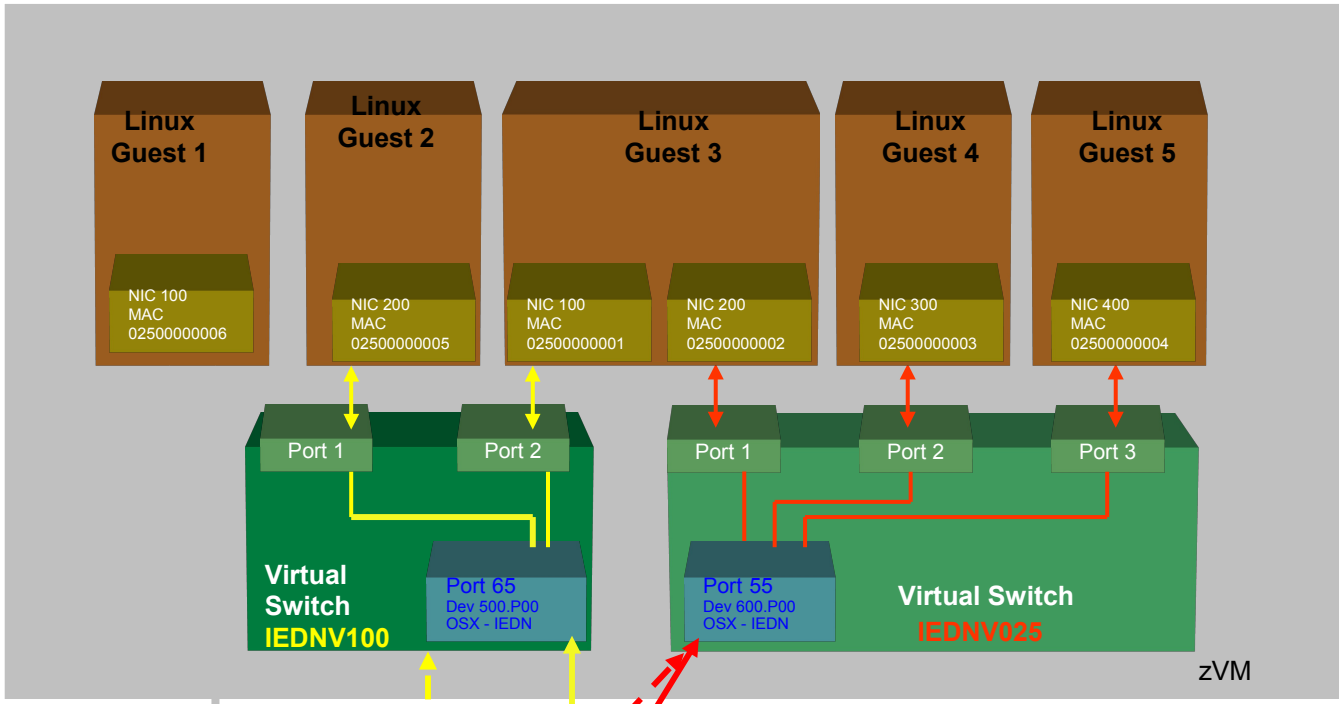

ISOLATE = Optional

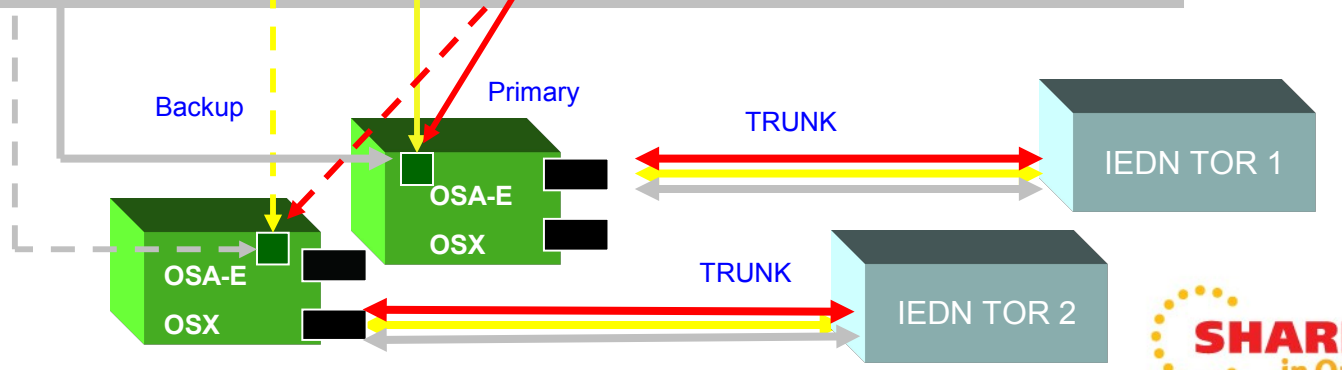

10 GbE Wrap Plugs

Required if no zBX!

Preparation for Enabling z/VM

- References
  - CP Planning and Admin Guide (SC24-6178-01) Chapter 16 for all the detailed installation steps
  - z/VM System Management Application Programming (SC24-6234-01)
- Software
  - zVM 6.1, current RSU, + VM64822, VM64904, VM64917, VM64956, VM64957
  - Check zVM prereq URL for the latest list
    - http://www.vm.ibm.com/service/vmrequrm.html
  - System z bundle 41z or higher
- Hardware
  - OSX and OSM CHPIDs genned and cabled

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
```

Setup IEDN & INMN VSWITCH Controllers

- Steps
  - Add directory entries
  - Format 191 “a” disks
  - Copy PROFILE EXEC from TCPMAINT
  - Define “server” configuration files on TCPMAINT for the controllers
  - Update obeyfile authorizations

DTCENS1 Directory Entry

```
00001 USER DTCENS1 NEWPASS 32M 128M BG
.
.
.
00019 MDISK 191 3390 161 5 LX4U1R MR READ
   WRITE MULTIPLE
```

DTCENS2 Directory Entry

```
00001 USER DTCENS2 NEWPASS 32M 128M G
.
.
.
00010 MDISK 191 3390 166 5 LX4U1R MR READ WRITE
   MULTIPLE
```

Configure DIRMAINT Authorizations

DIRMAINT AUTHFOR CONTROL Additions

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_ASUSER_NOPASS_FROM= VSMWORK1 *
ALLOW_ASUSER_NOPASS_FROM= VSMWORK2 *
ALLOW_ASUSER_NOPASS_FROM= VSMWORK3 *
ALLOW_ASUSER_NOPASS_FROM= VSMGUARD *

Enabling z/VM SMAPI and Manage Guest Server

- Ensemble_Port = “55555” needs to be added to DMSSICNF COPY file on MAINTS 193
- Add the VSMREQI6, VSMREQIM, ZVMMAPLX, and VSMGUARD directory entries
- Propagate sample PROFILE EXEC (VSMREQIN SAMPPROF) to the request servers (VSMREQIN, VSMREQI6, VSMREQIU, VSMPROXY, VSMREQIM)
- Propagate sample PROFILE EXEC (VSMWORK1 SAMPPROF) to Guard and Worker Servers (VSMGUARD, VSMWORK1, VSMWORK2, VSMWORK3)

- IMPORTANT! – You need to replace existing PROFILE EXECs with the updated version

ZVMMAPLX Directory Entry

```
00001 USER ZVMLXAPP NEWPASS 1024M 2048M G
00002 COMMAND SET D8ONECMD * OFF
00003 COMMAND SET RUN ON
00004 COMMAND TERM LINEND #
00005 CMD SET VSWITCH DTCINMN GRANT ZVMLXAPP OSDSIM ON
00006 CMD DEFINE VSWITCH SW2 TYPE INMN ETHERNET
00007 CMD SET VSWITCH SW2 GRANT ZVMLXAPP OSDSIM ON
00008 CMD SET VSWITCH SW2 UPLINK NIC ZVMLXAPP 200
00009 CMD DEFINE NIC 100 TYPE QDIO
00010 CMD DEFINE NIC 200 TYPE QDIO
00011 CMD COUPLE 100 TO SYSTEM DTCINMN
00012 CMD COUPLE 200 TO SYSTEM SW2
00013 COMMAND SPOOL CONS START *
.
.
.
00028 MDISK 0191 3390 2325 010 LX4W02 MR
```

Enabling z/VM … Next steps

- Install TCP DATA on A disk of AF_MGMT server (VSMREQIM)
- Update the TCP DATA to point to DTCENS1
- Authorize Management Guest (ZVMMAPLX) and VSMPROXY to perform all SMAPI functions

Authorize the Management Guest

1. Logon to the VSMWORK1 guest
2. Issue the following commands
3. #CP IPL CMS
4. acc (noprof
5. set filepool VMSYS
6. access VMSYS:VSMWORK1. B
7. xedit vsmwork1 authlist B

Authorize the Management Guest

```
===== * * * Top of File * * *
|...+....1....+....2....+....3....+....4....+....5....+....6....+....7...
===== DO.NOT.REMOVE
DO.NOT.RE
MOVE
===== MAINT                                          ALL
===== VSMPROXY                                ALL
===== ZVMLXAPP                                ALL
===== * * * End of File * * *
====>
```

Authorize the Management Guest

- 8. Repeat the VMSPROXY line and add ZVMLXAPP as shown
- 9. Issue the “file” subcommand to save the changes
- 10.Issue #CP IPL CMS to restart VSMWORK1
- 11.Issue #CP DISCONNECT
- Note: It is suggested to repeat an existing line in the file and alter the server name

Enabling z/VM … Next steps

- Update DMSSISVR NAMES as a local modification using the “automated local modification procedure”
- Issue enrolls and grants for VSMGUARD
- Authorize VSMGUARD as an ADMIN in DMSPARMS for VMSERVS 191
- Enroll VSMWORK and Request servers

z/VM Local Modification

- z/VM Guide for Automated Installation and Service - Appendix D
- Begin by ensuring the MAINT 512 disk is accessed as the D disk
- Next issue: **`localmod CMS DMSISVR NAMES`**
- Ensure all 13 entries shown in the CP Planning & Admin Guide Chapter 16 are merged in to the file

* Default AF_INET Server
* AF_INET6 Server
* Default AF_IUCV Server
* Default AF_SCLP Server
* Management Network Server
* Management Guest
* Guard Server
* Default Short Call Server
* Default Long Call Server 1
* Default Long Call Server 2
* Primary Vswitch Controller
* Backup Vswitch Controller
* Directory Manager

z/VM Local Modification

- After saving the changes
- Run: service CMS build
- When complete run: put2prod

Enroll & Grant VSMGUARD

- ENROLL USER VSMGUARD VMSYS:
- GRANT AUTHORITY VMSYS:VSMWORK1. TO VSMGUARD (WRITE NEWWRITE
- GRANT AUTHORITY VMSYS:VSMWORK1.DATA TO VSMGUARD (WRITE NEWWRITE
- GRANT AUTHORITY * * VMSYS:VSMWORK1.DATA TO VSMGUARD (WRITE
- GRANT AUTHORITY * * VMSYS:VSMWORK1. TO VSMGUARD (READ

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 * * *
```

Enrolling Request and Worker Servers

- enroll user vsmreqin vmsys:
- enroll user vsmreqi6 vmsys:
- enroll user vsmreqiu vmsys:
- enroll user vsmproxy vmsys:
- enroll user vsmreqim vmsys:
- enroll user vsmwork1 vmsys:
- enroll user vsmwork2 vmsys:
- enroll user vsmwork3 vmsys:

How to operate this new infrastructure?

