

 #SHAREorg

# Enabling Linux for System z for Ensemble Management

Richard Young  
IBM STG Lab Services

Thursday, August 9<sup>th</sup>, 2012  
15:00

Session 11937





## 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](http://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

- 1 Why make your Linux guest part of the Ensemble
- 2 Requirements and Planning Considerations
- 3 Defining Resources for your Virtual Servers
- 4 Defining a New the Virtual Server or Managing an Existing Server
- 5 Linux Installation Considerations
- 6 Installing/Enabling the GPMP
- 7 Enabling ARM

# Agenda

- 1 Why make your Linux guest part of the Ensemble**
- 2 Requirements and Planning Considerations**
- 3 Defining Resources for your Virtual Servers**
- 4 Defining a New the Virtual Server or Managing an Existing Server**
- 5 Linux Installation Considerations**
- 6 Installing/Enabling the GPMP**
- 7 Enabling ARM**

## Why participate in an Ensemble ?

- Hybrid computing
  - Host parts of applications that are not supported on Linux for System z on Linux for System x (or Windows / AIX)
  - Fit for purpose selection of different application components
- Bring application components in closer to System z
- Simplify the administration of the Ensemble components
- Simplify the production and disaster recovery deployment thru an “in the box” model.
- IEDN Network access to exchange information with a zBX blades (x86, Power, Datapower)
- Dedicated and isolated high speed virtualized networks
- Simplified HMC based graphical administration common across hardware architectures
- Potentially reduce network hops
- Firmware management thru the System z maintenance stream for the entire Ensemble
- CE services the hardware instead of the CRU model for zBX components
- Linux guest could be the network entry/egress point for the IEDN via methods such as packet forwarding or VPN



# Key Enhancements for Ensemble Participation



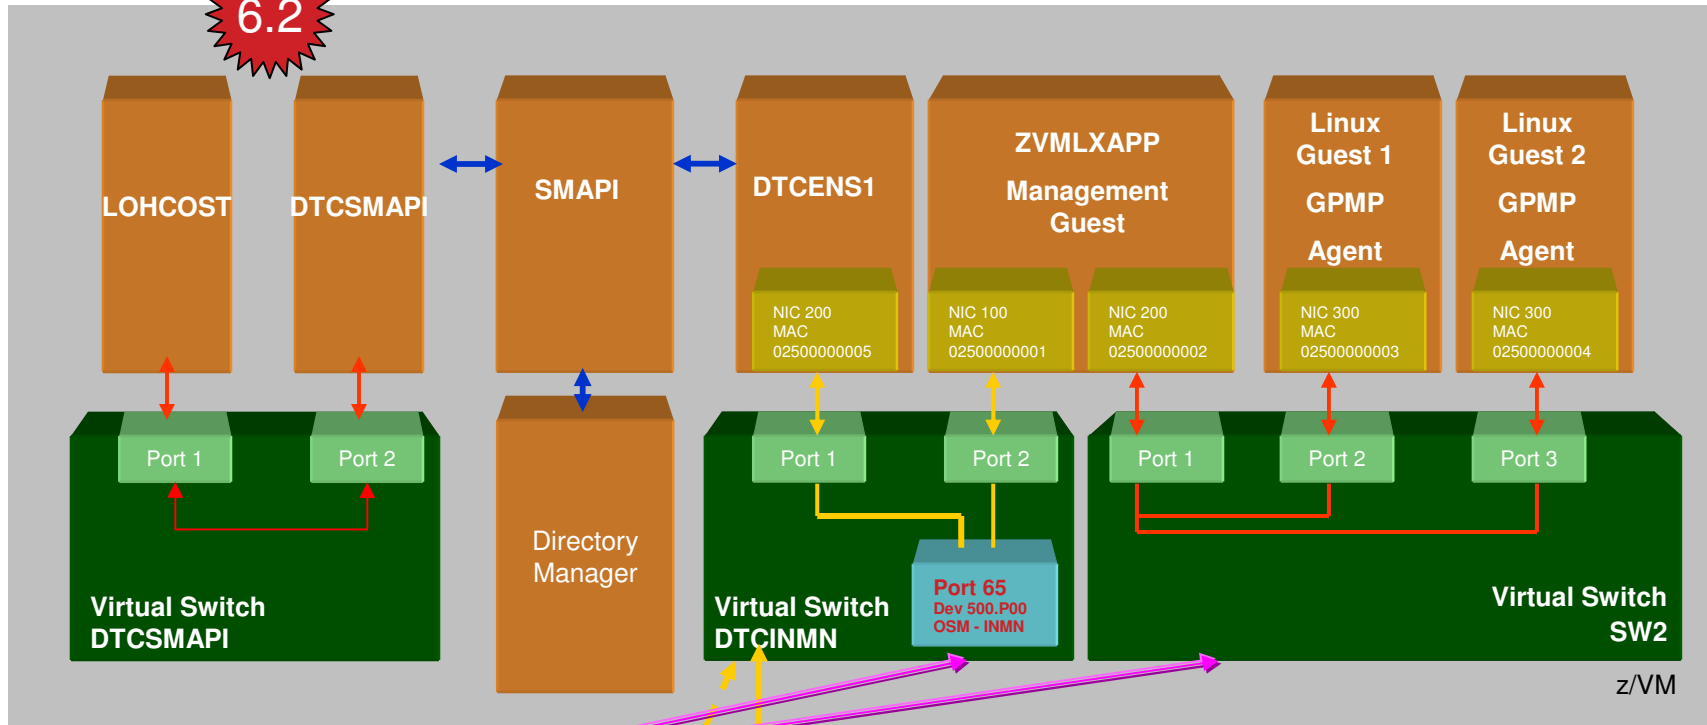
- Supported SLES and RHEL distributions
  - IEDN/INMN (OSX/OSM) NIC support
    - Kernel
    - UDEV
    - Tooling - YAST
  - Legacy NIC connection to IEDN or INMN via virtual switch possible via OSDSIM support in VSWITCH
  - Optional - Guest Platform Management Provider (GPMP)
  - Optional – Application Response Measurement (ARM) for middleware

# z/VM Ensemble INMN Infrastructure (System Defined)

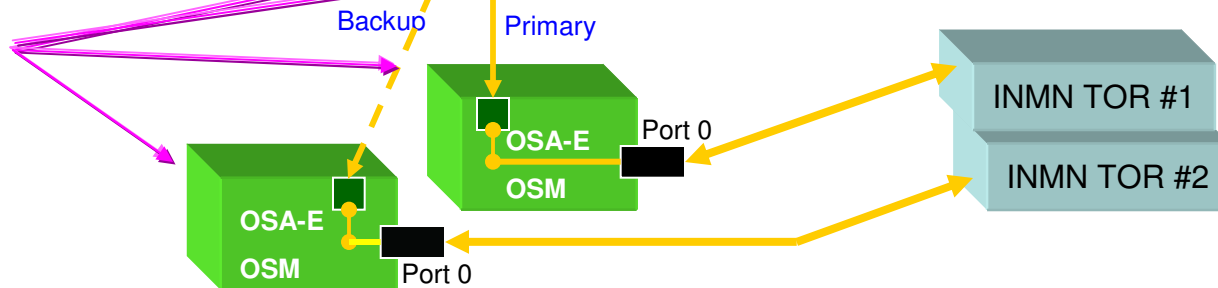


Required Infrastructure, but optional from a Linux OS perspective

6.2



ISOLATE = YES

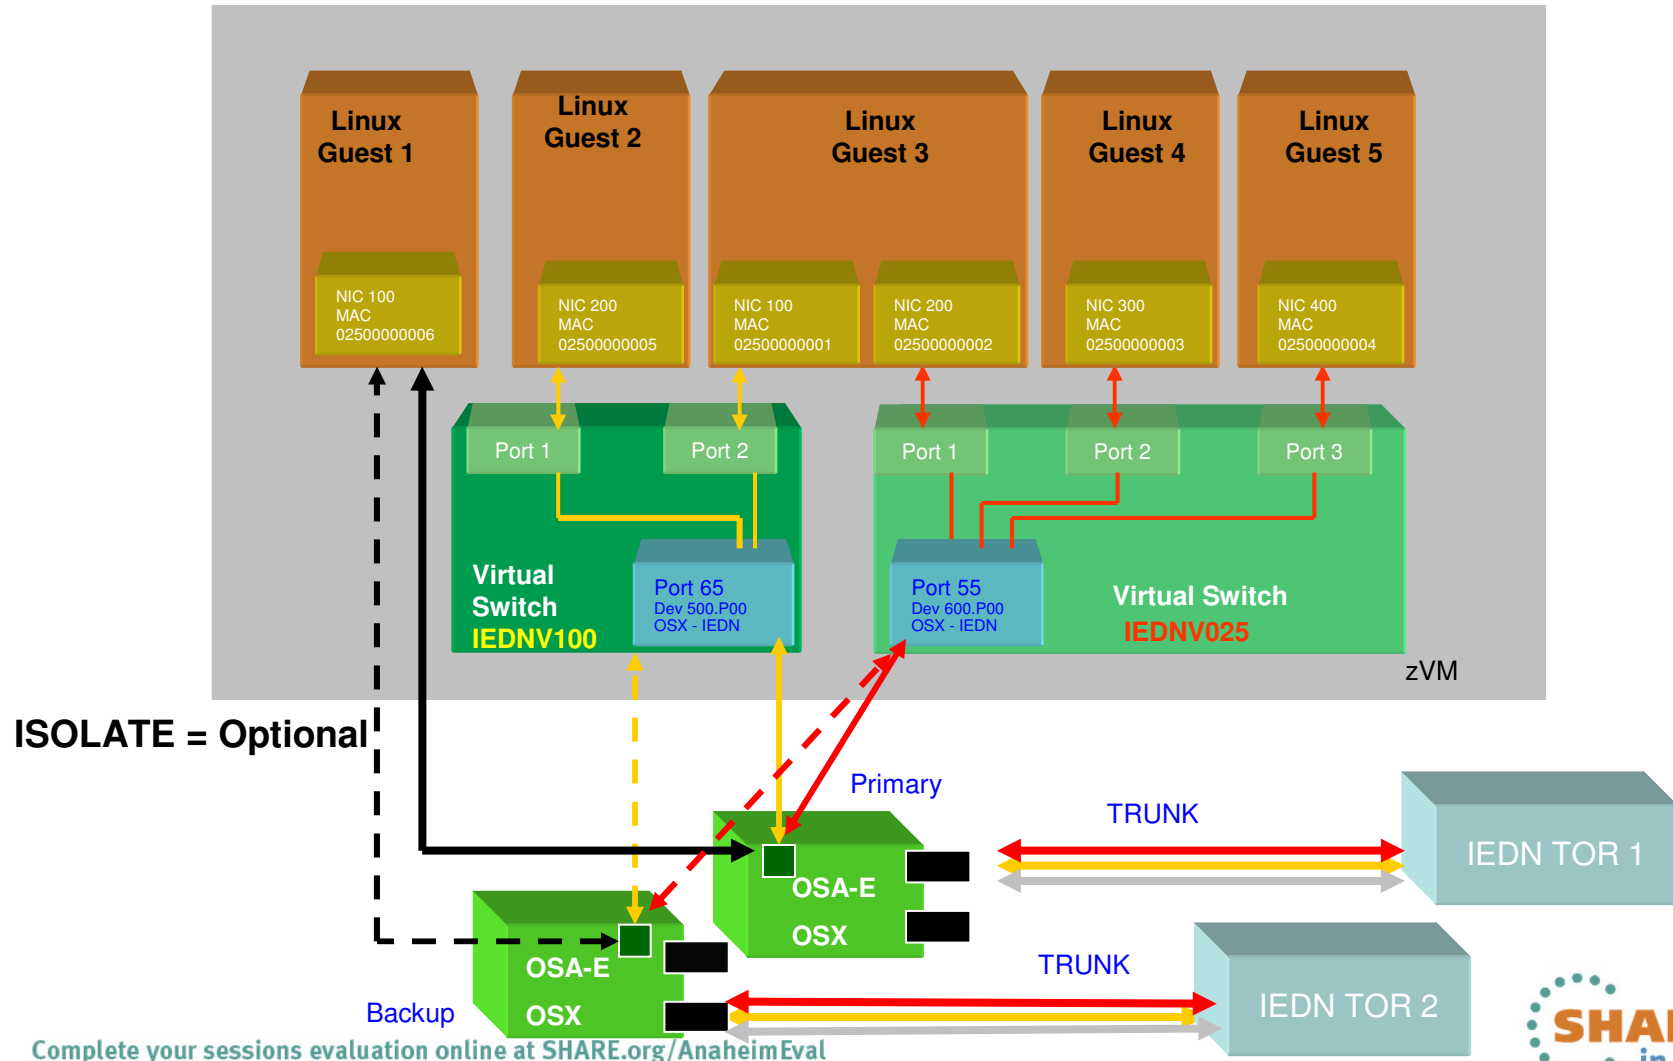




# z/VM Ensemble IEDN Infrastructure (You define)



Optional infrastructure, but commonly implemented from a Linux OS perspective





# Agenda

- 1 Why make your Linux guest part of the Ensemble
- 2 Requirements and Planning Considerations
- 3 Defining Resources for your Virtual Servers
- 4 Defining a New the Virtual Server or Managing an Existing Server
- 5 Linux Installation Considerations
- 6 Installing/Enabling the GPMP
- 7 Enabling ARM