- To start the SMAPI servers, XAUTOLOG VSMGUARD
- To automate, add it to the PROFILE EXEC of AUTOLOG1 or AUTOLOG2
- VSMGUARD will start the SMAPI servers and the Management Guest will start automatically.
- Do NOT add the new vswitch controllers to your AUTOLOGx or other automation, they will be automatically started when the management guest starts.
- ZVMLXAPP can be restarted via zManager

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:
  Hypervisor 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 configuration

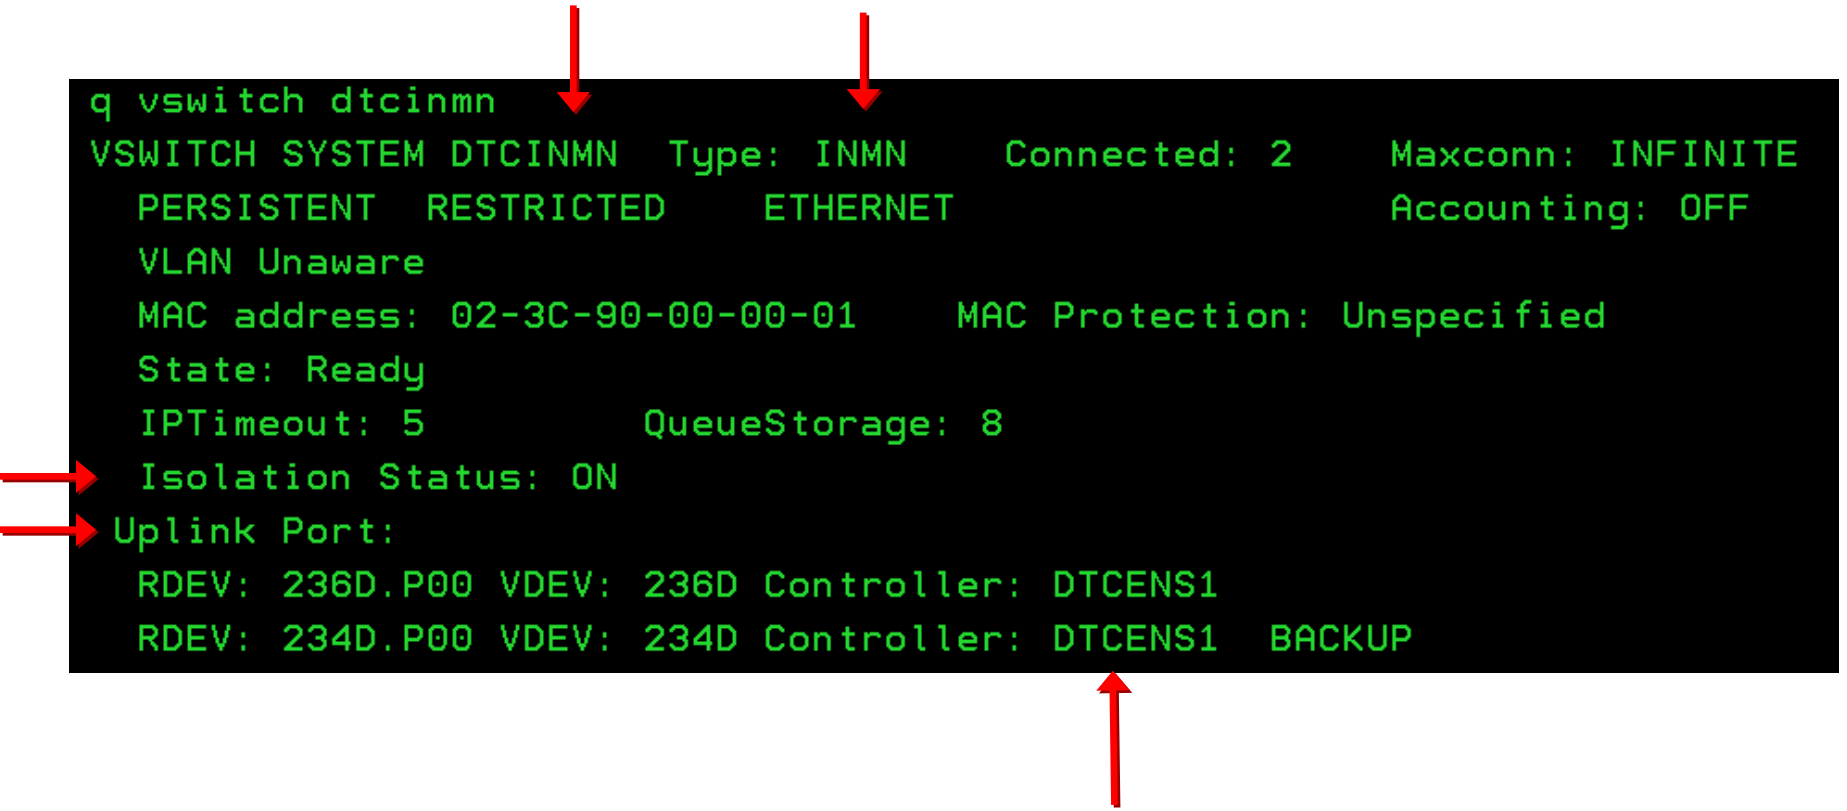

Validating the configuration

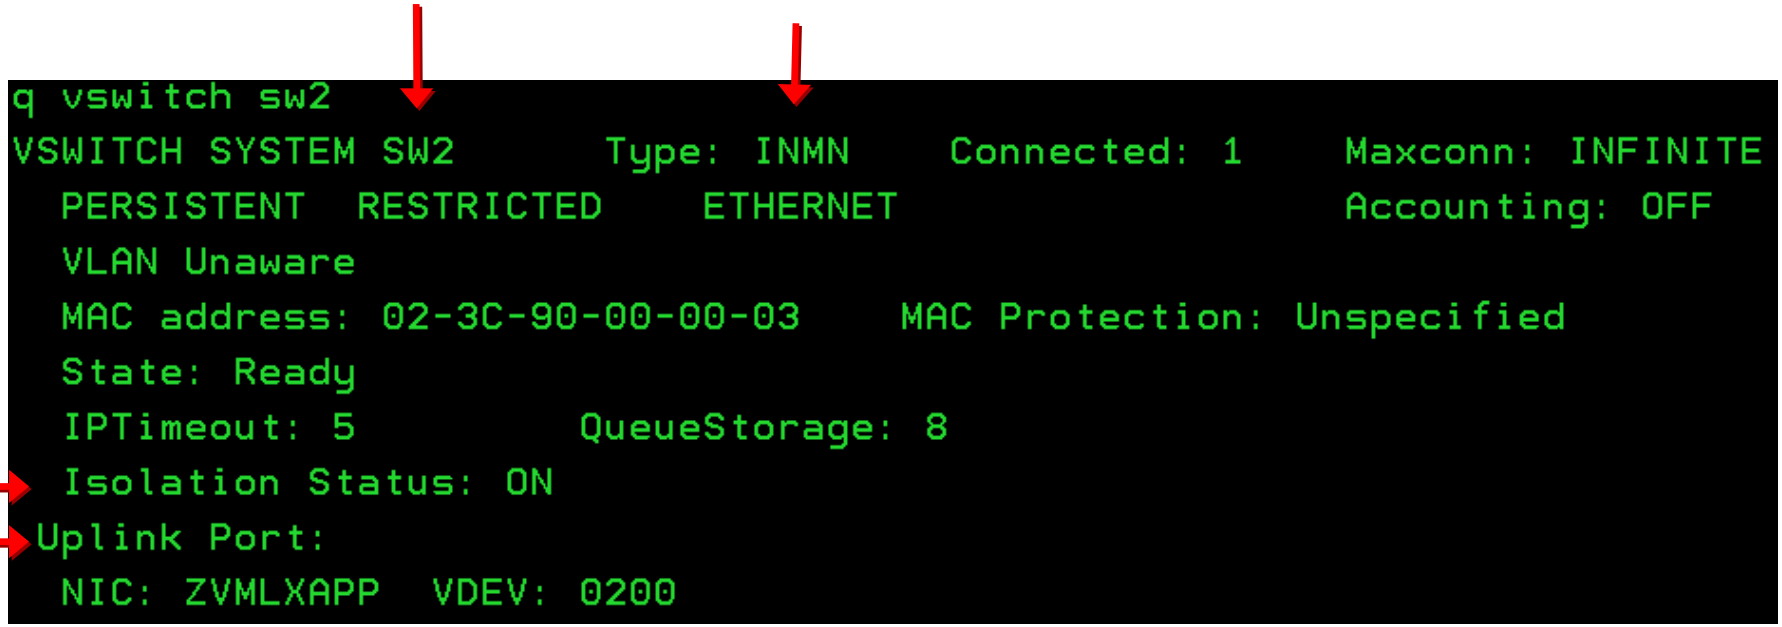

Validating the configuration

```
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: <list>    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: <list>    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 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 configuration