# Planning Resources



- Manuals
  - Ensemble Planning and Configuration Guide GC27-2608
  - Ensemble Performance Management Guide GC27-2607
  - HMC Web Services API SC27-2616
- Red Book
  - IBM zEnterprise Unified Resource Manager SG24-7921
- IBM Resource Link
  - Driver Restriction Letter
- Required Maintenance
  - <http://www.vm.ibm.com/service/vmrequrm.html>

## Planning Considerations – Before you begin

- Not all guests in a z/VM need be Ensemble managed
- Not all LPARs on a CEC need be Ensemble managed
- Multiple paths for a guest to become “Ensemble Managed”
  - Existing guests can be “Migrated” to ensemble management at the HMC.
  - New guests can be directly created at the HMC and new Linux image put in place.
  - Create a new guest container at the HMC, but IPL an existing Linux guests disk in the new container.
- A Linux virtual server can be a member of the Ensemble and does not have to be run under z/VM, however Ensemble management of it is severely restricted.
  - No disk storage management
  - No INMN network
  - No GPMP
  - Only IEDN trunk connections supported (requires manual edits of Linux sysconfig files)

## Planning Considerations – Before you begin

- z/VM Linux guests may have their Ensemble MAC change every time the guest is started
  - Implications for DHCP and some provisioning products
- x86 Ensemble Linux virtual servers persist the Ensemble MAC across restarts
- z/VM Linux guests persist their network devices via the virtual device number and you might need to manually setup the udev entry.
- x86 Ensemble Linux virtual servers persist their network devices via the Ensemble virtual MAC and you must manually setup the udev entry if you have more than one nic.
- x86 Ensemble virtual server have multiple network and disk device drivers to pick from at the HMC. For z/VM Linux guests there are no choices at the HMC.

# Planning Considerations – Before you begin

- No support for Live Guest Mobility/Migration by the Unified Resource Manager
- Supported Levels of Linux:
  - System z (z196 or z114) - 64 bit distros only
    - RHEL 5.6 6.0, 6.1, 6.2
    - SLES 10 SP3, 11 SP1, SP2
  - x86\_64 (zBX HX5) – 64 bit distros only
    - RHEL 5.5, 5.6, 6.0
    - SLES 10 SP4, 11 SP1
- OSX Support
  - Red Hat RHEL 5.6 and RHEL 6.0 and above
  - SUSE SLES 10 SP4 and SLES 11 SP1 with a maintenance update
- GPMP Support (System z)
  - Red Hat Enterprise Linux (RHEL) 6, 5.5, 5.4, or 5.3
  - SUSE Linux Enterprise Server (SLES) 11 or 10
- Native LPAR
  - There is no guest platform management provider support for Linux on IBM System z running natively in an LPAR.



# Planning Considerations

- Access to the IEDN and OSX connections must be configured using Unified Resource Manager
  - Ensemble will reject “out-of-band” connection attempts
- Unified Resource Manager Workload Management
  - Use only one resource manager at a time
  - If you are managing a guest with z/VM VMRM don’t manage it with the Unified Resource Manager at the same time.
  - Enable one or the other, not both
  - Requires an INMN connection to the guest and that the GPMP be installed
- Next Steps
  - Create IEDN Virtual Switches to connect to the IEDN virtual networks
  - Define Ensemble managed disk storage resources
  - Define virtual server containers for Linux guests or migrate existing guests as needed to be managed by the Unified Resource Manager

# Agenda

- 1 Why make your Linux guest part of the Ensemble
- 2 Requirements and Planning Considerations
- 3 Defining Resources for your Virtual Servers
- 4 Defining a New the Virtual Server or Managing an Existing Server
- 5 Linux Installation Considerations
- 6 Installing/Enabling the GPMP
- 7 Enabling ARM



# Agenda

1 Why make your Linux guest part of the Ensemble

2 Requirements and Planning Considerations

3 Defining Resources for your Virtual Servers



Ensemble Storage Resources – Discovery, Definition, and Assignment

Ensemble Virtual Networking Resources and Virtual Switch Management

4 Defining a New the Virtual Server or Managing an Existing Server

5 Linux Installation Considerations

6 Installing/Enabling the GPMP

7 Enabling ARM

## Ensemble Storage Resources for System z

- The Unified Resource Manager requires storage resources be defined to it and assigned to each hypervisor
- Storage Types
  - ECKD
  - SCSI EDEVs (Emulated Devices)
  - No Ensemble support for dedicated FCP devices
- Discovery of both ECKD DASD and SCSI LUN is supported
- Presented to Linux as either
  - Fullpack Minidisk
  - Non-Fullpack Minidisk via storage group
  - Minidisk Links to other z/VM guests/virtual servers

# Ensemble Storage Resources

Ensemble Management > ITSO Ensemble > Members > SCZP301

Virtual Servers | Hypervisors | Blades | Topology

Filter: [ ] Tasks: [ ] Views: [ ]

Select	Name	Status	Processors	Memory (MB)	Type	Auto Start	Shutdown Timeout (s)
<input checked="" type="checkbox"/>	A12 (LBSSSI:VMLINUX9)	Operating			z/VM	-	
<input type="checkbox"/>	A17 (VMLINUX4)	Operating			z/VM	-	
<input type="checkbox"/>	B.1.01	Operating	1	32,768	PowerVM	✓	
<input type="checkbox"/>	B.1.02	Operating	1	32,768	PowerVM	✓	
<input type="checkbox"/>	B.1.03	Operating	1	32,768	PowerVM	✓	
<input type="checkbox"/>	B.1.04	Operating	1	32,768	PowerVM	✓	
<input type="checkbox"/>	B.1.05	Operating	1	32,768	PowerVM	✓	
<input type="checkbox"/>	B.1.10	Operating	1	32,768	PowerVM	✓	
<input type="checkbox"/>	B.1.11	Operating	2	131,072	x Hyp	✓	
<input type="checkbox"/>	B.1.12	Operating	2	131,072	x Hyp	✓	

Max Page Size: 30 Total: 12 Filtered: 12 Selected: 1

---

Tasks: A12 (LBSSSI:VMLINUX9)

- Image Details
  - Toggle Lock
- Daily
  - Activate
  - Deactivate
  - Grouping
  - Hardware Messages
  - Operating System Messages
  - Reset Normal
- Recovery
  - Access Removable Media
  - Integrated 3270 Console
  - Integrated ASCII Console
  - Load
  - Load from Removable Media or Server
  - PSW Restart
  - Reset Clear
  - Start All
  - Stop All
- Service
- Operational Customization
  - Configure Channel Path On/Off
  - Customize/Delete Activation Profiles
  - Logical Processor Add
  - View Activation Profiles
- Configuration
  - Choose z/VM Virtual Servers to Manage
  - Manage Storage Resources**
  - Manage Virtual Switches
  - New Virtual Server

/bonsai/action/T8243#

# Ensemble Storage Resources - Discovery



Discovery can provide a quicker and more accurate way to get new storage resources defined to the Ensemble compared to manual entry

Manage Storage Resources - ITSO Ensemble

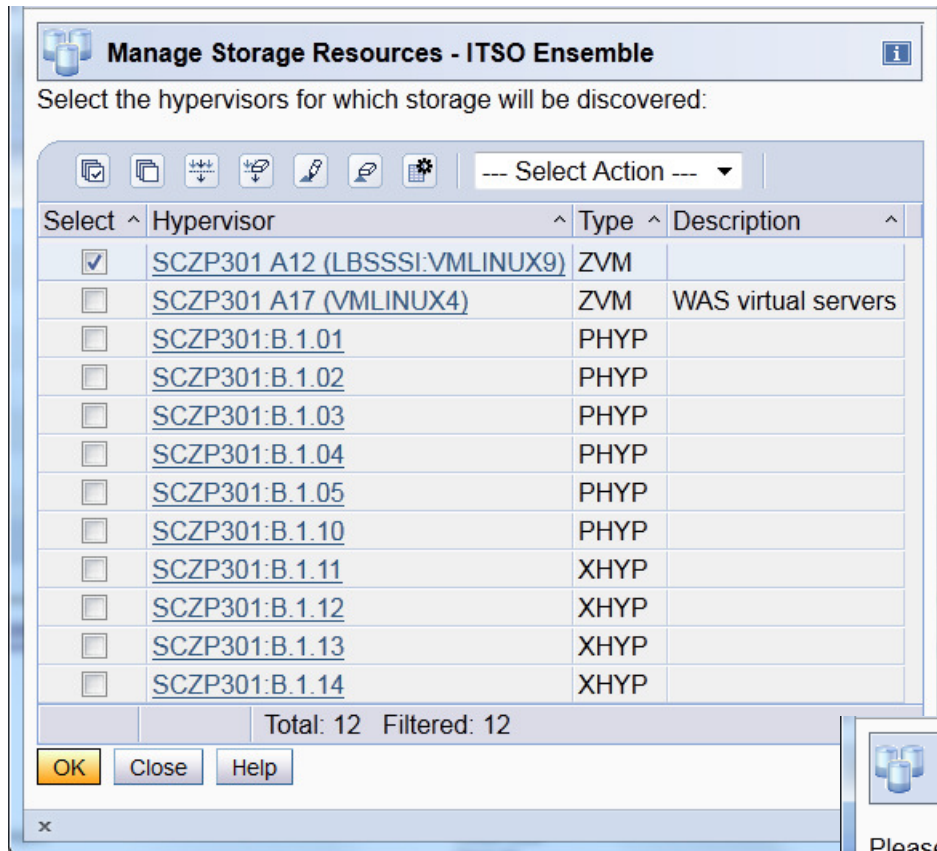
Storage Resources Virtual Disks

Select	Name	Hypervisor	Group
<input type="checkbox"/>	HTTPB2 disk2	SCZP301:B	B
<input type="checkbox"/>	HTTPB2 disk	SCZP301:B	B
<input type="checkbox"/>	HTTPG1 disk1	SCZP301:B	B
<input type="checkbox"/>	HTTPG1 disk	SCZP301:B	B
<input type="checkbox"/>	HTTPG2 disk1	SCZP301:B	B
<input type="checkbox"/>	HTTPG2 disk	SCZP301:B	B
<input type="checkbox"/>	HTTPS1 disk1	SCZP301:B	B
<input type="checkbox"/>	HTTPS1 disk	SCZP301:B	B
<input type="checkbox"/>	HTTPS2 disk1	SCZP301:B	B
<input type="checkbox"/>	HTTPS2 disk	SCZP301:B	B
<input type="checkbox"/>	LBSxBLUN2	SCZP301:B	B
<input type="checkbox"/>	LUN1	SCZP301:B	B
<input type="checkbox"/>	LX9880	SCZP301 A	cyl
<input type="checkbox"/>	LX9881	SCZP301 A	cyl \$3390\$
<input type="checkbox"/>	LX9883	SCZP301 A	cyl
<input type="checkbox"/>	LX9884	SCZP301 A12 (LBSSSI:VMLINUX9)	no ECKD 10017 cyl \$3390\$
<input type="checkbox"/>	LX9885	SCZP301 A12 (LBSSSI:VMLINUX9)	yes ECKD 10017 cyl
<input type="checkbox"/>	LXCB47	SCZP301 A17 (VMLINUX4)	no ECKD 3339 cyl
<input type="checkbox"/>	LXCB48	SCZP301 A17 (VMLINUX4)	no ECKD 3339 cyl
<input type="checkbox"/>	pBlade disk 2	SCZP301:B.1.05	yes FCP 15.0 GB
<input type="checkbox"/>	SAP AIX04 disk1	SCZP301:B.1.04	yes FCP 50.0 GB
<input type="checkbox"/>	WASB1 disk1	SCZP301 A17 (VMLINUX4)	yes FCP 30.0 GB
<input type="checkbox"/>	WASG1 disk1	SCZP301 A17 (VMLINUX4)	yes FCP 30.0 GB
<input type="checkbox"/>	WASS1 disk1	SCZP301 A17 (VMLINUX4)	yes FCP 30.0 GB
<input type="checkbox"/>	xBladeLun1	SCZP301:B.1.11	yes FCP 15.0 GB