```
q auth vmsys:vsmwork1.data.
  Directory = VMSYS:VSMWORK1.DATA
  Grantee   R   W   NR  NW
  MAINT     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
  VSMREQIU  X   X   X   X
  VSMREQI6  X   X   X   X
  VSMWORK2  X   X   X   X
  VSMWORK3  X   X   X   X
```

Validating your 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

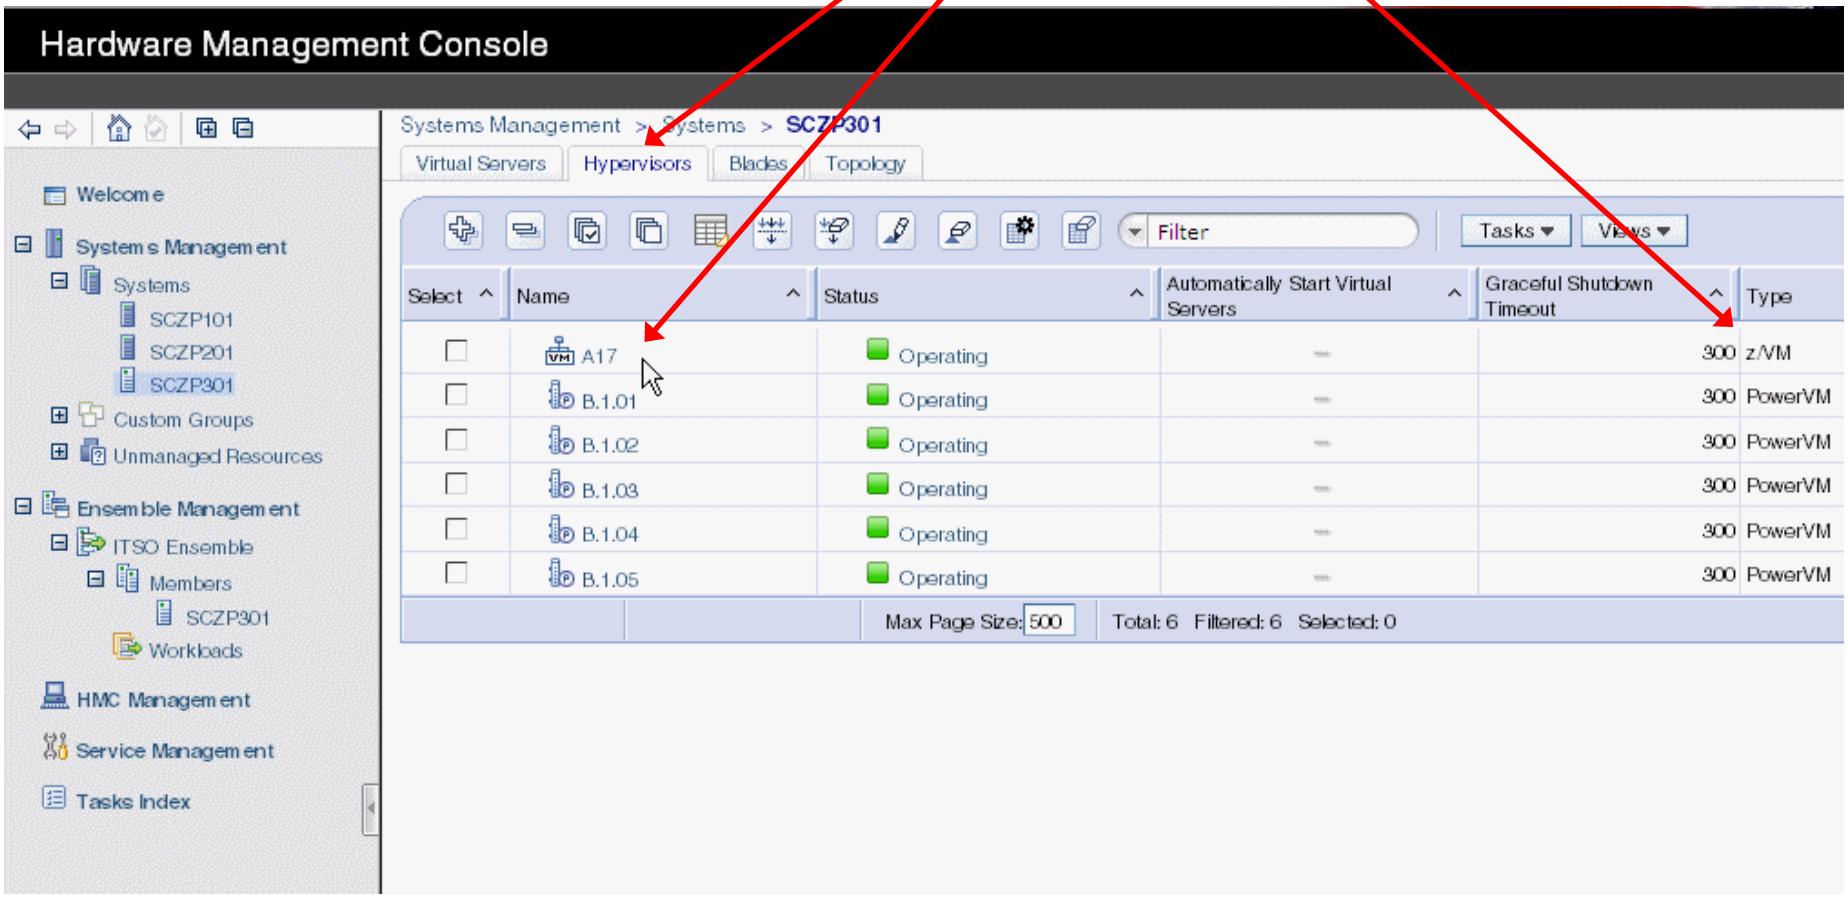

Validating your configuration

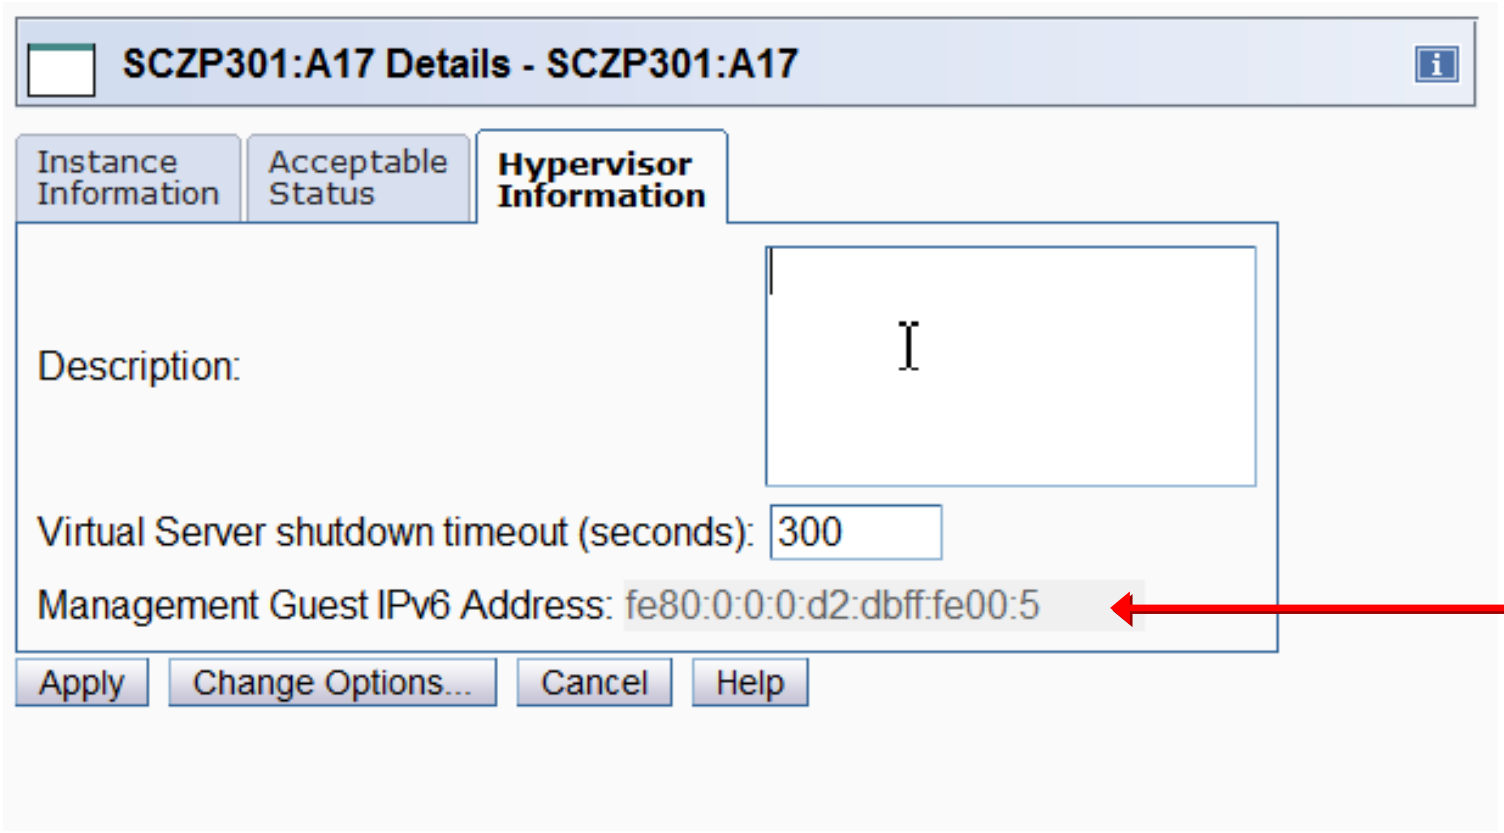

- The IPV6 IP address will display when the management guest is activated and zVM is part of the Ensemble

Implementation Tips

- Existing SMAPI servers need their existing PROFILE EXECs updated
- You can NOT manually define via CP commands an IEDN vswitch and attach to a guest. This must happen via zManager
- z/VM Ensemble configuration and logging in VMSYS file system. Back it up with the rest of your system.
- Resources must work without zManager if they are going to work with zManager. (ie 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.

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
- zManager is not a RACF security administration application. DIRMAINT actions driven by zManager can still invoke the DIRMAINT RACF interface.
- If ZVMLXAPP does not start, the other SMAPI service machines will not be started.

Once setup, Possible Next Steps

- Define IEDN Virtual Switches via Unified Resource Manager
- Define disk storage resources in the Unified Resource Manager
- Define virtual server containers for Linux guests or migrate existing guest
- Manage guest resources via Unified Resource Manager

Managing guest priorities from zManager

- Only one resource manager at a time
  - If you are managing a guest with VMRM don't enable zManager to also manage its resources at the same time
  - Enable one or the other, not both

Ensemble Enabling Linux on System z Guests

Linux considerations for residing in an Ensemble

- Native OSX interfaces are supported by the more recent kernel levels (RHEL 6.1, SLES 10 SP3, SLES 11 SP1).
- Keep this in mind when installing, you won't find that OSX OSA, unless the kernel has OSX support
- Might need to add a udev entry for OSX devices
- Utilize OSDSIM support when you can't get to the latest kernel level
- If your Linux network configuration contains a MAC, remove it.
  - The ensemble could assign a different MAC next time
  - If that macs don't match, you won't be able to communicate

OSX Interface Defined to Virtual Server

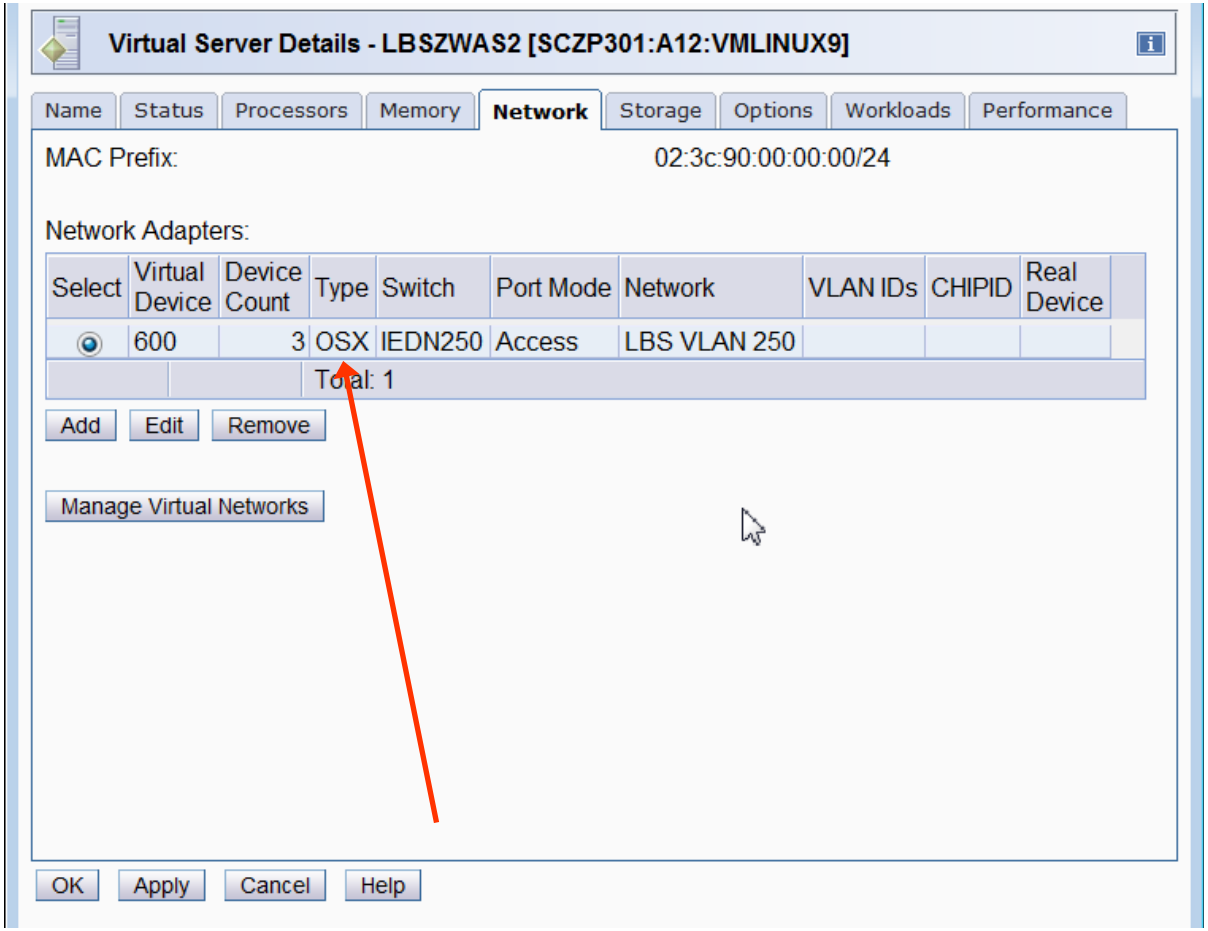

Installer Boot With OSX Interface

```
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
15:27:03 1) IBM IUCV
15:27:03 2) IBM IUCV
15:27:03 3) IBM IUCV
15:27:03
15:27:03 >
```

- The SLES 11 SP1 Installer System will NOT find an OSX interface

- Install with OSD Interface and add or convert existing after supporting kernel level is in place

Boot with same interface defined as OSD

```
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
15:29:48
15:29:48 >
```

If the Network Adapter is redefined as OSD instead of OSX (Utilizing OSDSIM) the OSA devices are discovered

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

Sample script with MAC coded

```
16:00:40 cat ifcfg-eth0
16:00:41 BOOTPROTO='static'
16:00:41 IPADDR='172.27.250.7/24'
16:00:41 BROADCAST='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 #
```

The LLADDR can be removed

```
16:03:56 cp ifcfg-eth0 backup-ifcfg-eth0
16:03:56 lbxzwas1:/etc/sysconfig/network # 16:04:36 sed '/LLADDR/d' backup-ifcfg
-eth0 > ifcfg-eth0
16:04:36 sed '/LLADDR/d' backup-ifcfg-eth0 > ifcfg-e <work # sed '/LLADDR/d' bac
kup-ifcfg-eth0 > ifcfg-et                                                  h0
16:04:36 lbxzwas1:/etc/sysconfig/network #
```

Desired Script with LLADDR Removed

```
16:05:04 cat ifcfg-eth0
16:05:05 BOOTPROTO='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 #
```

Migrating an Existing Virtual Server

- 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 migrate all guests. You may chose to just migrate the ones you want.

Migrating an Existing Virtual Server

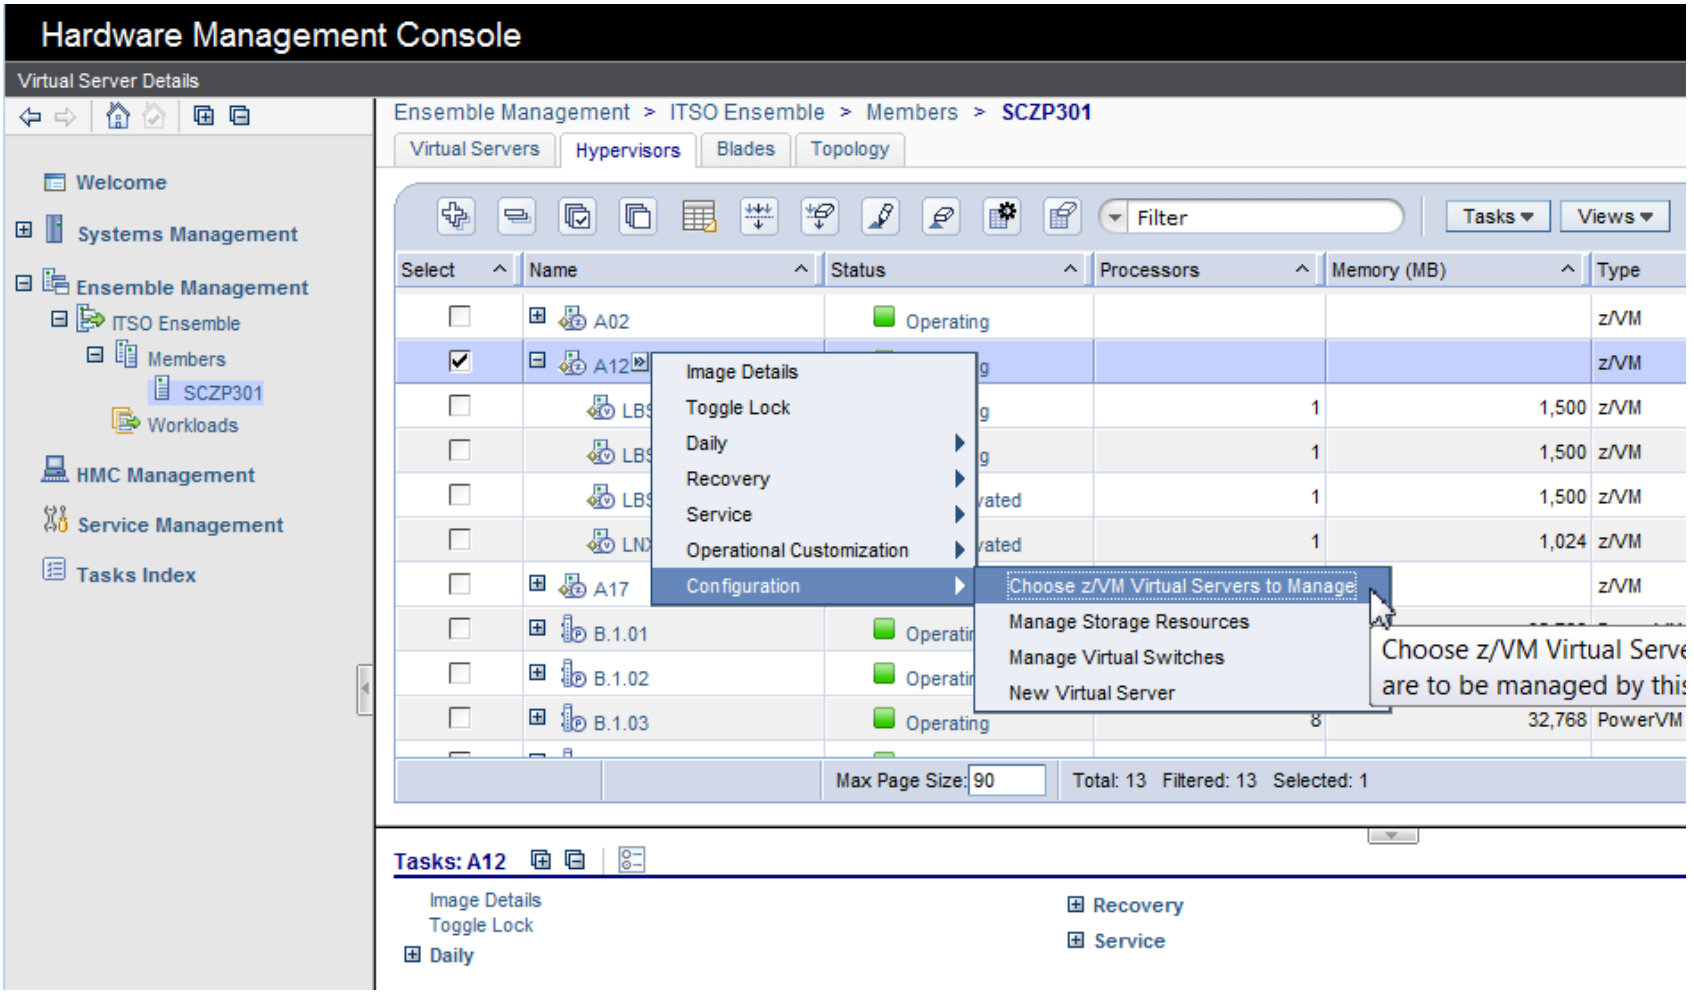

Migrating an Existing Virtual Server

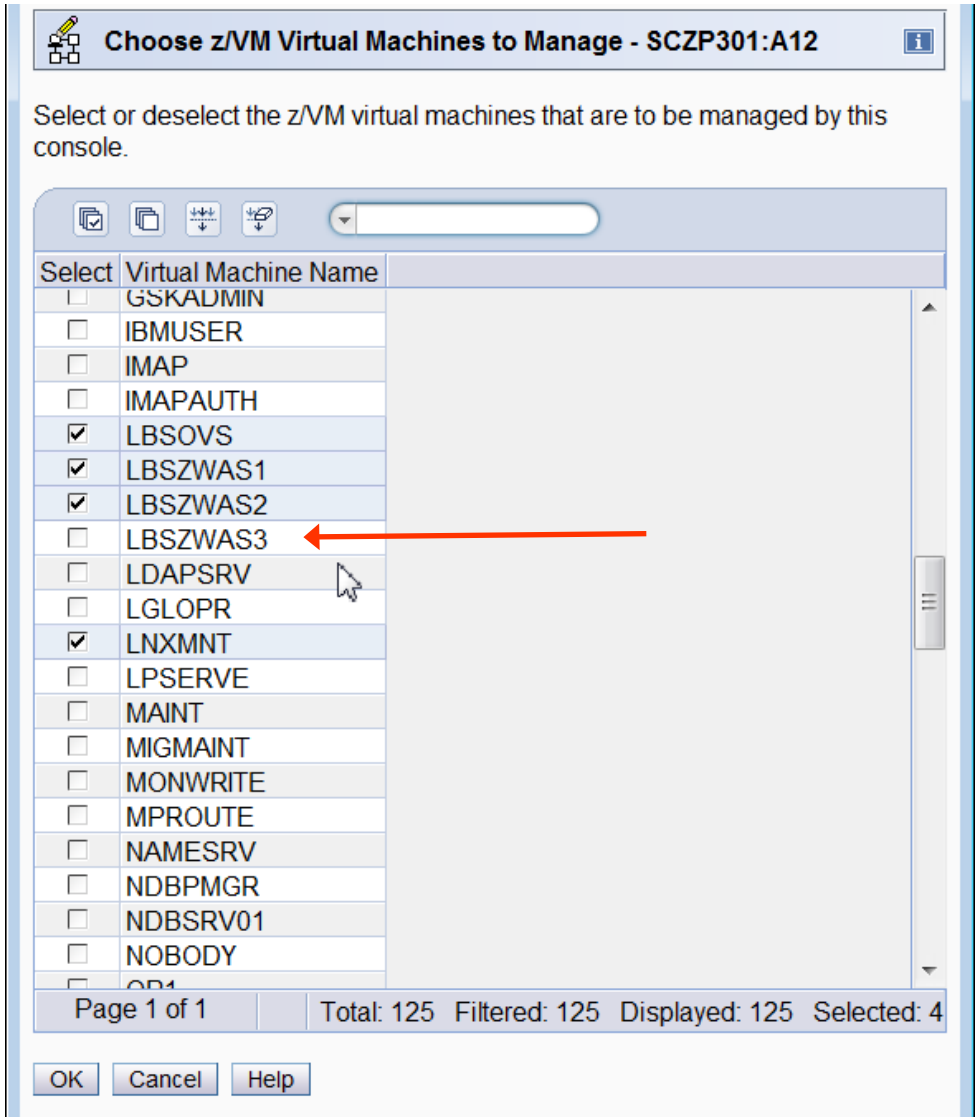

Migrating an Existing Virtual Server

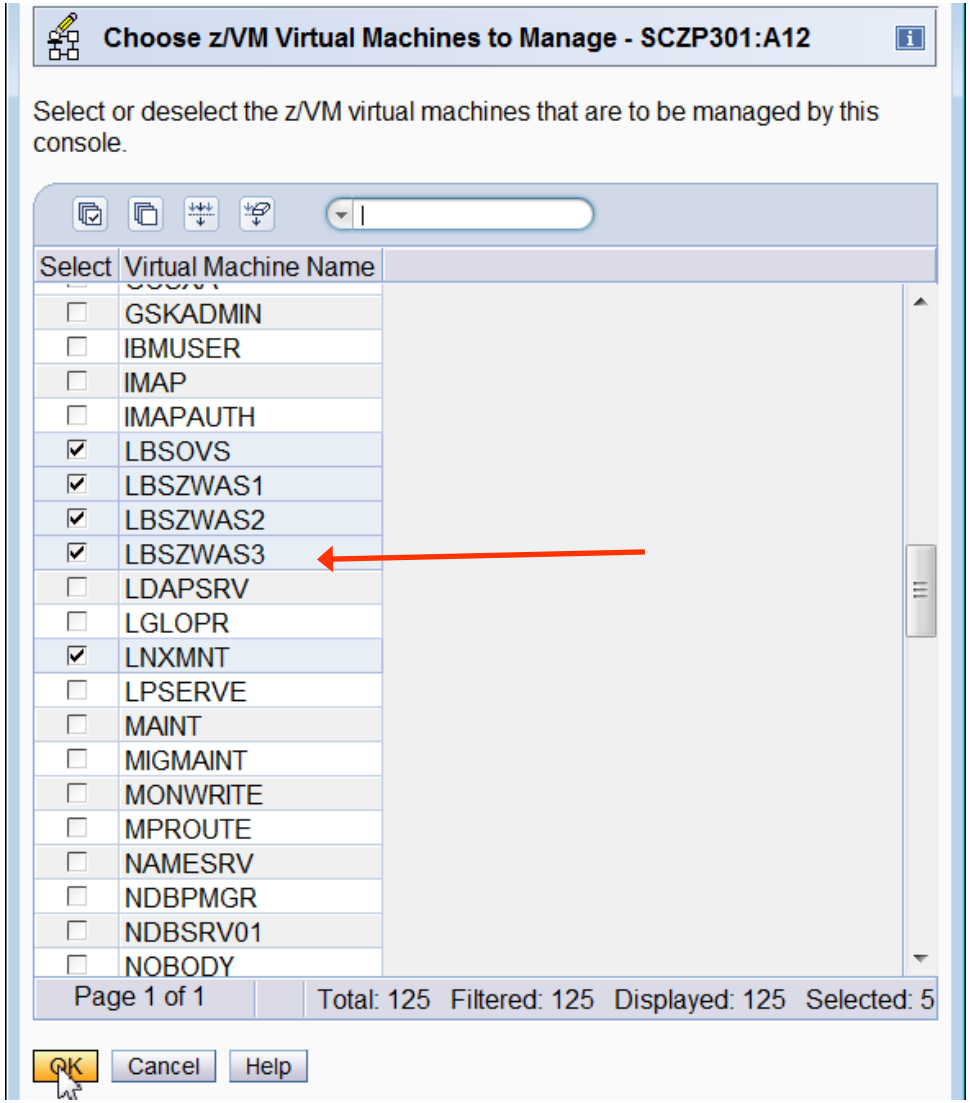

Migrating an Existing Virtual Server

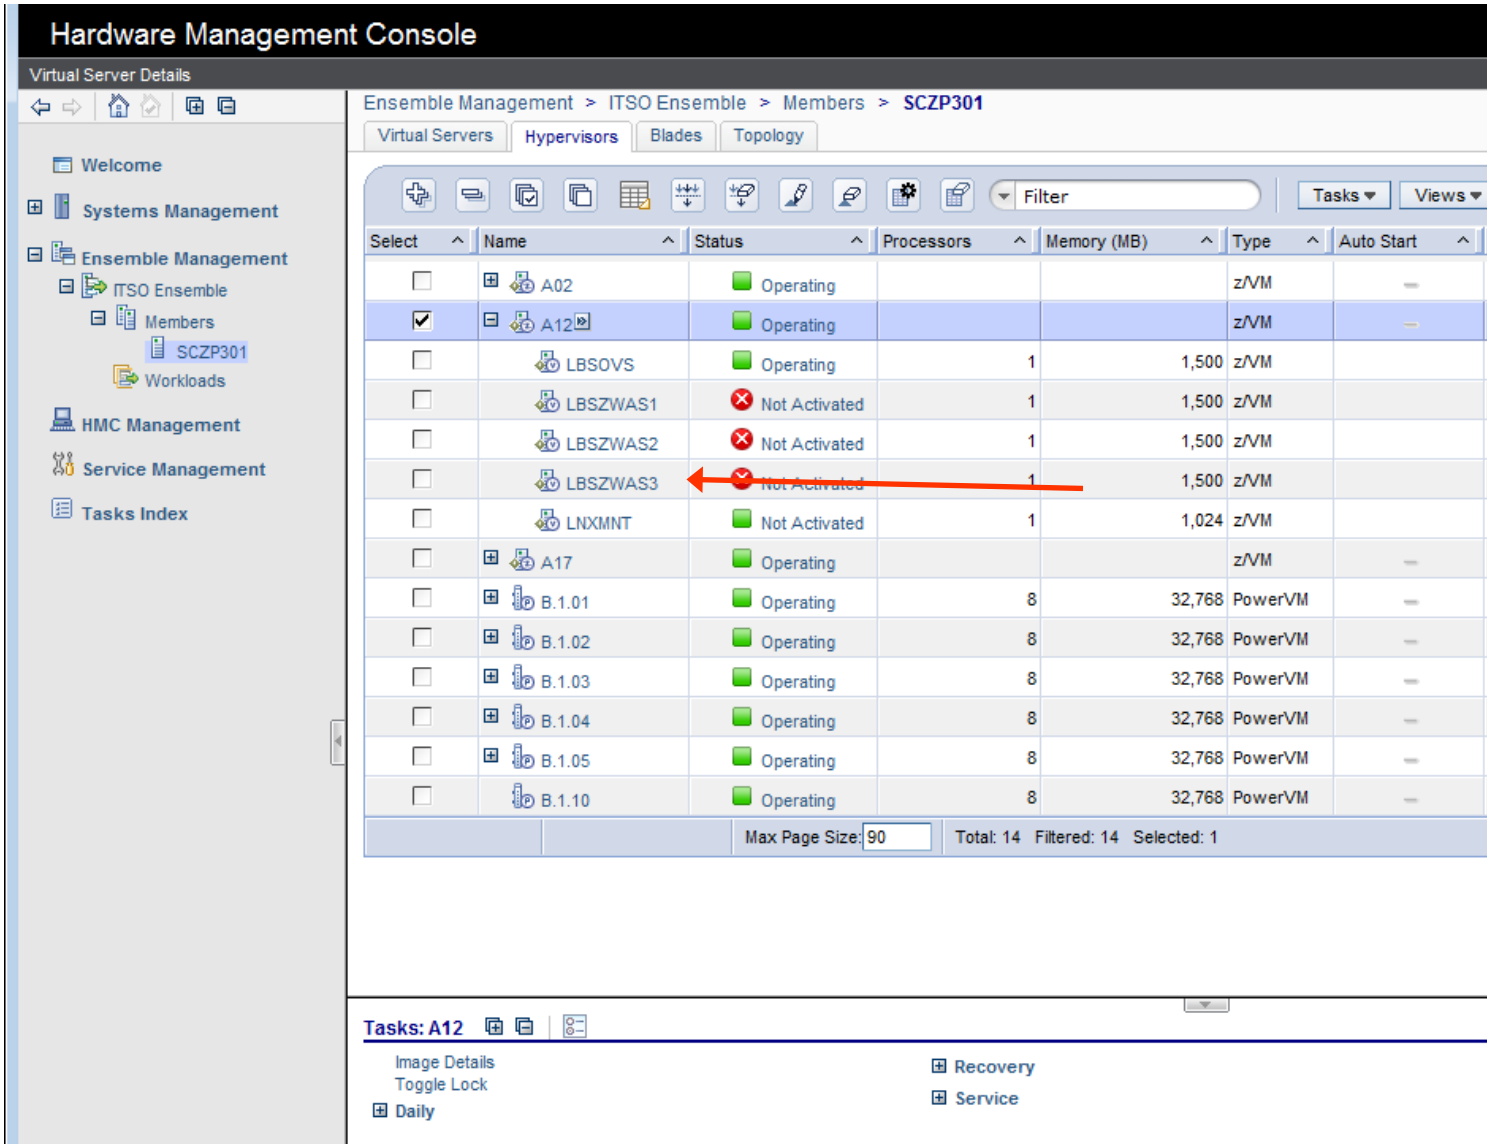

Migrating an Existing Virtual Server

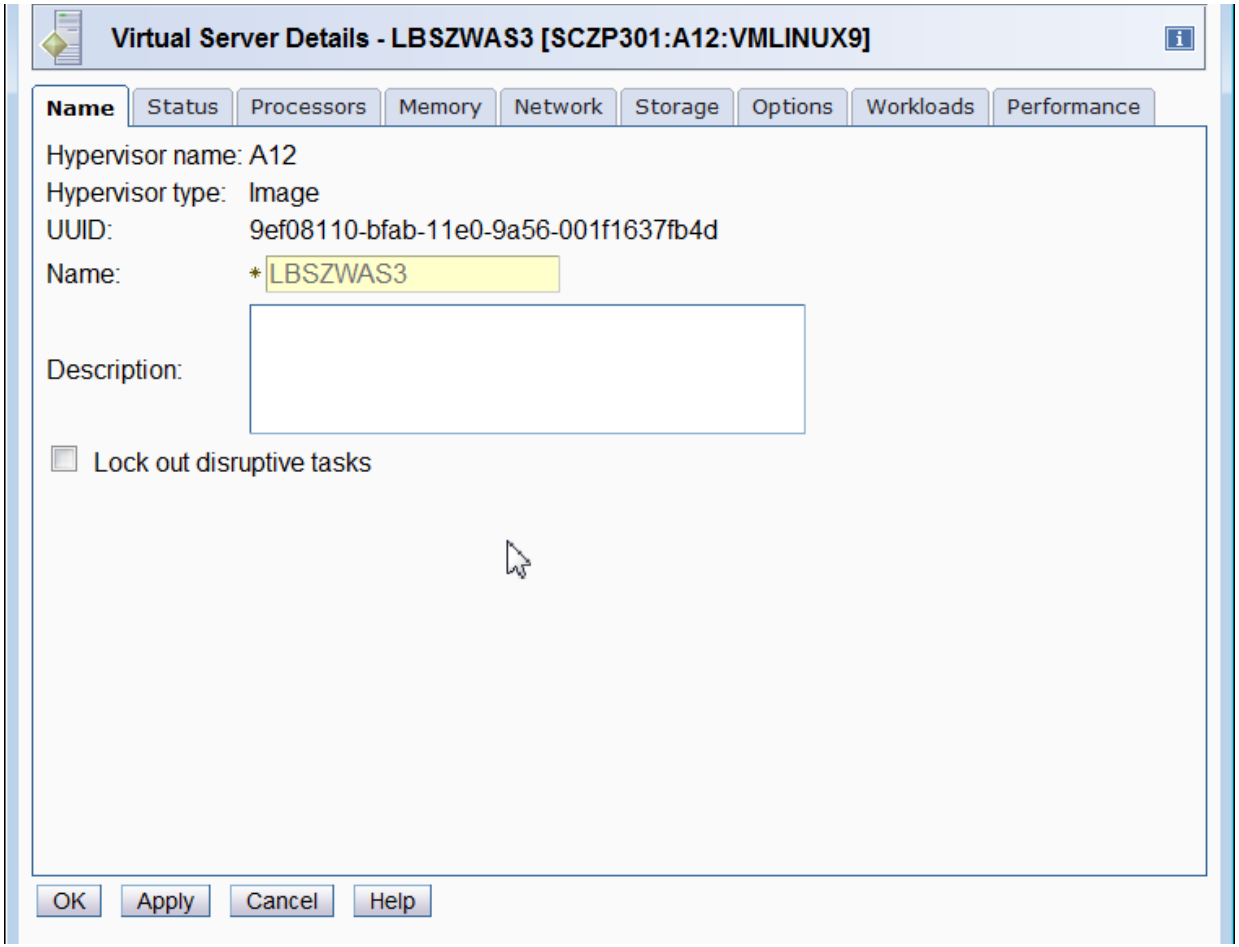