Page 1 of 1 Max Page Size: 250 Total: 40 Filtered: 40 Displayed: 40

Close Help

# Ensemble Storage Resources - Discovery



Select the hypervisors for which storage will be discovered:


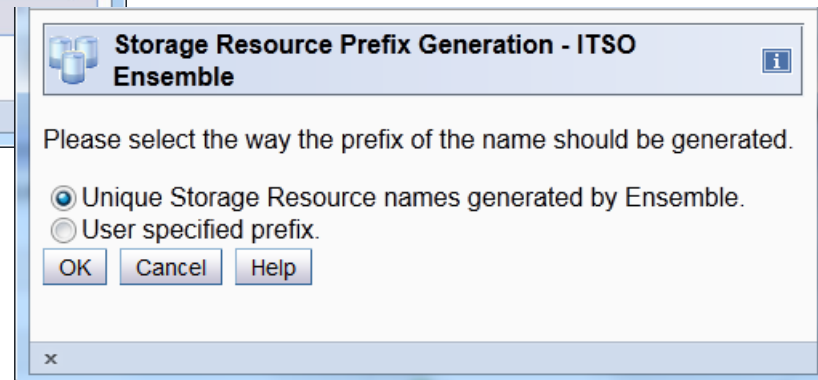
Select	Hypervisor	Type	Description
<input checked="" type="checkbox"/>	SCZP301 A12 (LBSSSI:VMLINUX9)	ZVM	
<input type="checkbox"/>	SCZP301 A17 (VMLINUX4)	ZVM	WAS virtual servers
<input type="checkbox"/>	SCZP301:B.1.01	PHYP	
<input type="checkbox"/>	SCZP301:B.1.02	PHYP	
<input type="checkbox"/>	SCZP301:B.1.03	PHYP	
<input type="checkbox"/>	SCZP301:B.1.04	PHYP	
<input type="checkbox"/>	SCZP301:B.1.05	PHYP	
<input type="checkbox"/>	SCZP301:B.1.10	PHYP	
<input type="checkbox"/>	SCZP301:B.1.11	XHYP	
<input type="checkbox"/>	SCZP301:B.1.12	XHYP	
<input type="checkbox"/>	SCZP301:B.1.13	XHYP	
<input type="checkbox"/>	SCZP301:B.1.14	XHYP	

Total: 12 Filtered: 12

OK Close Help

Select only the hypervisors for which discovery is required to be run.

Selecting extra will extend the discovery process

Please select the way the prefix of the name should be generated.

Unique Storage Resource names generated by Ensemble.  
 User specified prefix.

OK Cancel Help



## Ensemble Storage Resources - Discovery

- The discovered z/VM storage can be exported to a CSV file (export option shown on next slide)
- These entries can be edited if needed and then imported to define the storage (step not shown)
- In Unified Resource Manager discovery of SCSI LUNs under z/VM requires a free FCP device / path with no LUNs defined to it

# Ensemble Storage Resources - Discovery



Discovered Storage Resources - ITSO Ensemble

x Hyp/PowerVM z/VM

--- Select Action ---

- Select Action ---
- zVM Export
- Table Actions**
- Configure Columns

Name	Type	Hypervisor	Volume Serial	Device Number	Host WWPN	Target WWPN	LUN	Defined
ITSO-Ens	cyl	SCZP301:A12	DKCF47	CF47				new
ITSO-Ens	cyl	SCZP301:A12	LX9PG2	CF49				new
ITSO-Ensemble_SR_000003	0 cyl	SCZP301:A12	LX9882	9882				new
ITSO-Ensemble_SR_000004	0 cyl	SCZP301:A12	LX9SPL	CF41				new
ITSO-Ensemble_SR_000006	0 cyl	SCZP301:A12	ML1RES	9081				new
ITSO-Ensemble_SR_000007	0 cyl	SCZP301:A12	ML1P01	9083				new
ITSO-Ensemble_SR_000008	0 cyl	SCZP301:A12	DKCF48	CF48				new
ITSO-Ensemble_SR_000011	0 cyl	SCZP301:A12	LX9TD1	CF46				new
ITSO-Ensemble_SR_000013	0 cyl	SCZP301:A12	ML1S01	9082				new
ITSO-Ensemble_SR_000014	0 cyl	SCZP301:A12	LX9W01	CF43				new
ITSO-Ensemble_SR_000016	0 cyl	SCZP301:A12	LX9U1R	CF45				new
ITSO-Ensemble_SR_000017	0 cyl	SCZP301:A12	LBSRL1	9080				new
ITSO-Ensemble_SR_000018	0 cyl	SCZP301:A12	LX9W02	CF44				new
ITSO-Ensemble_SR_000019	0 cyl	SCZP301:A12	LX9RES	CF40				new
ITSO-Ensemble_SR_000020	0 cyl	SCZP301:A12	LX9PG1	CF42				new
LX9880	0 cyl	SCZP301:A12	LBSCM1	9880				new
LX9881	0 cyl	SCZP301:A12	LX9881	9881				new
LX9883	0 cyl	SCZP301:A12	LX9883	9883				new
LX9884	0 cyl	SCZP301:A12	LX9884	9884				new
LX9885	0 cyl	SCZP301:A12	0X0200	9885				new
Total: 20								

Close Help



# Storage Resources– Manually add Storage Resource



**Manage Storage Resources - ITSO Ensemble**

Select the hypervisor to which the new storage resource will be added:

--- Select Action ---

Select ^	Hypervisor ^	Type ^	Description ^
<input checked="" type="radio"/>	SCZP301 A12 (LBSSSI:VMLINUX9)	ZVM	
<input type="radio"/>	SCZP301 A17 (VMLINUX4)	ZVM	WAS virtual servers
<input type="radio"/>	SCZP301:B.1.01	PHYP	
<input type="radio"/>	SCZP301:B.1.02	PHYP	
<input type="radio"/>	SCZP301:B.1.03	PHYP	
<input type="radio"/>	SCZP301:B.1.04	PHYP	
<input type="radio"/>	SCZP301:B.1.05	PHYP	
<input type="radio"/>	SCZP301:B.1.10	PHYP	
<input type="radio"/>	SCZP301:B.1.11	XHYP	
<input type="radio"/>	SCZP301:B.1.12	XHYP	
<input type="radio"/>	SCZP301:B.1.13	XHYP	
<input type="radio"/>	SCZP301:B.1.14	XHYP	

Total: 12 Filtered: 12

OK Close Help

# Storage Resources – Manually add Storage Resource



**Add Storage Resource to in Ensemble - A12**

Name: \* NEWVOL


Size: \* 3339 cylinders:

Device Number: \* CF49

Volser: \* NEWVOL

Description:

OK Cancel Help



**Add Storage Resource to Hypervisor - A12**

The storage resource NEWVOL was successfully added to hypervisor SCZP301 A12 (LBSSSI:VMLINUX9).

OK

# Ensemble Storage Resources – Storage Group



- Guest storage allocations not using groups will result in a devno mdisk allocation
- To allocate less than a full pack you need to use a “Storage Group”
- Three predefined groups exist:
  - 3390
  - 3380
  - FCP

Manage Storage Resources - ITSO Ensemble

Storage Resources Virtual Disks

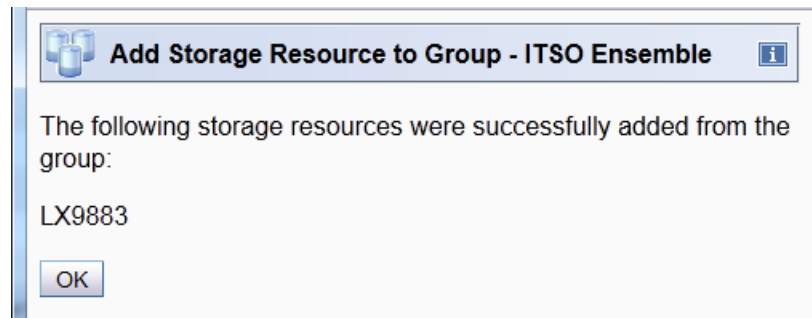


--- Select Action ---

Select	Name	Hypervisor	Group
<input type="checkbox"/>	HTTPB2 disk2	SCZP301:B	B
<input type="checkbox"/>	HTTPB2 disk	SCZP301:B	B
<input type="checkbox"/>	HTTPG1 disk1	SCZP301:B	B
<input type="checkbox"/>	HTTPG1 disk	SCZP301:B	B
<input type="checkbox"/>	HTTPG2 disk1	SCZP301:B	B
<input type="checkbox"/>	HTTPG2 disk	SCZP301:B	B
<input type="checkbox"/>	HTTPS1 disk1	SCZP301:B	B
<input type="checkbox"/>	HTTPS1 disk	SCZP301:B	B
<input type="checkbox"/>	HTTPS2 disk1	SCZP301:B	B
<input type="checkbox"/>	HTTPS2 disk	SCZP301:B	B
<input type="checkbox"/>	LBSxBLUN2	SCZP301:B	B
<input type="checkbox"/>	LUN1	SCZP301:B	B
<input type="checkbox"/>	LX9881	SCZP301 A	cyl \$3390\$
<input checked="" type="checkbox"/>	LX9883	SCZP301 A	cyl
<input type="checkbox"/>	LX9884	SCZP301 A	cyl \$3390\$
<input type="checkbox"/>	LX9885	SCZP301 A12 (LBSSSI.VMLINUX9)	yes ECKD 10017 cyl
<input type="checkbox"/>	LXCB47	SCZP301 A17 (VMLINUX4)	no ECKD 3339 cyl
<input type="checkbox"/>	LXCB48	SCZP301 A17 (VMLINUX4)	no ECKD 3339 cyl
<input type="checkbox"/>	NEWVOL	SCZP301 A12 (LBSSSI.VMLINUX9)	no ECKD 3339 cyl
<input type="checkbox"/>	pBlade disk 2	SCZP301:B.1.05	yes FCP 15.0 GB
<input type="checkbox"/>	SAP AIX04 disk1	SCZP301:B.1.04	yes FCP 50.0 GB
<input type="checkbox"/>	WASB1 disk1	SCZP301 A17 (VMLINUX4)	yes FCP 30.0 GB
<input type="checkbox"/>	WASG1 disk1	SCZP301 A17 (VMLINUX4)	yes FCP 30.0 GB
<input type="checkbox"/>	WASS1 disk1	SCZP301 A17 (VMLINUX4)	yes FCP 30.0 GB
<input type="checkbox"/>	xBladeLun1	SCZP301:B.1.11	yes FCP 15.0 GB

Page 1 of 1 Max Page Size: 250 Total: 40 Filtered: 40 Displayed: 40

Close Help

# Ensemble Storage Resources– Storage Group



# Agenda

1 Why make your Linux guest part of the Ensemble

2 Requirements and Planning Considerations

3 Defining Resources for your Virtual Servers

Ensemble Storage Resources – Discovery, Definition, and Assignment



Ensemble Virtual Networking Resources and Virtual Switch Management

4 Defining a New the Virtual Server or Managing an Existing Server

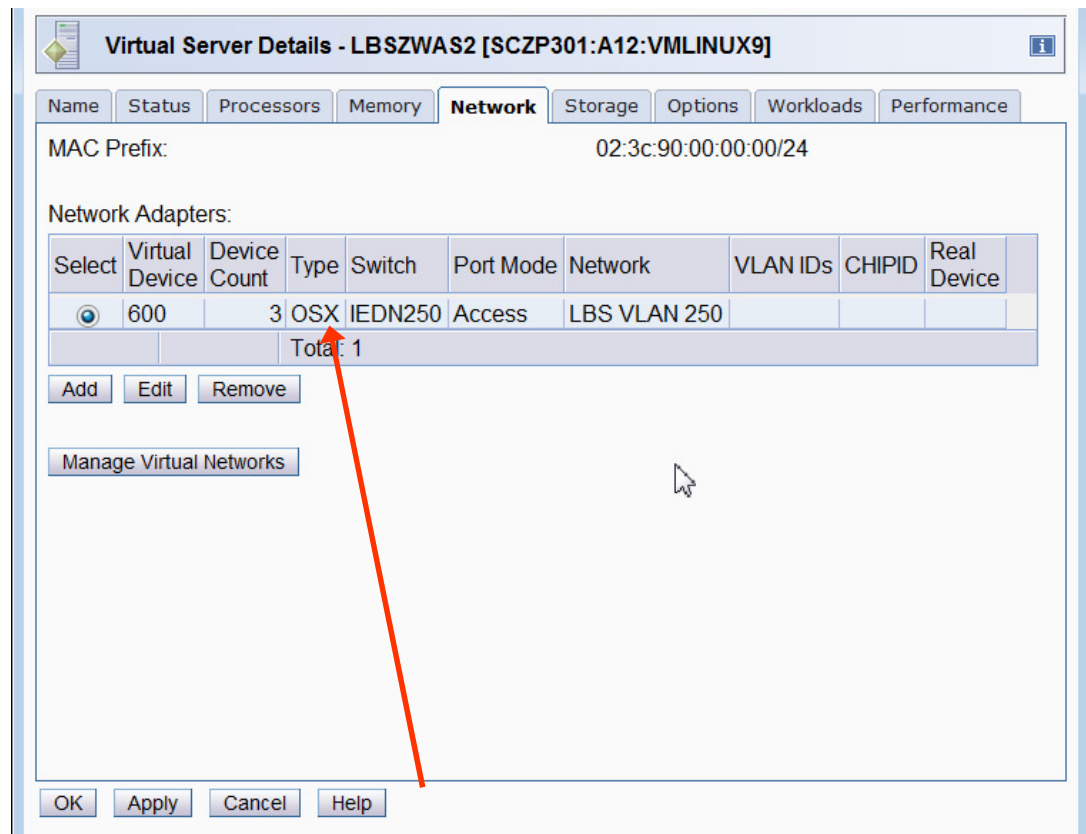
5 Linux Installation Considerations

6 Installing/Enabling the GPMP

7 Enabling ARM

# Linux Network 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 OSX device without kernel support
- Latest levels (such as SLES 11 SP2) have full Yast and installer system support as well
- VSWITCH can “simulate” OSD for older systems that need OSX access
  - 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
- Linux in a native LPAR is also supported
  - No GPMP or workload management