Once a guest is known to the Unified Resource Manager as a Virtual Server, you can change its configuration from the Unified Resource Manager

Migrating an Existing Virtual Server

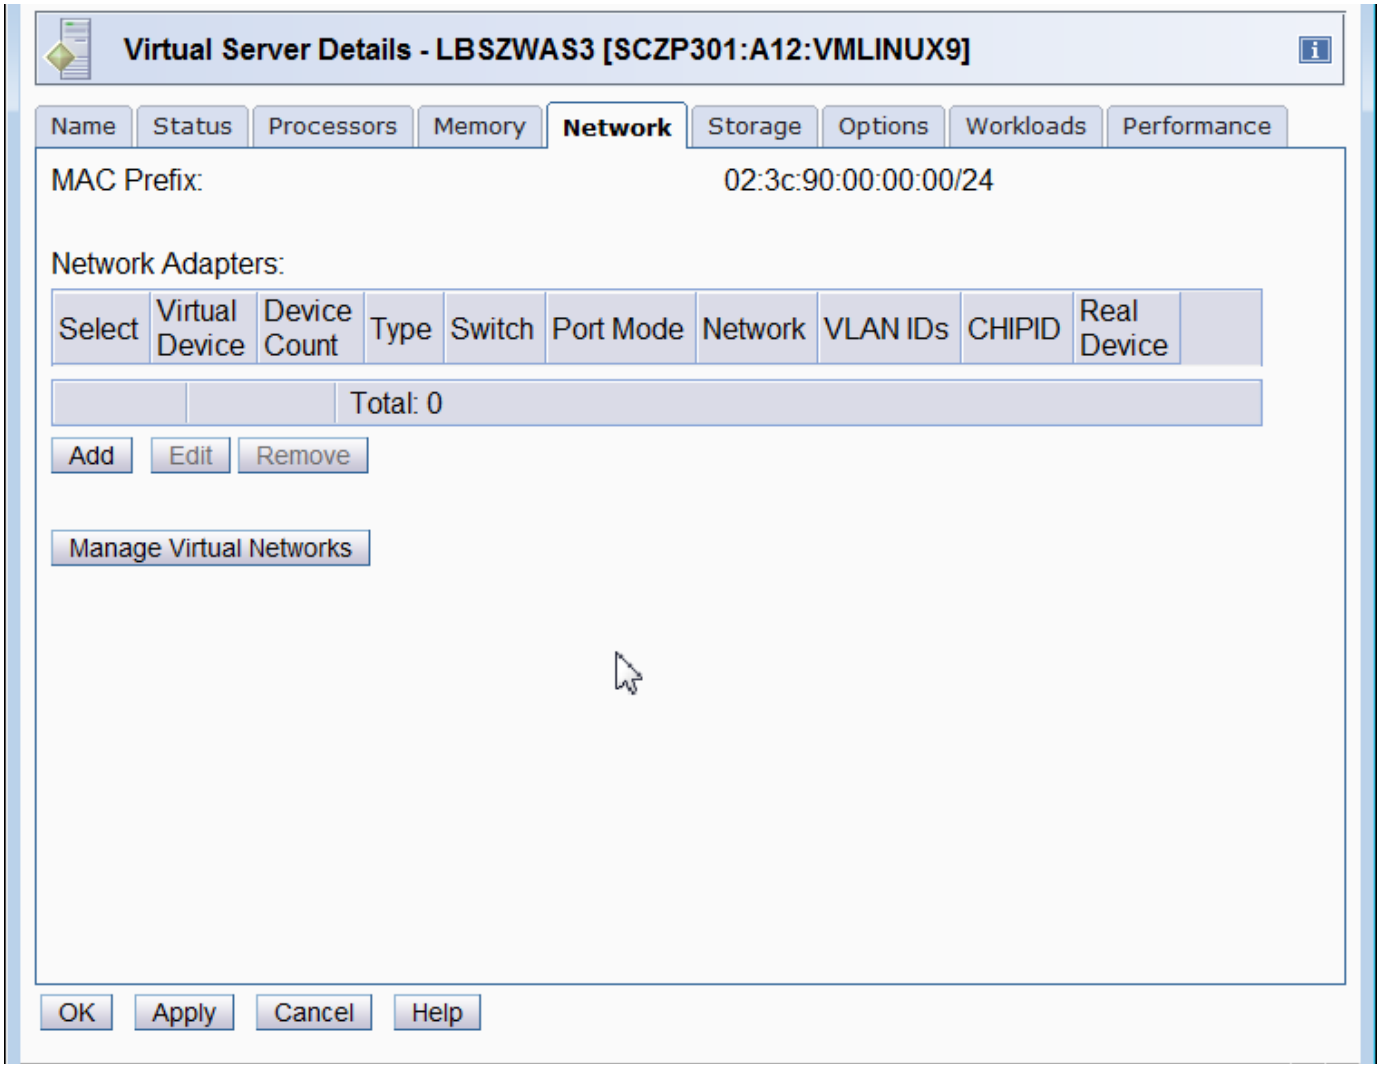

When migrated this guest had no NICs defined to it, so we will add two of them via the Unified Resource Manager

Migrating an Existing Virtual Server

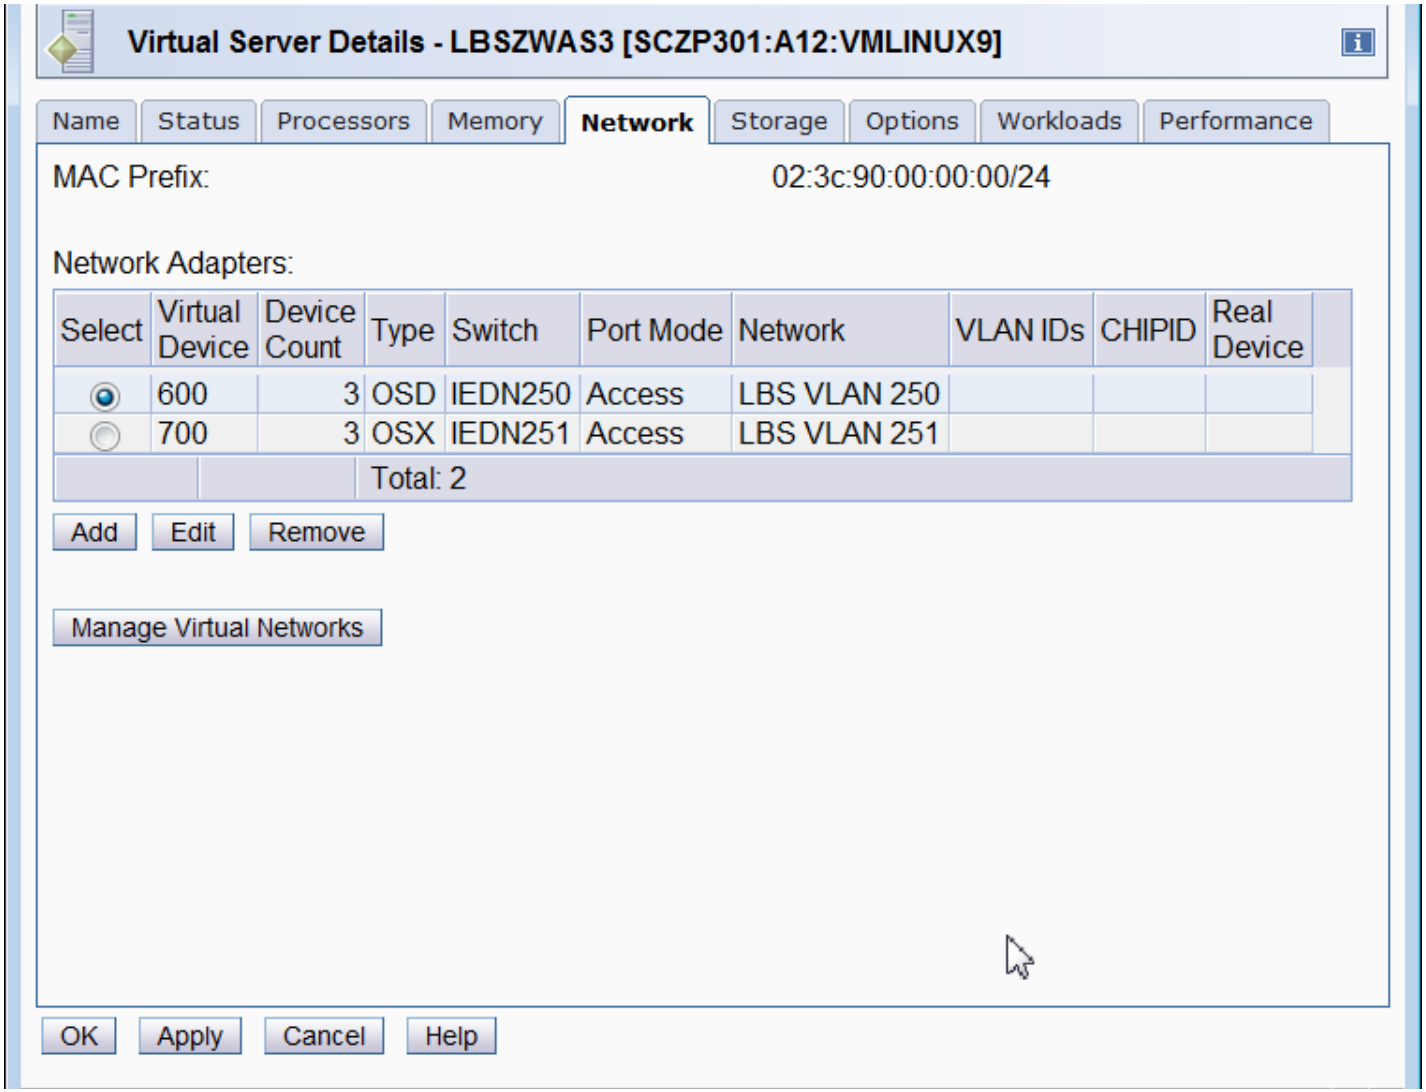

The Guest Platform Management Provider

- Provides more detailed level performance data from the guest operating system
- Connects to the INMN
- Optional, but is required if you want to feed Application Response Measurement data (ARM) to the zManager
- Is provided via the zManager code stream

Enabling INMN on Linux for System z

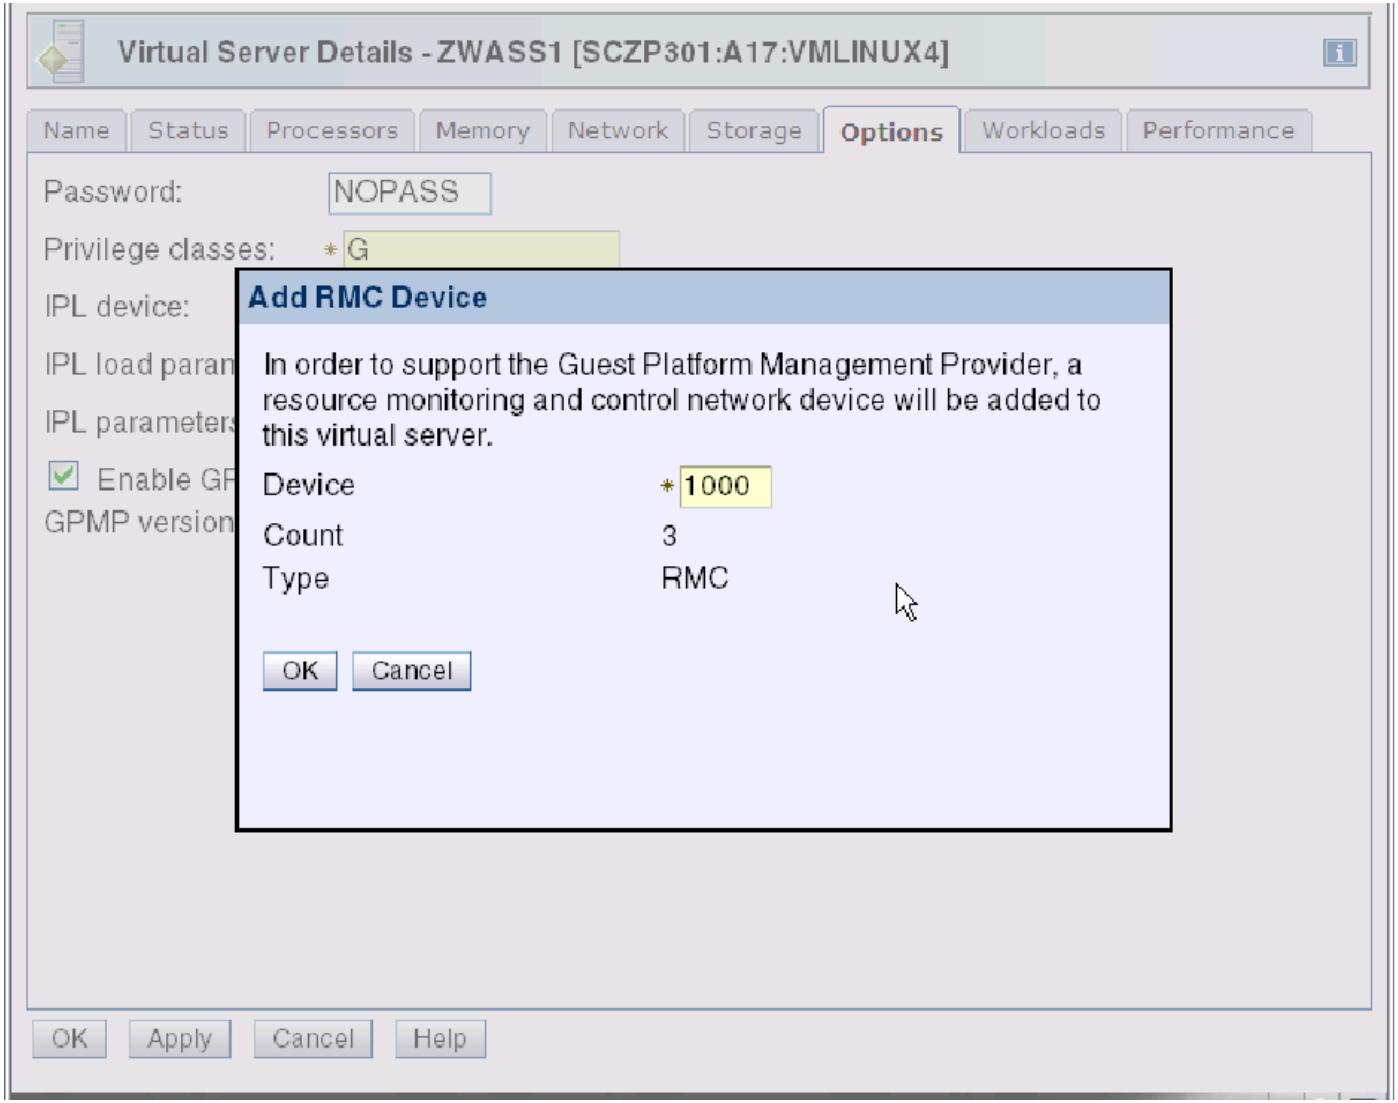

Enabling INMN on Linux – Yast Example

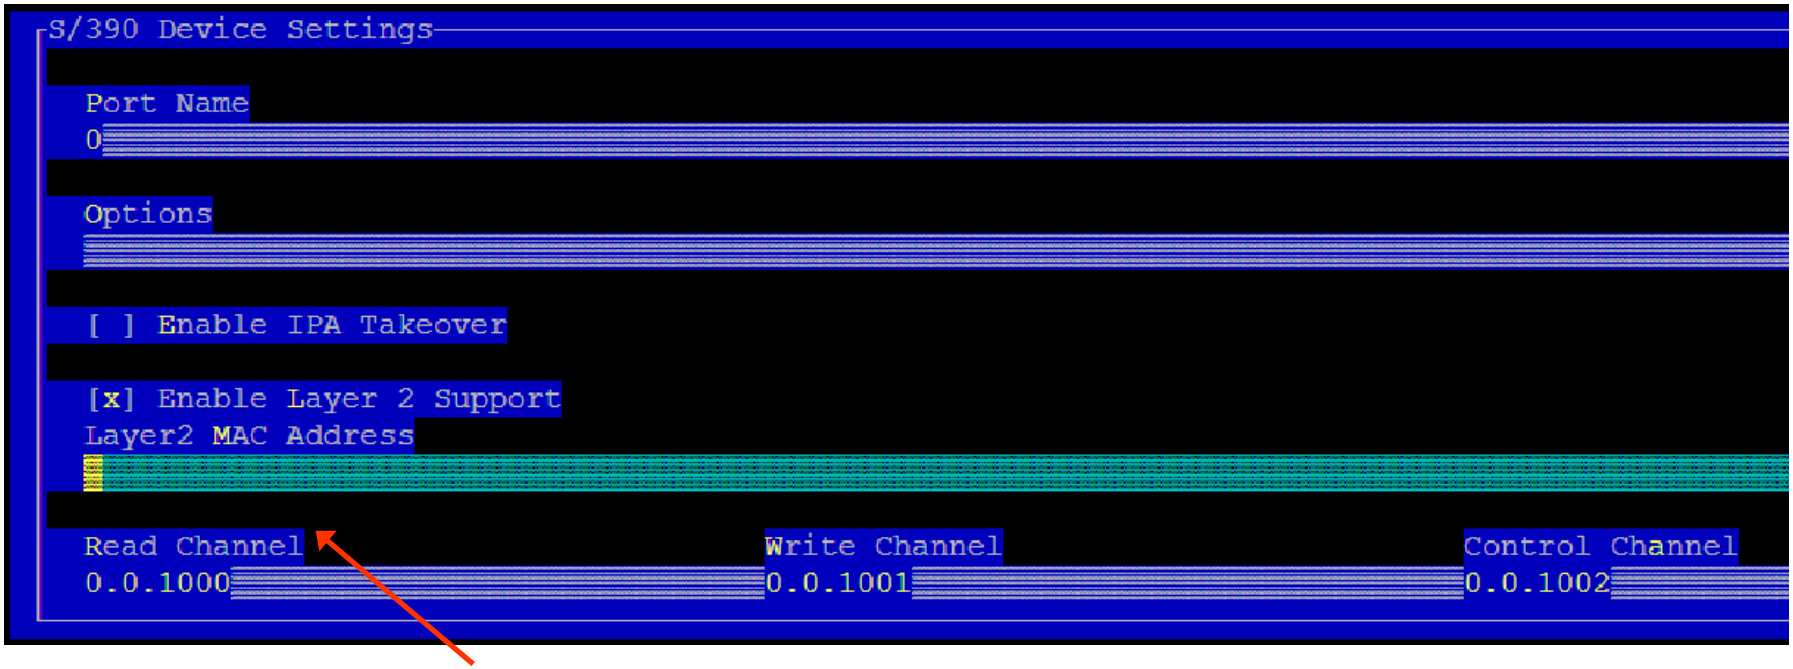

- Remember we do not want to code a MAC address in the configuration files. The Ensemble could assign a different virtual MAC the next time this guest is started. Communications would fail if we tried to assign a different MAC

Enabling INMN on Linux for System z

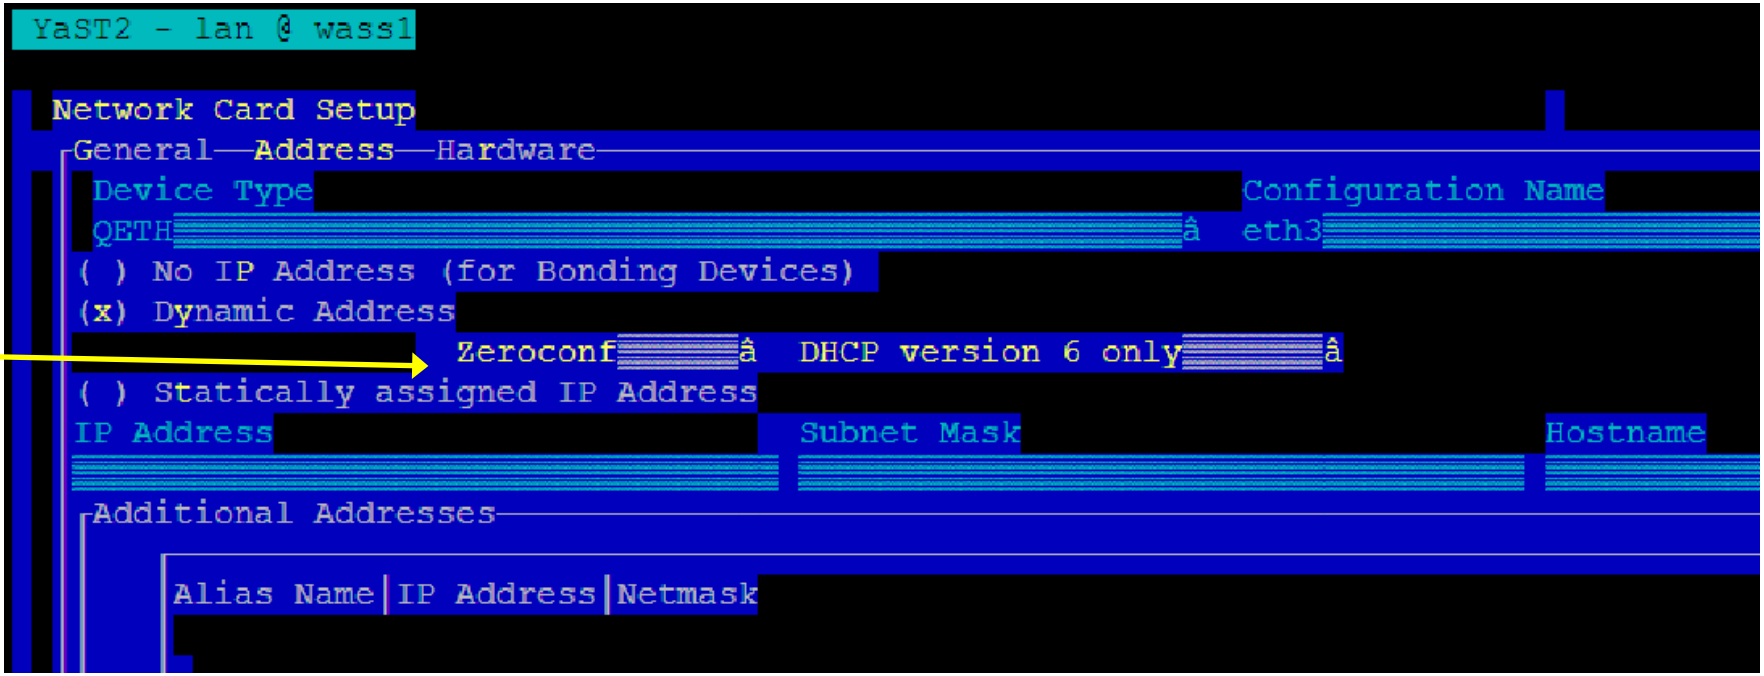

- Utilize “Zeroconf” to dynamically assign the IP address
- The assignment is made by the Ensemble

Validating INMN on Linux for System z

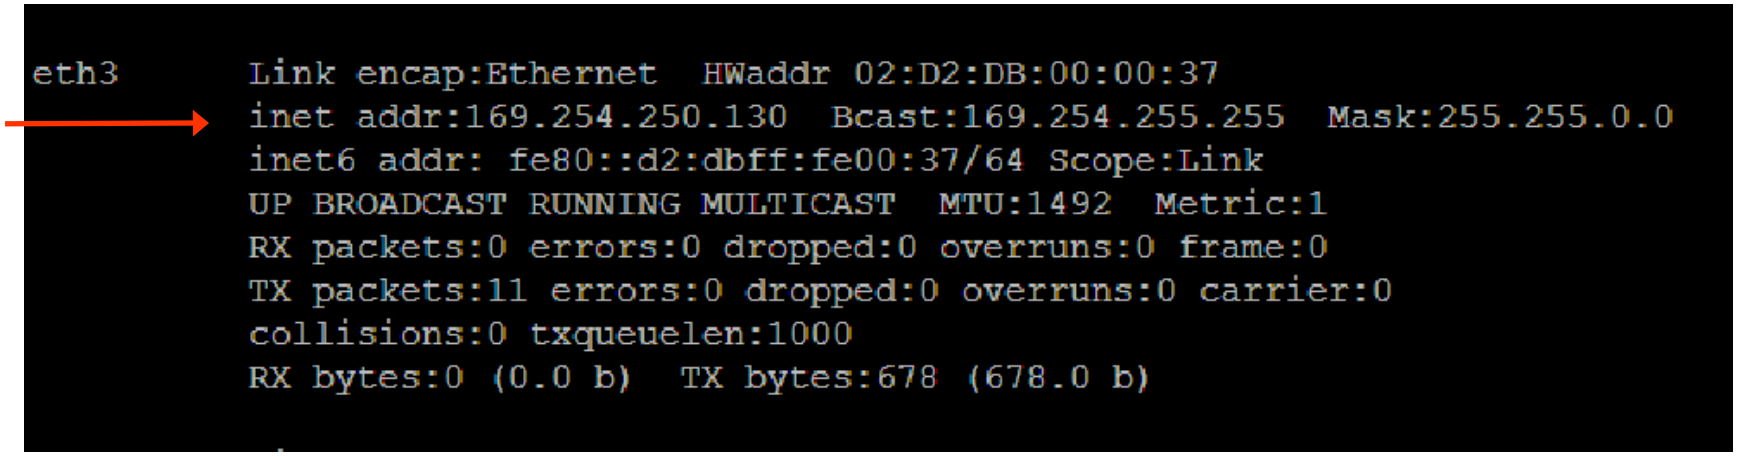

Validating INMN on Linux for System z

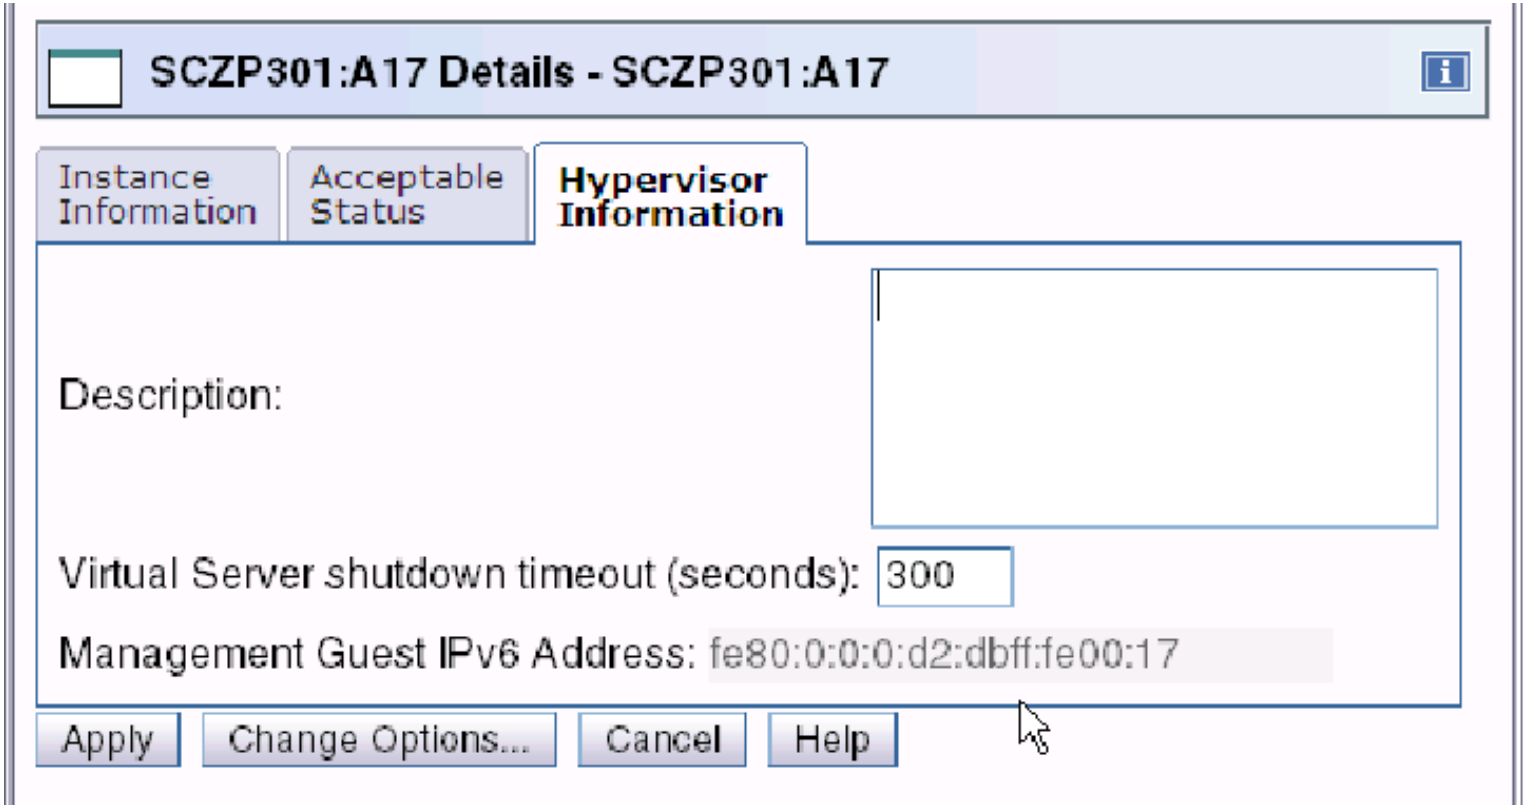

Validating INMN on Linux for System z