Virtual Server Details - LBSZWAS2 [SCZP301:A12:VMLINUX9]

MAC Prefix: 02:3c:90:00:00:00/24

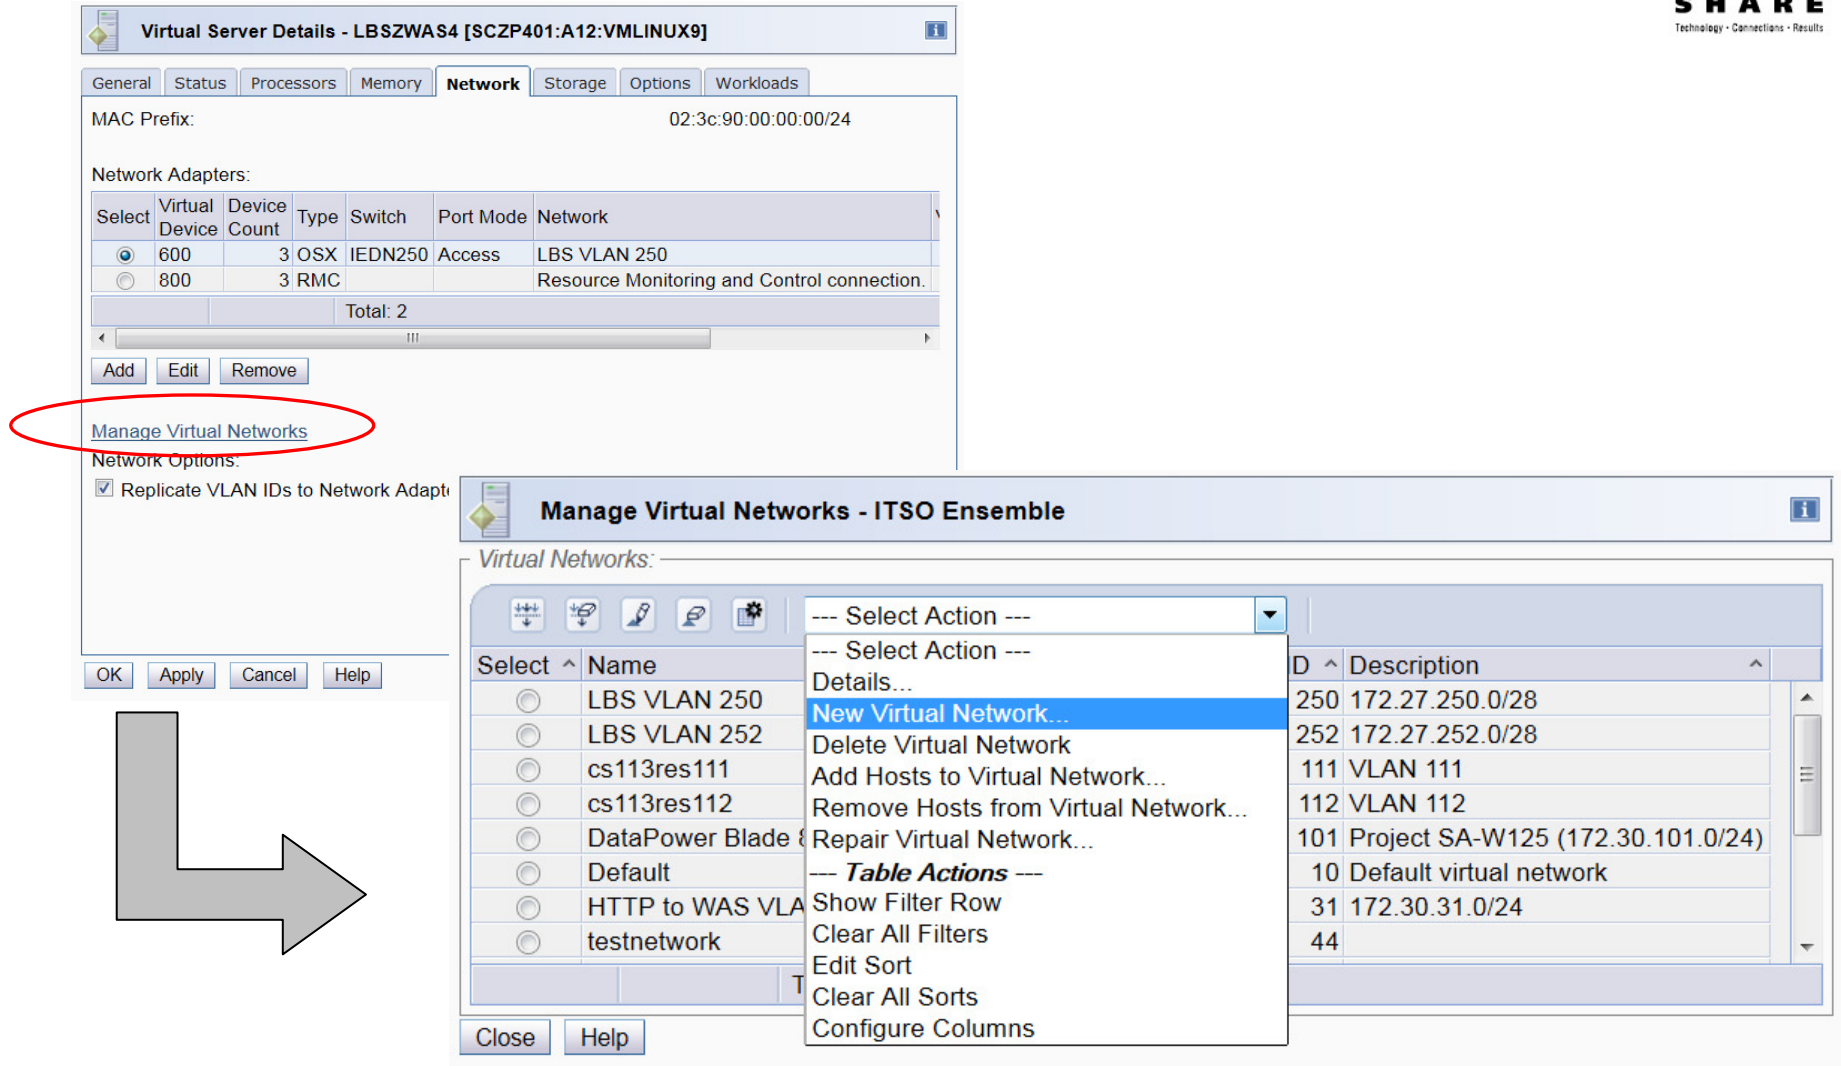
Network Adapters:

Select	Virtual Device	Device Count	Type	Switch	Port Mode	Network	VLAN IDs	CHIPID	Real Device
<input checked="" type="radio"/>	600	3	OSX	IEDN250	Access	LBS VLAN 250			
			Total: 1						

Buttons: Add, Edit, Remove, Manage Virtual Networks, OK, Apply, Cancel, Help



# Creating an Ensemble Virtual Network



**Virtual Server Details - LBSZWAS4 [SCZP401:A12:VMLINUX9]**

General | Status | Processors | Memory | **Network** | Storage | Options | Workloads

MAC Prefix: 02:3c:90:00:00:24

Network Adapters:

Select	Virtual Device	Device Count	Type	Switch	Port Mode	Network
<input checked="" type="radio"/>	600	3	OSX	IEDN250	Access	LBS VLAN 250
<input type="radio"/>	800	3	RMC			Resource Monitoring and Control connection.
Total: 2						

Buttons: Add | Edit | Remove

[Manage Virtual Networks](#)

Network Options:  
 Replicate VLAN IDs to Network Adapters

Buttons: OK | Apply | Cancel | Help

---

**Manage Virtual Networks - ITSO Ensemble**

Virtual Networks:

Select	Name	Description
<input type="radio"/>	LBS VLAN 250	250 172.27.250.0/28
<input type="radio"/>	LBS VLAN 252	252 172.27.252.0/28
<input type="radio"/>	cs113res111	111 VLAN 111
<input type="radio"/>	cs113res112	112 VLAN 112
<input type="radio"/>	DataPower Blade S	101 Project SA-W125 (172.30.101.0/24)
<input type="radio"/>	Default	10 Default virtual network
<input type="radio"/>	HTTP to WAS VLA	31 172.30.31.0/24
<input type="radio"/>	testnetwork	44

Context Menu (over LBS VLAN 250):

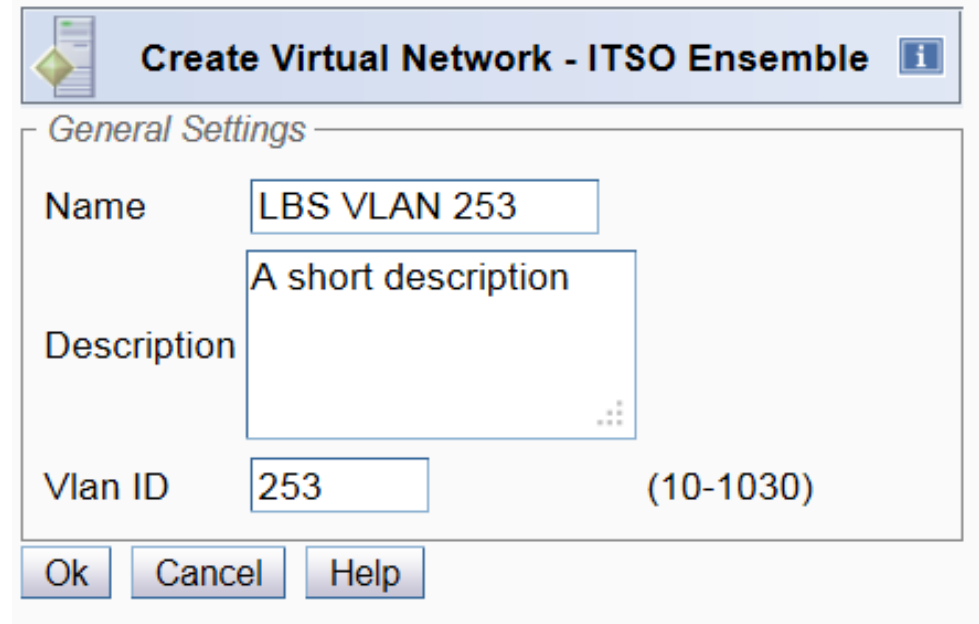
- Select Action ---
- Select Action ---
- Details...
- New Virtual Network...**
- Delete Virtual Network
- Add Hosts to Virtual Network...
- Remove Hosts from Virtual Network...
- Repair Virtual Network...
- *Table Actions* ---
- Show Filter Row
- Clear All Filters
- Edit Sort
- Clear All Sorts
- Configure Columns

Buttons: Close | Help



# Creating an Ensemble Virtual Network

- Defining an Ensemble virtual network is simple.
- Supply the name, description, and VLAN ID number
- All required changes to the top of rack switches are transparently handled for you



**Create Virtual Network - ITSO Ensemble** ⓘ

*General Settings*

Name

Description

Vlan ID  (10-1030)

Ok Cancel Help

# Ensemble Virtual Networking – Virtual Switches



Ensemble Management > ITSO Ensemble > Members > SCZP301

Virtual Servers | Hypervisors | Blades | Topology

Filter: [ ] Tasks: [ ] Views: [ ]

Select	Name	Status	Processors	Memory (MB)	Type	Auto Start	Shutdown Timeout (s)
<input checked="" type="checkbox"/>	A12 (LBSSSI:VMLINUX9)	Operating			z/VM	-	300
<input type="checkbox"/>	LBSOVS	Operating	1	1,500	z/VM		
<input type="checkbox"/>	LBSZWAS1	Not Activated	1	1,500	z/VM		
<input type="checkbox"/>	LBSZWAS3	Not Activated			z/VM		
<input type="checkbox"/>	LNXMNT	Not Activated	1	1,024	z/VM		
<input type="checkbox"/>	A17 (VMLINUX4)	Operating			z/VM	-	300
<input type="checkbox"/>	B.1.01	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.02	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.03	Operating	1	32,768	PowerVM	✓	300

Max Page Size: 30 Total: 16 Filtered: 16 Selected: 1

Tasks: A12 (LBSSSI:VMLINUX9)

- Image Details
  - Toggle Lock
- Daily
  - Activate
  - Deactivate
  - Grouping
  - Hardware Messages
  - Operating System Messages
  - Reset Normal
- Recovery
  - Access Removable Media
  - Integrated 3270 Console
  - Integrated ASCII Console
  - Load
  - Load from Removable Media or Server
  - PSW Restart
  - Reset Clear
  - Start All
  - Stop All
- Service
- Operational Customization
  - Configure Channel Path On/Off
  - Customize/Delete Activation Profiles
  - Logical Processor Add
  - View Activation Profiles
- Configuration
  - Choose z/VM Virtual Servers to Manage
  - manage Storage Resources
  - Manage Virtual Switches**
  - New Virtual Server

'bonsai/action/T8243#

# Ensemble Virtual Networking – Virtual Switches

**Manage Virtual Switches - SCZP301:A12**

--- Select Action ---

Select	Name	Type	VLAN	Port Type	GVRP	Layer Mode	Router	Queue Size	IP Timeout
<input type="radio"/>	IEDN250	IEDN	AWARE	ACCESS		ETH		8	5
<input type="radio"/>	IEDN251	IEDN	AWARE	ACCESS		ETH		8	5
<input type="radio"/>	IEDN252	IEDN	AWARE	ACCESS		ETH		8	5
<input type="radio"/>	IEDN250B	IEDN	AWARE	ACCESS		ETH		8	5
<input type="radio"/>	VSWITCH1	QDIO	UNAWARE		false	IP	NONROUTER	8	5
Total: 5									

OK Cancel Help

**Manage Virtual Switches - SCZP301:A12**

--- Select Action ---

Select	Name	Type	VLAN	Port Type	GVRP	Layer Mode	Router	Queue Size	IP Timeout
<input type="radio"/>	IEDN250	IEDN	AWARE	ACCESS		ETH		8	5
<input type="radio"/>	IEDN251	IEDN	AWARE	ACCESS		ETH		8	5
<input type="radio"/>	IEDN252	IEDN	AWARE	ACCESS		ETH		8	5
<input type="radio"/>	IEDN250B	IEDN	AWARE	ACCESS		ETH		8	5
<input type="radio"/>	VSWITCH1	QDIO	UNAWARE		false	IP	NONROUTER	8	5
Total: 5									

OK Cancel Help

# Ensemble Virtual Networking – Virtual Switches



Creating a Virtual Switch from the Unified Resource Manager is easy

Select “ETH” for a layer 2 or “IP” for layer 3 virtual switch. (we would typically select ETH for IEDN access)

The OSX CHPID and device numbers are the remaining required bits of information

**Create IEDN Virtual Switch - SCZP301:A12**

Name: \* IEDN300

Layer Mode: \* ETH

Router: NONROUTER

Queue size (Mbytes): \* 8

IP timeout: \* 5

Uplinks:

	OSX CHPID	Device Number
Uplink 1:	OSX 1.18	2300
Uplink 2:	OSX 1.19	2400
Uplink 3:	OSX 1.18	

Connect uplinks:

Bridge Port:

Connect Port:

Device Number:

Primary Port:

MTU Enforcement: EXTERNAL

Controller:

Use any available:

Name	Available	VDEV Range	IP	ETHERNET	VLAN_ARP	GVRP	LINKAGG	ISOLATION
DTCENS1	true	0600-F000	true	true	true	true	true	true
DTCENS2	true	0600-F000	true	true	true	true	true	true
		Total: 2						

OK Cancel Help

# Agenda

- 1 Why make your Linux guest part of the Ensemble
- 2 Requirements and Planning Considerations
- 3 Defining Resources for your Virtual Servers
- 4 Defining a New the Virtual Server or Managing an Existing Server
- 5 Linux Installation Considerations
- 6 Installing/Enabling the GPMP
- 7 Enabling ARM



# New Virtual Server Guest Definition



Ensemble Management > ITSO Ensemble > Members > SCZP301

Virtual Servers | Hypervisors | Blades | Topology

Filter [ ] Tasks [ ] Views [ ]

Sel.	Name	Status	Processors	Memory (MB)	Type	Auto Start	Shutdown Timeout (s)
<input checked="" type="checkbox"/>	A12 (LBSSSI:VMLINUX9)	Operating			z/VM	—	300
<input type="checkbox"/>	LBSOVS	Operating	1	1,500	z/VM		
<input type="checkbox"/>	LBSZWAS1	Not Activated	1	1,500	z/VM		
<input type="checkbox"/>	LBSZWAS3	Not Activated			z/VM		
<input type="checkbox"/>	LNXMNT	Not Activated	1	1,024	z/VM		
<input type="checkbox"/>	A17 (VMLINUX4)	Operating			z/VM	—	300
<input type="checkbox"/>	B.1.01	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.02	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.03	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.04	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.05	Operating	1	32,768	PowerVM	✓	300

Max Page Size: 30 | Total: 16 | Filtered: 16 | Selected: 1

Tasks: A12 (LBSSSI:VMLINUX9)

- Image Details
  - Toggle Lock
- Daily
  - Activate
  - Deactivate
  - Grouping
  - Hardware Messages
  - Operating System Messages
  - Reset Normal
- Recovery
  - Access Removable Media
  - Integrated 3270 Console
  - Integrated ASCII Console
  - Load
  - Load from Removable Media or Server
  - PSW Restart
  - Reset Clear
  - Start All
  - Stop All
- Service
- Operational Customization
  - Configure Channel Path On/Off
  - Customize/Delete Activation Profiles
  - Logical Processor Add
  - View Activation Profiles
- Configuration
  - Choose z/VM Virtual Servers to Manage
  - Manage Storage Resources
  - Manage Virtual Switches
  - [New Virtual Server](#)

/bonsai/action/T8243#

# New Virtual Server Guest Definition



**New Virtual Server - SCZP301:A12**

- ✓ Welcome
- **Enter Name**
- Assign Processors
- Specify Memory
- Add Network
- Add Storage
- Specify Options
- Select Workloads
- Performance Management
- Summary

**Enter Name**  
Enter in a name and description for the virtual server.

Hypervisor name: A12  
Hypervisor type: z/VM

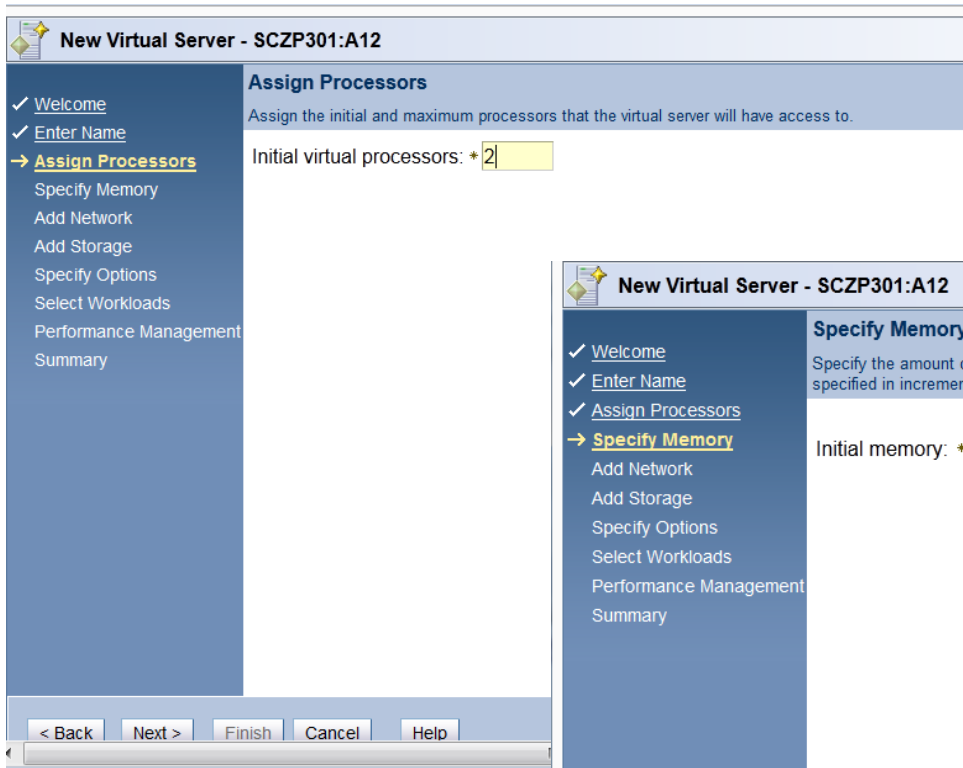
Name: \*

Description:

< Back   Next >   Finish   Cancel   Help



# New Virtual Server Guest Definition



**New Virtual Server - SCZP301:A12**

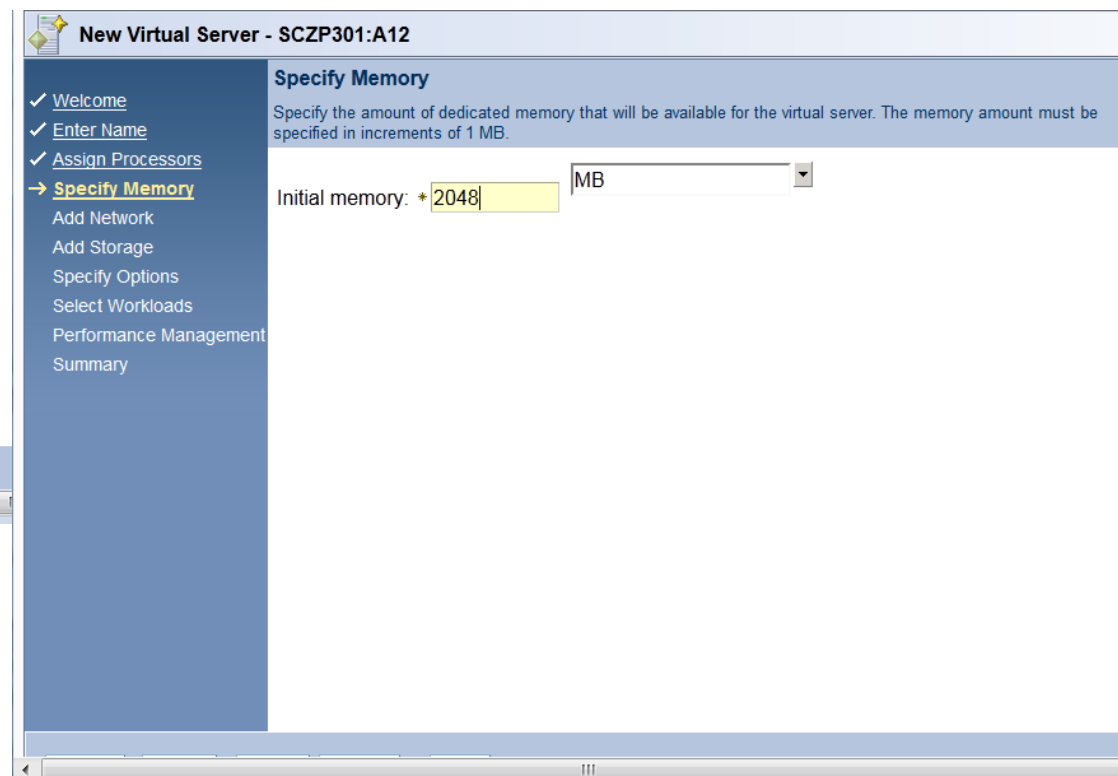
- ✓ Welcome
- ✓ Enter Name
- **Assign Processors**
- Specify Memory
- Add Network
- Add Storage
- Specify Options
- Select Workloads
- Performance Management
- Summary