```
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

Enabling the GPMP on Linux for System z

```
wass1:~ # rpm -ivh http://[fe80:0:0:0:d2:dbff:fe00:17%eth3]:80/gpmp.s390x.rpm
Retrieving http://[fe80:0:0:0:d2:dbff:fe00:17%eth3]:80/gpmp.s390x.rpm
Preparing...                ########################################### [100%]
   1:gpmp                   ########################################### [100%]
Making group ibmlarm
Creating user ibmlarm with default group ibmlarm
Making group ibmgpmp
Creating user ibmgpmp with default group ibmgpmp
Adding user ibmgpmp to group ibmlarm
Running ldconfig to set up ARM libraries... Done with ldconfig.
Fixing permissions
chown ibmgpmp:ibmgpmp /opt/ibm/gpmp
chmod 555 /opt/ibm/gpmp
```

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

Enabling the GPMP on Linux for System z

```
chown ibmgpmp:ibmgpmp /opt/ibm/gpmp/CollectFFDC.sh /opt/ibm/gpmp/armsad /opt/ibm/gpm
k /opt/ibm/gpmp/gpmpmain /opt/ibm/gpmp/gpmpsad /opt/ibm/gpmp/gpmpshm /opt/ibm/gpmp/j
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/gpmpmain /opt/ibm/gpmp/gpmpsad /opt/ibm/gpmp/gpmpshm /opt/ibm/gpmp/java /opt/ibm
/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 ibmgpmp:ibmgpmp /opt/ibm/gpmp/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/gpmp/gpmpshm
chgrp ibmlarm /opt/ibm/gpmp/gpmpshm
chmod 6550 /opt/ibm/gpmp/gpmpshm
chown root /opt/ibm/gpmp/gpasetuid
chmod 4550 /opt/ibm/gpmp/gpasetuid
Adding the gpmpcheck process to crontab for user ibmgpmp.

Note: Run /opt/ibm/gpmp/post-install-config to grant permissions
to other users to access GPMP and ARM components. Currently, only the
ibmgpmp 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 userid
- The second command will cause it to “autostart” on subsequent IPLs
- You can not start the GPMP under the root userid

Enabling the GPMP on Linux for System z

```
root      2200     1  0 11:51 ?        00:00:00 /usr/sbin/cupsd
root      2209     1  0 11:51 ?        00:00:00 /usr/sbin/nscd
root      2271     1  0 11:51 ?        00:00:00 /usr/lib/postfix/master
root      2299     1  0 11:51 ?        00:00:00 /usr/sbin/gdm --no-console
root      2308     1  0 11:51 ?        00:00:00 /usr/sbin/cron
root      2323     1  0 11:51 ?        00:00:00 /usr/sbin/xinetd -pidfile /var/run/xinet
root      2325     2  0 11:51 ?        00:00:00 [flush-253:1]
root      2332     1  0 11:52 ttyS0    00:00:00 /sbin/mingetty --noclear /dev/ttyS0 dumb
root      2414  2192  0 12:41 ?        00:00:00 sshd: root@pts/0
root      2417  2414  0 12:41 pts/0    00:00:00 -bash
root      2449     2  0 12:41 ?        00:00:00 [flush-94:0]
root      4072     1  0 12:43 ?        00:00:00 /sbin/autoip -B eth3
postfix   4544  2271  0 12:43 ?        00:00:00 pickup -l -t fifo -u
postfix   4545  2271  0 12:43 ?        00:00:00 qmgr -l -t fifo -u
root      4659     2  0 12:45 ?        00:00:00 [flush-253:0]
ibmgpmp   4668     1  0 12:45 ?        00:00:00 gpmpmain daemon
root      4681  2417  0 12:46 pts/0    00:00:00 ps -ef
wass1:~ #
```

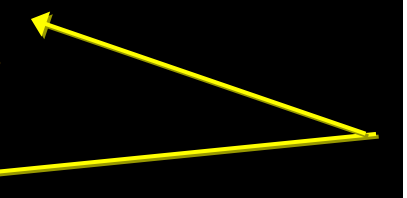

Enabling ARM, WebSphere on Linux for System z