**Assign Processors**  
Assign the initial and maximum processors that the virtual server will have access to.

Initial virtual processors: \* 2

< Back   Next >   Finish   Cancel   Help

Additional CPU and memory options are available for configuration after the virtual server has been defined



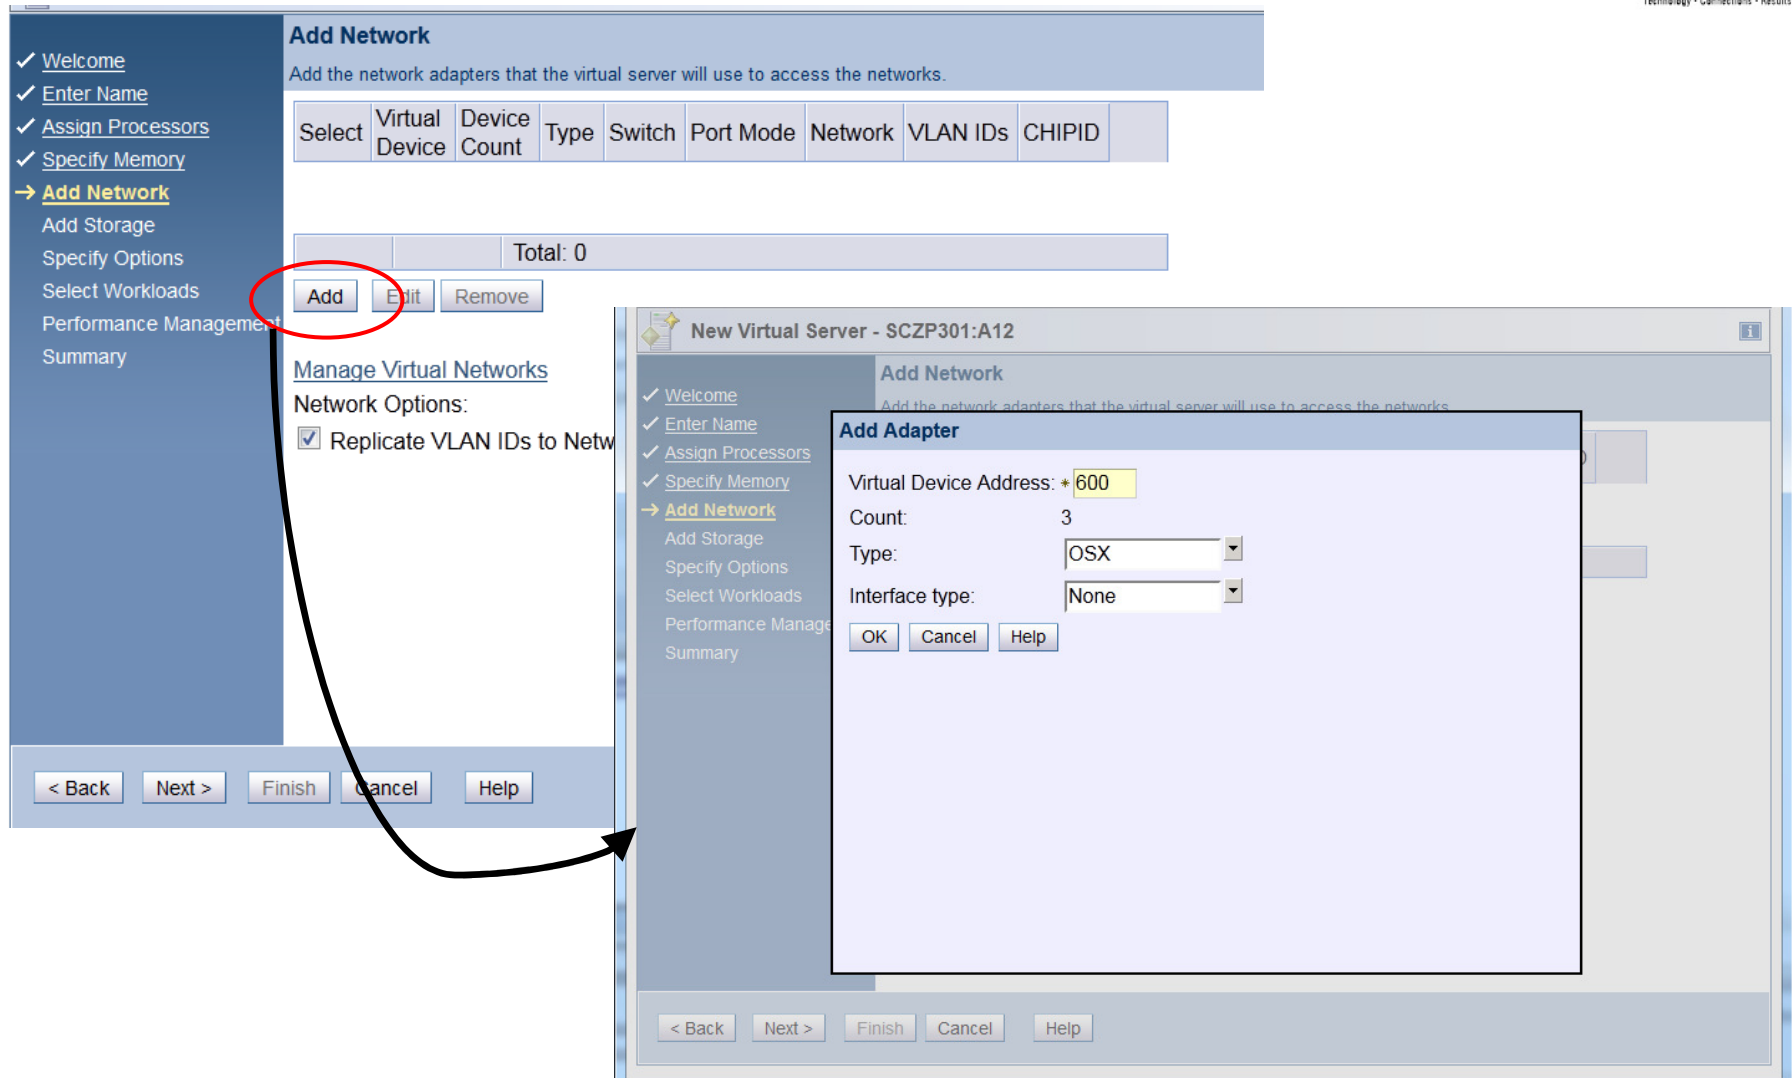
**New Virtual Server - SCZP301:A12**

- ✓ Welcome
- ✓ Enter Name
- ✓ Assign Processors
- **Specify Memory**
- Add Network
- Add Storage
- Specify Options
- Select Workloads
- Performance Management
- Summary

**Specify Memory**  
Specify the amount of dedicated memory that will be available for the virtual server. The memory amount must be specified in increments of 1 MB.

Initial memory: \* 2048 MB

# New Virtual Server Guest Definition



**Add Network**  
Add the network adapters that the virtual server will use to access the networks.

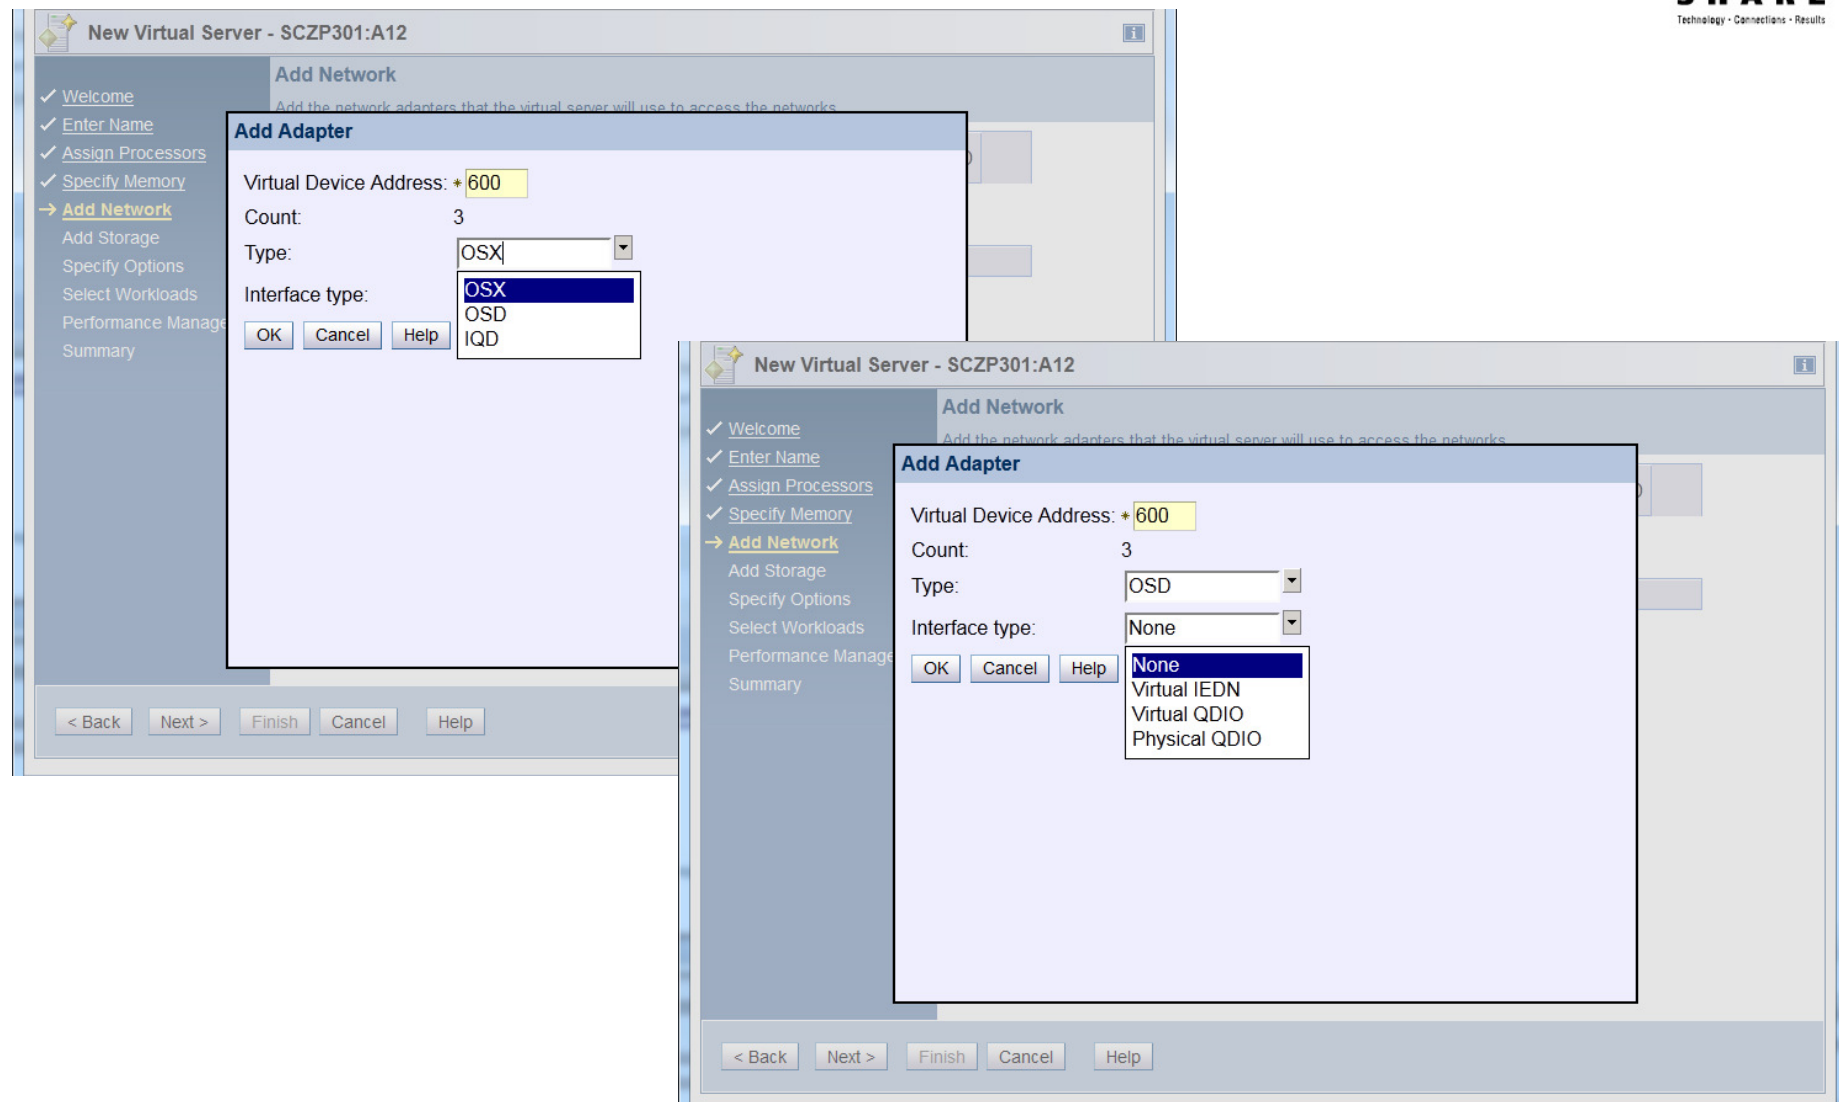
Select	Virtual Device	Device Count	Type	Switch	Port Mode	Network	VLAN IDs	CHIPID
Total: 0								

Replicate VLAN IDs to Network

**Add Adapter**

Virtual Device Address: \*600  
 Count: 3  
 Type: OSX  
 Interface type: None

# New Virtual Server Guest Definition

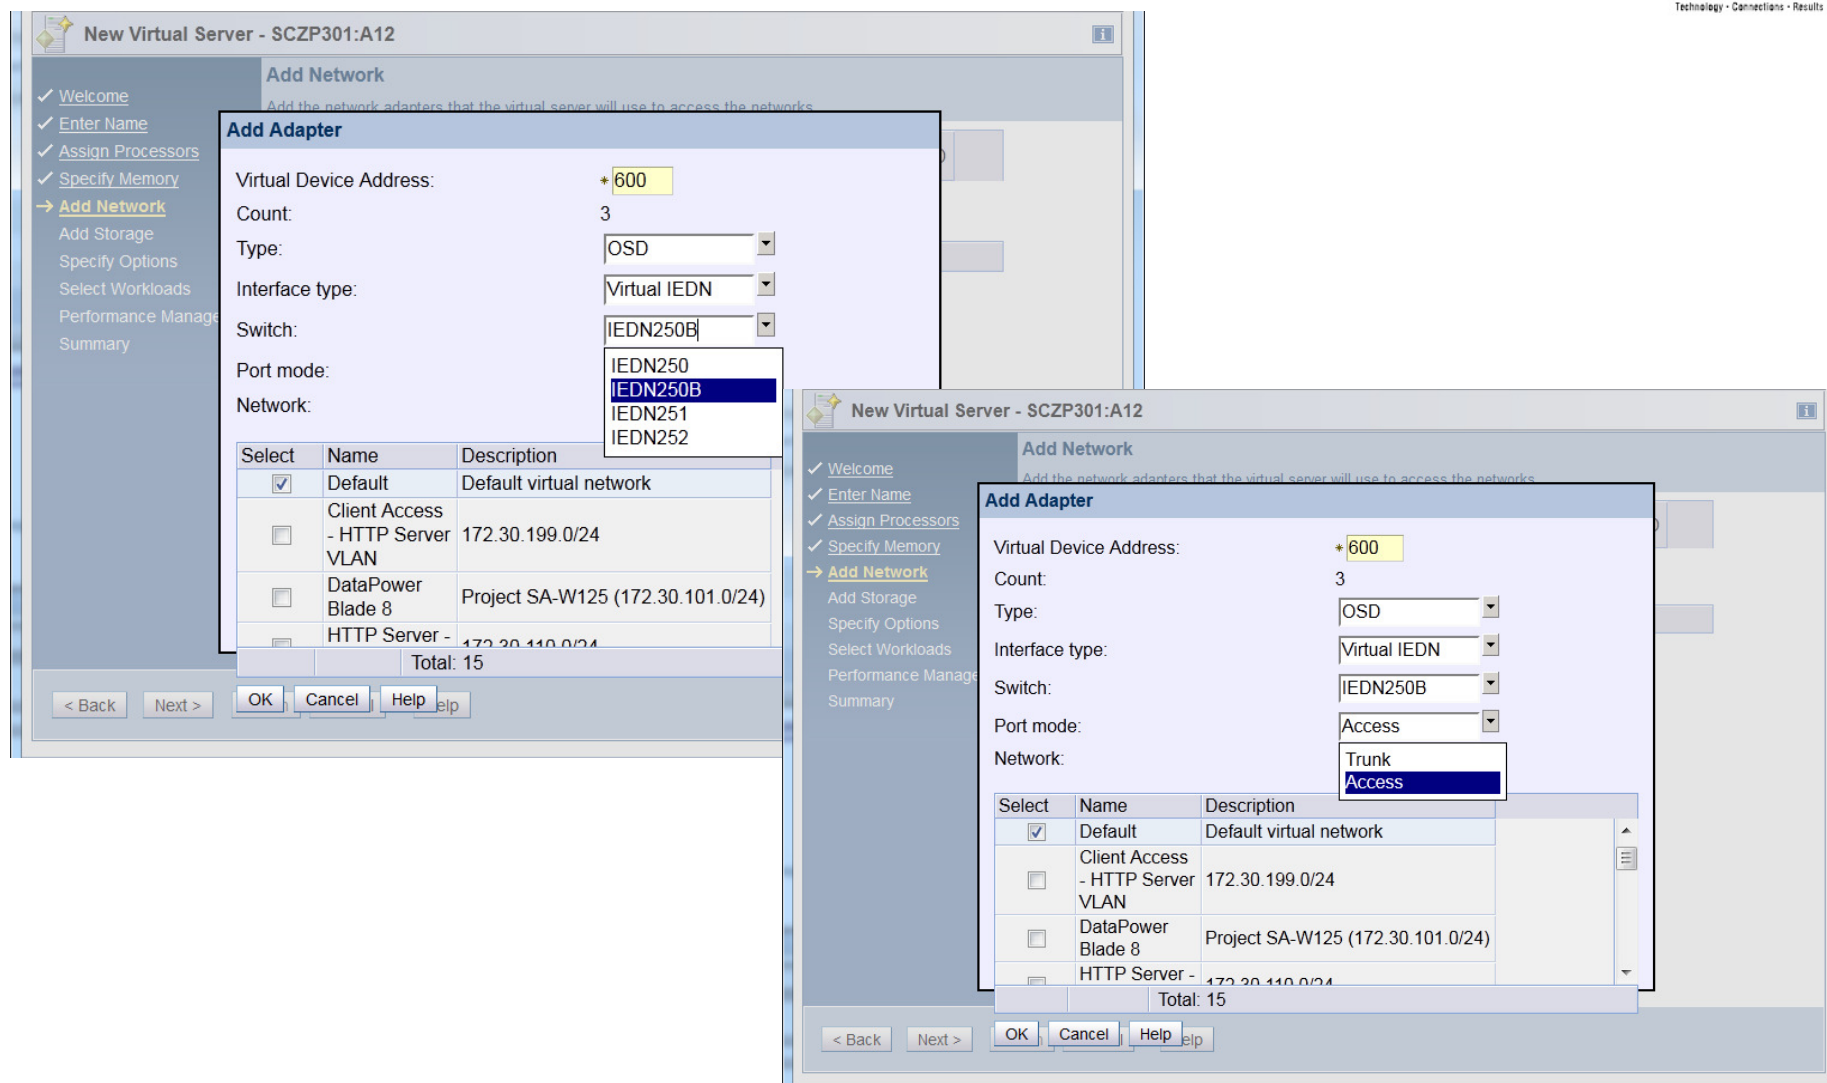


The image displays two screenshots of the 'New Virtual Server - SCZP301:A12' wizard, specifically the 'Add Network' step. Both screenshots show an 'Add Adapter' dialog box with the following fields:

- Virtual Device Address: \*600
- Count: 3
- Type: OSD
- Interface type: [Dropdown]

In the left screenshot, the 'Interface type' dropdown is open, showing 'OSX' as the selected option. In the right screenshot, the 'Interface type' dropdown is open, showing 'None' as the selected option. The background wizard window shows a progress list on the left with 'Add Network' highlighted, and navigation buttons ('< Back', 'Next >', 'Finish', 'Cancel', 'Help') at the bottom.

# New Virtual Server Guest Definition



**Add Adapter**

Virtual Device Address: +600  
 Count: 3  
 Type: OSD  
 Interface type: Virtual IEDN  
 Switch: IEDN250B  
 Port mode: IEDN250B  
 Network: IEDN250B

Select	Name	Description
<input checked="" type="checkbox"/>	Default	Default virtual network
<input type="checkbox"/>	Client Access - HTTP Server VLAN	172.30.199.0/24
<input type="checkbox"/>	DataPower Blade 8	Project SA-W125 (172.30.101.0/24)
<input type="checkbox"/>	HTTP Server -	172.30.110.0/24

Total: 15

< Back Next > OK Cancel Help

**Add Adapter**

Virtual Device Address: +600  
 Count: 3  
 Type: OSD  
 Interface type: Virtual IEDN  
 Switch: IEDN250B  
 Port mode: Access  
 Network: Access

Select	Name	Description
<input checked="" type="checkbox"/>	Default	Default virtual network
<input type="checkbox"/>	Client Access - HTTP Server VLAN	172.30.199.0/24
<input type="checkbox"/>	DataPower Blade 8	Project SA-W125 (172.30.101.0/24)
<input type="checkbox"/>	HTTP Server -	172.30.110.0/24

Total: 15

< Back Next > OK Cancel Help

# New Virtual Server Guest Definition

New Virtual Server - SCZP301:A12

Add Network

Add the network adapters that the virtual server will use to access the networks

**Add Adapter**

Virtual Device Address: \*600  
 Count: 3  
 Type: OSD  
 Interface type: Virtual IEDN  
 Switch: IEDN250B  
 Port mode: Access

Network:

Select	Name	Description
<input type="checkbox"/>	Default	Default virtual network
<input type="checkbox"/>	Client Access - HTTP Server VLAN	172.30.199.0/24
<input type="checkbox"/>	DataPower Blade 8	Project SA-W125 (172.30.101.0/24)
<input type="checkbox"/>	HTTP Server - 172.30.110.0/24	
Total: 15		

< Back   Next >   OK   Cancel   Help <sub>elp</sub>

New Virtual Server - SCZP301:A12

Add Network

Add the network adapters that the virtual server will use to access the networks

**Add Adapter**

Virtual Device Address: \*600  
 Count: 3  
 Type: OSD  
 Interface type: Virtual IEDN  
 Switch: IEDN250B  
 Port mode: Access

Network:

Select	Name	Description
<input type="checkbox"/>	HTTP Server - WAS VLAN	172.30.110.0/24
<input type="checkbox"/>	HTTP to WAS VLAN31	172.30.31.0/24
<input checked="" type="checkbox"/>	LBS VLAN 250	172.27.250.0/28
<input type="checkbox"/>	LBS VLAN 251	172.27.251.0/28
Total: 15		

< Back   Next >   OK   Cancel   Help <sub>elp</sub>

# New Virtual Server Guest Definition



**New Virtual Server - SCZP301:A12**

- ✓ Welcome
- ✓ Enter Name
- ✓ Assign Processors
- ✓ Specify Memory
- **Add Network**
- Add Storage
- Specify Options
- Select Workloads
- Performance Management
- Summary

### Add Network

Add the network adapters that the virtual server will use to access the networks.

Select	Virtual Device	Device Count	Type	Switch	Port Mode	Network	VLAN IDs	CHIPID
<input checked="" type="radio"/>	600	3	OSD	IEDN250B	Access	LBS VLAN 250	-1	
			Total: 1					

### Manage Virtual Networks

Network Options:

Replicate VLAN IDs to Network Adapters

< Back   **Next >**   Finish   Cancel   Help



# New Virtual Server Guest Definition

**New Virtual Server - SCZP301:A12**

- ✓ Welcome
- ✓ Enter Name
- ✓ Assign Processors
- ✓ Specify Memory
- ✓ Add Network
- **Add Storage**
- Specify Options
- Select Workloads
- Performance Management
- Summary

< Back   **Next >**   Finish   Cancel   Help

**Add Storage**

Add the storage drives that the virtual server will use to access the storage resources.

Select	Device	Name	Description	Type	Resource Name	Mode	Size
Total: 0							

Add   Edit   Remove

**Add Drive**

Device Address: \* 200

Name: \* OSVOL

Description: Core Linux OS

Type: Fullpack

Storage resource:

Select	Name	Description	Size
<input checked="" type="radio"/>	LX9880		10017 cyl
<input type="radio"/>	LX9883		10018 cyl
<input type="radio"/>	NEWVOL		3339 cyl
Total: 3			

Mode: Read-Write

Read Password:

Confirm Read Password:

Write Password:

Confirm Write Password:

Multi-Write Password:

Confirm Multi-Write Password:

< Back   Next >

# New Virtual Server Guest Definition



**New Virtual Server - SCZP301:A12**

- ✓ [Welcome](#)
- ✓ [Enter Name](#)
- ✓ [Assign Processors](#)
- ✓ [Specify Memory](#)
- ✓ [Add Network](#)
- **[Add Storage](#)**
- [Specify Options](#)
- [Select Workloads](#)
- [Performance Management](#)
- [Summary](#)

### Add Storage

Add the storage drives that the virtual server will use to access the storage resources.