```
wasg1:/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin # /opt/ibm/gpmp/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.
```

Enabling ARM, WebSphere on Linux for System z

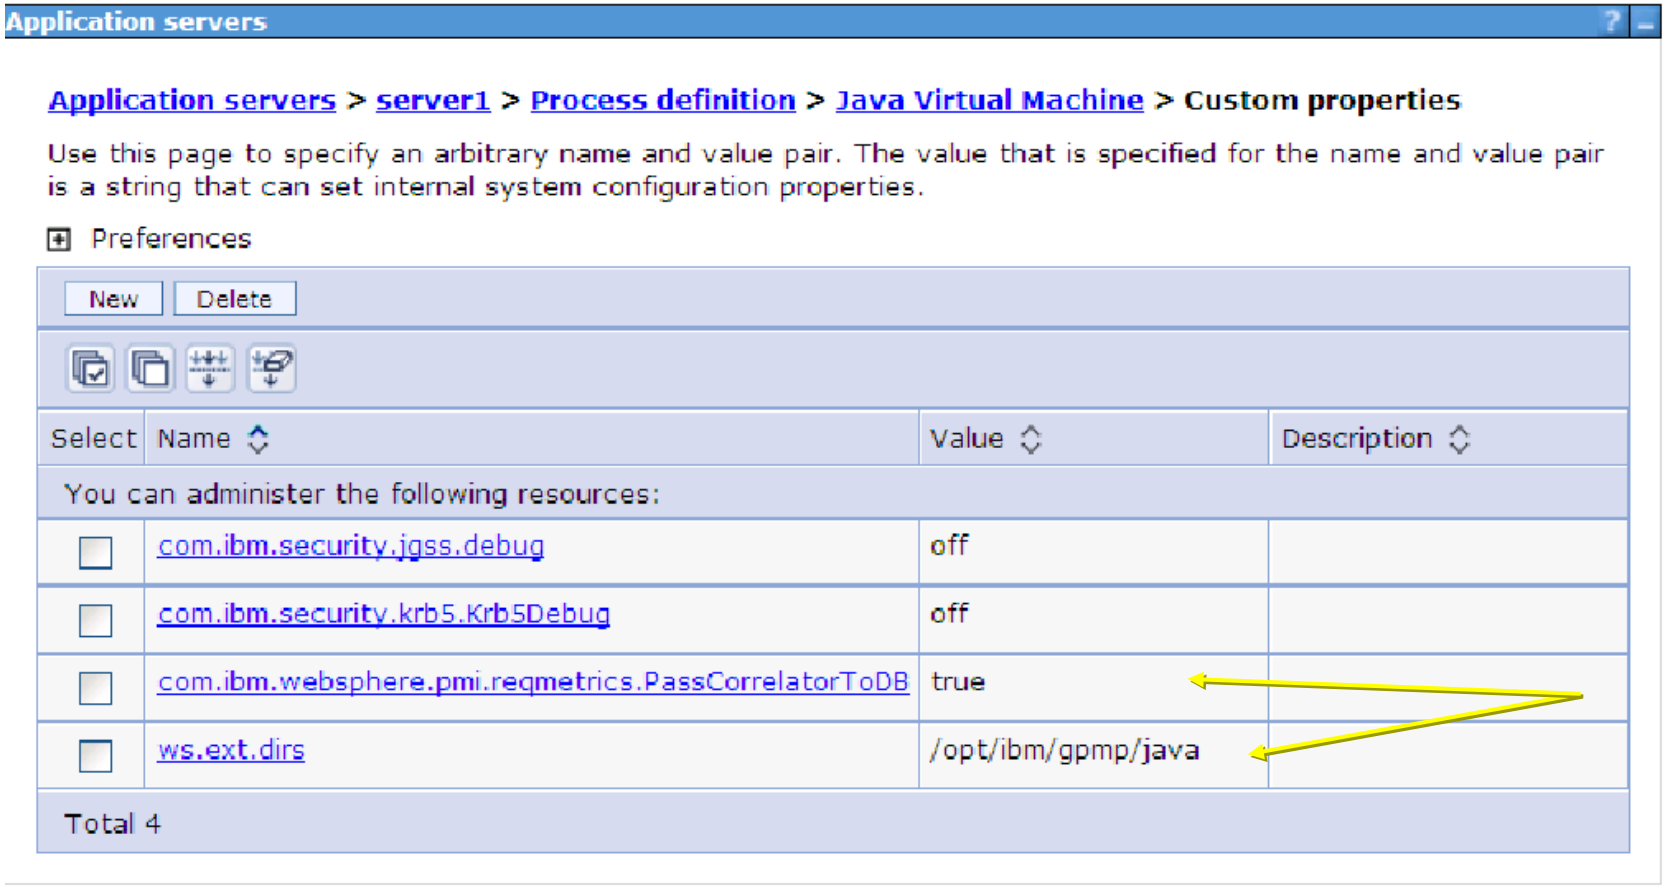

Enabling ARM, WebSphere on Linux for System z

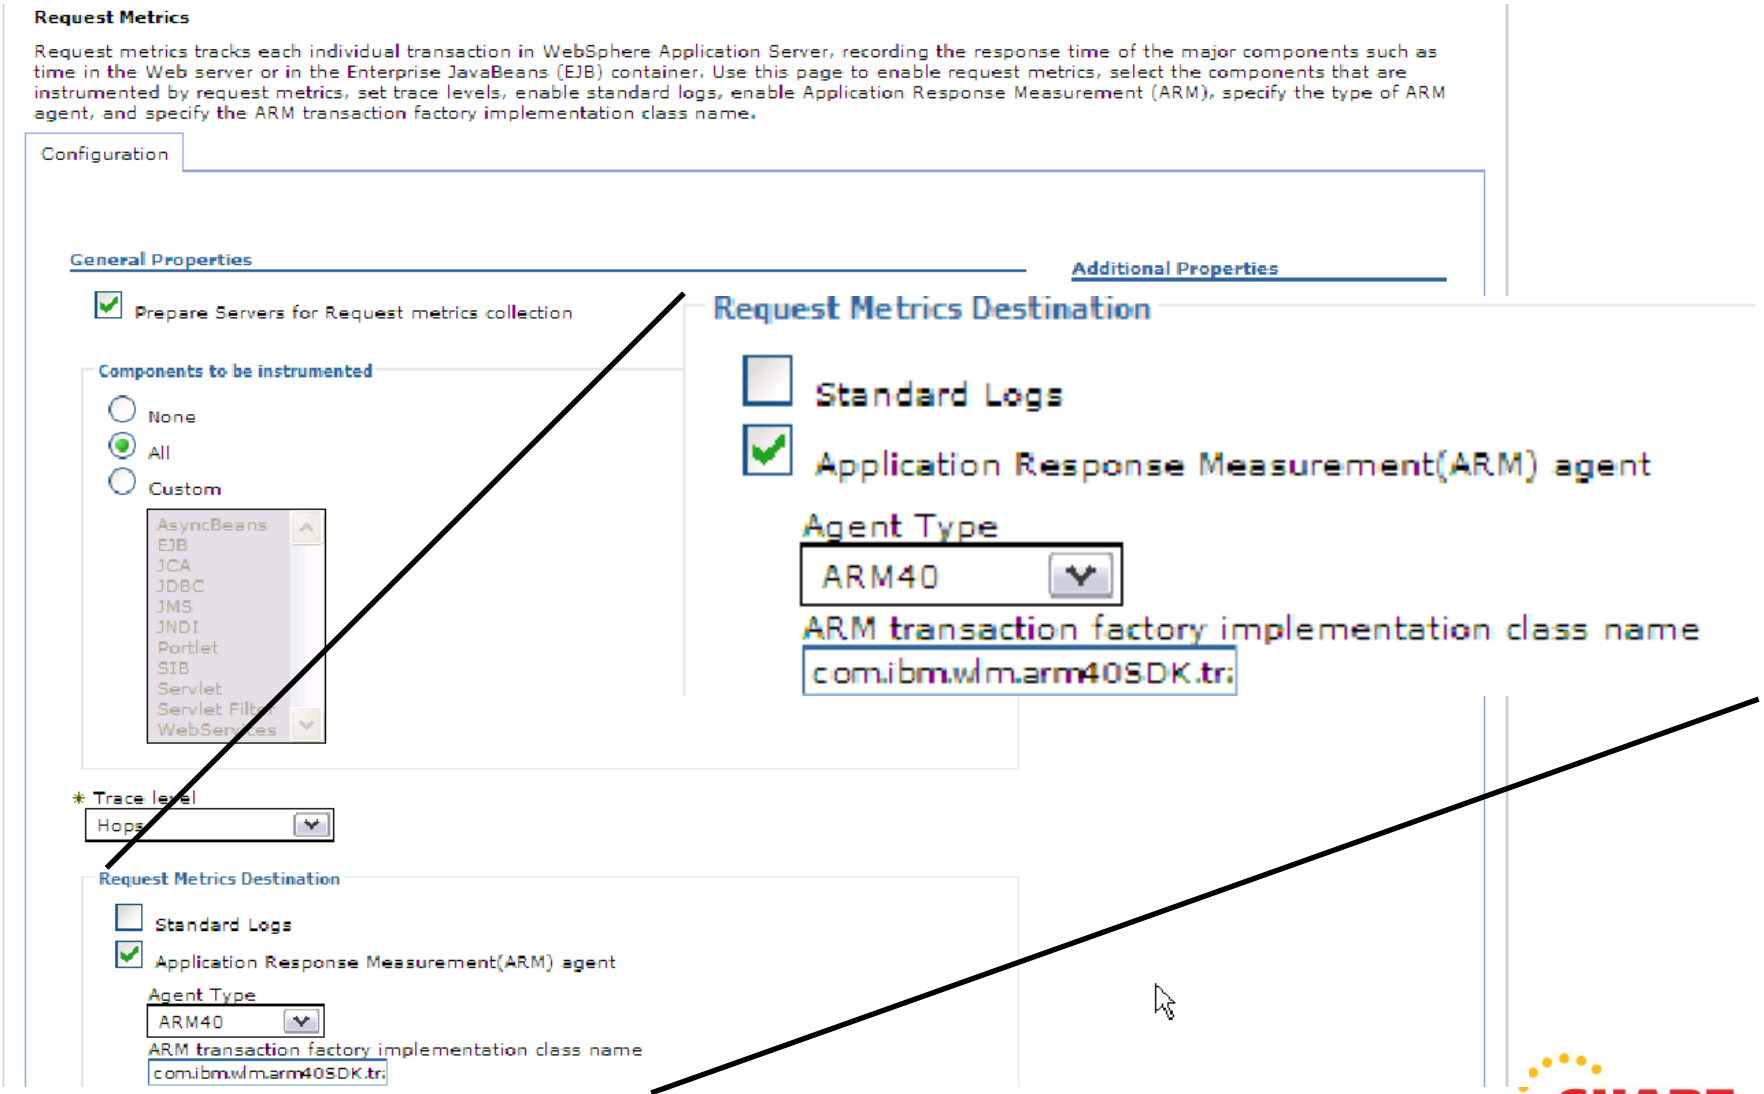

Enabling ARM on Linux for System z

```
case $PLATFORM in

  AIX)

    WAS_LIBPATH="$WAS_HOME"/bin
    NLSPATH=/usr/lib/nls/msg/%L/%N:/usr/lib/nls/msg/en_US/%N:${N
#     WAS_BOOTCLASSPATH=
    ;;

  Linux)

    WAS_LIBPATH="$WAS_HOME"/bin:/usr/lib64
    NLSPATH=/usr/lib/locale/%L/LC_MESSAGES/%N:${NLSPATH:-}
    JAVA_HIGH_ZIPFDS=200
#     WAS_BOOTCLASSPATH=
    ;;

  SunOS)

:
```

- **Update WAS setupCmdLine.sh**
- **Add :/usr/lib64**
- **To WAS_LIBPATH**

Enabling ARM, WebSphere on Linux for System z

```
wasg1:/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin # ./startServer.sh server1
ADMU0116I: Tool information is being logged in file
           /opt/IBM/WebSphere/AppServer/profiles/AppSrv01/logs/server1/startServer.log
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 6240
wasg1:/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin # /usr/sbin/lsarm -a
FEW6046I APPL: WebSphere:APPLICATION_SERVER
wasg1:/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin #
```

References

z/VM CP Planning and Administration Guide SC24-6178-01
z/VM CP Commands and Utilities Reference SC24-6175-01
z/VM Directory Maintenance Facility Commands Reference SC24-6188-01
zEnterprise Ensemble Performance Management Guide GC27-2607-01
zEnterprise Ensemble Planning and Configuration Guide GC27-2608-01
IBM zEnterprise Unified Resource Manager Redbook SG24-7921

धन्यवाद

Hindi

多謝

Traditional Chinese

감사합니다

Korean

Спасибо

Russian

Gracias

Spanish

Thank You

English

شكراً

Arabic

Obrigado

Brazilian Portuguese

Danke

German

Grazie

Italian

多谢

Simplified Chinese

Merci

French

நன்றி

Tamil

ありがとうございました

Japanese

ขอบคุณ

Thai

Questions?

Richard G. Young

Certified I/T Specialist

IBM STG Lab Services

zVM & Linux on z Team Lead

IBM

777 East Wisconsin Ave

Milwaukee, WI 53202

Tel 414 921 4276

Fax 414 921 4276

Mobile 262 893 8662

Email: ryoung1@us.ibm.com