Select	Device	Name	Description	Type	Resource Name	Mode	Size
<input checked="" type="radio"/>	200	OSVOL	Linux OS	Fullpack	LX9883	Read-Write	10018 cyl
Total: 1							

[Manage Storage Resources](#)

< Back   Next >   Finish   Cancel   Help

# New Virtual Server Guest Definition



New Virtual Server - SCZP301:A12

**Specify Options**

Choose the boot source for your virtual server.

Password: [.....]

Confirm password: [.....]

Privilege classes: \*G

IPL device: \*200

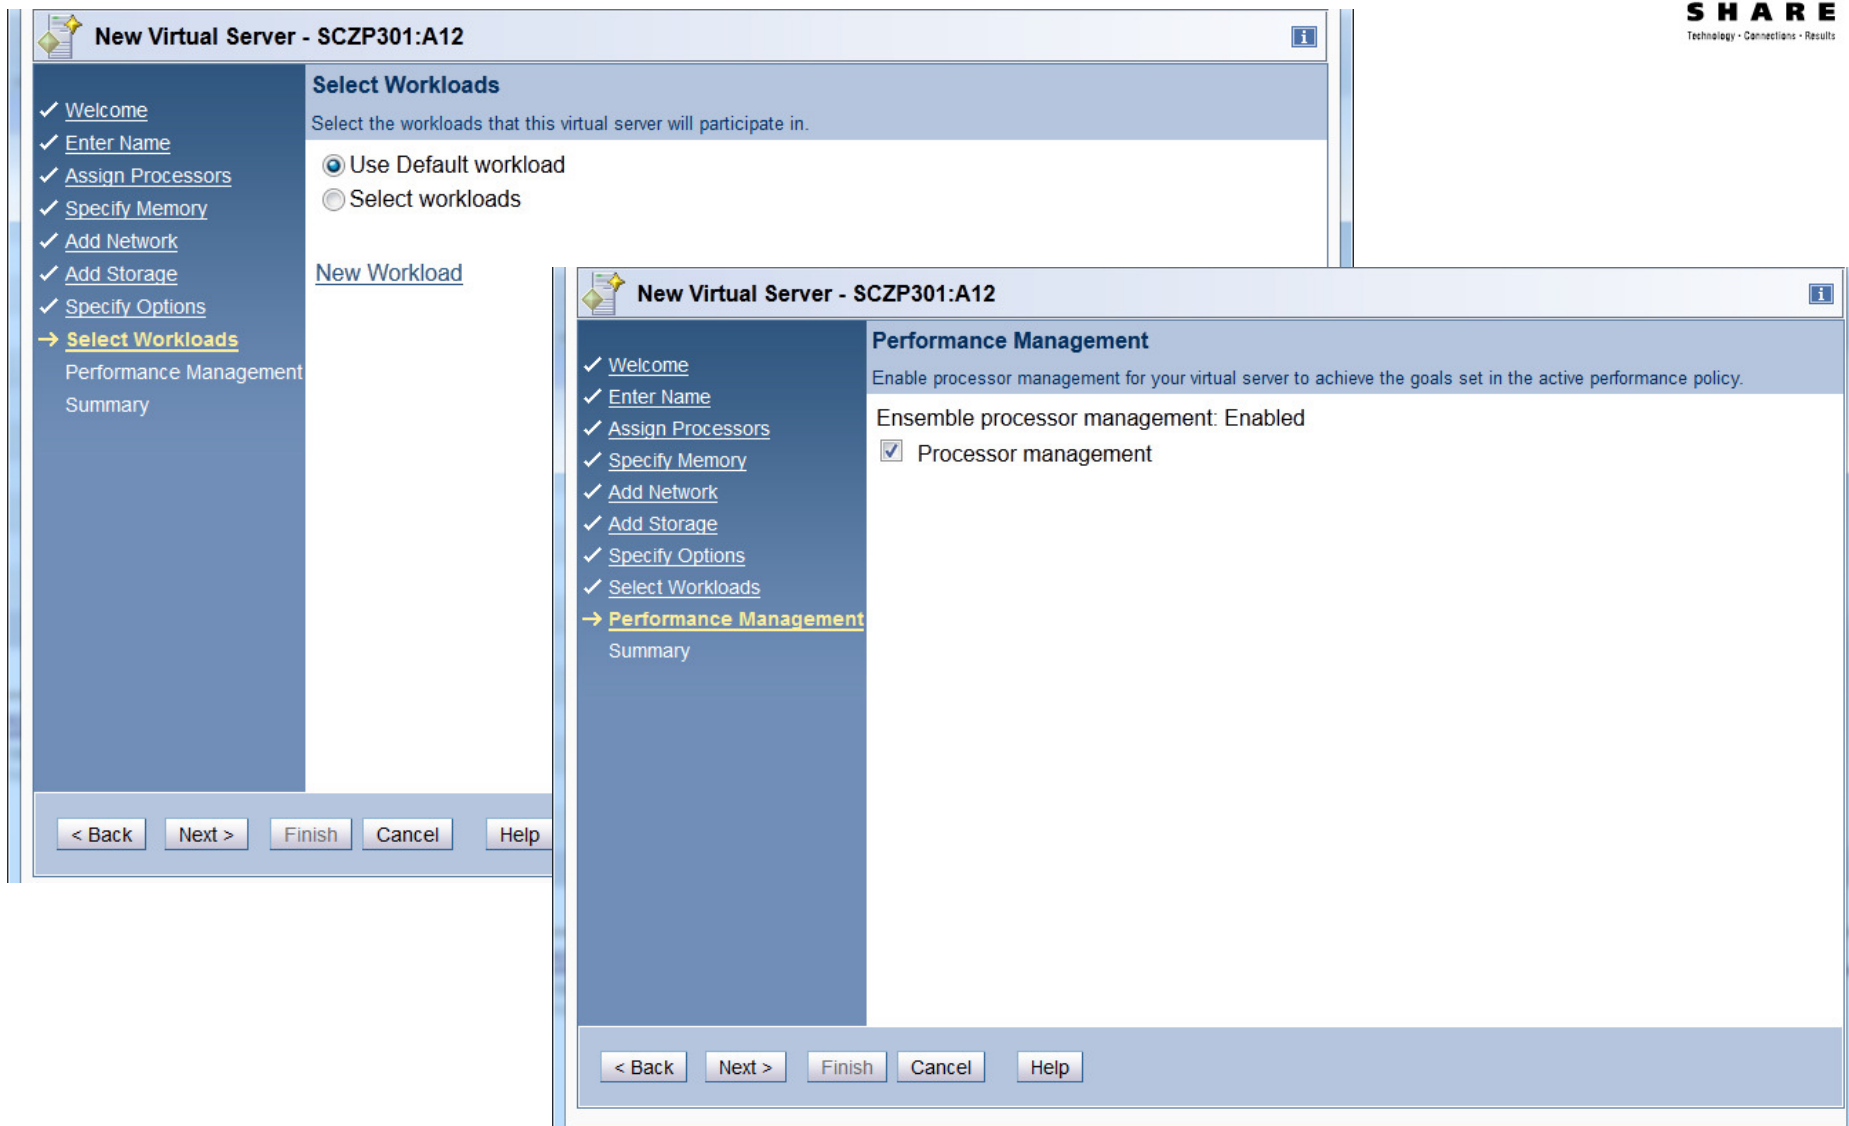
IPL load parameters: [ ]

IPL parameters: AUTOOCR

Enable GPMP support

< Back   Next >   Finish   Cancel   Help

# New Virtual Server Guest Definition



The screenshot displays two overlapping windows from the 'New Virtual Server - SCZP301:A12' wizard. The background window is at the 'Select Workloads' step, and the foreground window is at the 'Performance Management' step.

**Background Window: Select Workloads**

- Navigation pane: Welcome, Enter Name, Assign Processors, Specify Memory, Add Network, Add Storage, Specify Options, **Select Workloads**, Performance Management, Summary.
- Section: **Select Workloads**
- Text: Select the workloads that this virtual server will participate in.
- Options:
  - Use Default workload
  - Select workloads
- Section: **New Workload**
- Buttons: < Back, Next >, Finish, Cancel, Help.

**Foreground Window: Performance Management**

- Navigation pane: Welcome, Enter Name, Assign Processors, Specify Memory, Add Network, Add Storage, Specify Options, Select Workloads, **Performance Management**, Summary.
- Section: **Performance Management**
- Text: Enable processor management for your virtual server to achieve the goals set in the active performance policy.
- Text: Ensemble processor management: Enabled
- Option:
  - Processor management
- Buttons: < Back, Next >, Finish, Cancel, Help.

# New Virtual Server Guest Definition



New Virtual Server - SCZP301:A12

- ✓ Welcome
- ✓ Enter Name
- ✓ Assign Processors
- ✓ Specify Memory
- ✓ Add Network
- ✓ Add Storage
- ✓ Specify Options
- ✓ Select Workloads
- ✓ Performance Management
- **Summary**

### Summary

Verify the information below before completing the wizard.

Name:	LBSZWAS2
Description:	
Initial virtual processors:	2
Assigned dedicated memory:	2048 MB
Network Devices:	600: OSD; LBS VLAN 250; IEDN250B
Storage Devices:	200: OSVOL, LX9883
Password:	ZWAS2
Privilege classes:	G
IPL device:	200
IPL load parameters:	
IPL parameters:	AUTOOCR
Workloads:	Default
Processor management:	Enabled

< Back   Next >   **Finish**   Cancel   Help

# New Virtual Server Guest Definition



Ensemble Management > ITSO Ensemble > Members

Members Virtual Servers Hypervisors Blades Topology

Filter Tasks Views

S ^	Name ^	Memb... ^	Status ^	Processo... ^	Mem... (MB) ^	Type ^	Auto St... ^	Shutd... Timeout (s) ^
<input checked="" type="checkbox"/>	A12 (LBSSSI:VMLINUX9)	SCZP301	Operating			z/VM	—	300
<input type="checkbox"/>	LBSOVS	SCZP301	Operating	1	1,500	z/VM		
<input type="checkbox"/>	LBSZWAS1	SCZP301	Not Active	1	1,500	z/VM		
<input type="checkbox"/>	LBSZWAS2	SCZP301	Not Active	1	2,048	z/VM		
<input type="checkbox"/>	LBSZWAS3	SCZP301	Not Active			z/VM		
<input type="checkbox"/>	LNXMNT	SCZP301	Not Active	1	1,024	z/VM		
<input type="checkbox"/>	A17 (VMLINUX4)	SCZP301	Operating			z/VM	—	300
<input type="checkbox"/>	B.1.01	SCZP301	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.02	SCZP301	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.03	SCZP301	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.04	SCZP301	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.05	SCZP301	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.10	SCZP301	Operating	1	32,768	PowerVM	✓	300
<input type="checkbox"/>	B.1.11	SCZP301	Operating	2	131,072	x Hyp	✓	300
<input type="checkbox"/>	B.1.12	SCZP301	Operating	2	131,072	x Hyp	✓	300

Max Page Size: 30 Total: 18 Filtered: 18 Selected: 1

Tasks: A12 (LBSSSI:VMLINUX9)

- Image Details
- Toggle Lock
- Daily

- Recovery
- Service

Operational Customization

Configuration

- Choose z/VM Virtual Servers to Manage
- Manage Storage Resources
- Manage Virtual Switches
- New Virtual Server



## Enable an Existing Virtual Server to be Managed

- You could either create new guest containers and copy or point them at existing disk storage
- Or you can make an existing guest be “Managed” by the Unified Resource Manager
- If you decide to start managing existing virtual servers, the pre-existing virtual network definitions will be reflected in the virtual server details but the “unmanaged” storage definitions will not.
- You can decide later to “unmanage” a virtual server, but it must have no managed resources assigned to it.

# Enable an Existing Virtual Server to be Managed



**Hardware Management Console**

Virtual Server Details

Ensemble Management > ITSO Ensemble > Members > SCZP301

Virtual Servers | Hypervisors | Blades | Topology

Filter: [ ] Tasks: [ ] Views: [ ]

Select	Name	Status	Processors	Memory (MB)	Type
<input type="checkbox"/>	A02	Operating			z/VM
<input checked="" type="checkbox"/>	A12	Operating			z/VM
<input type="checkbox"/>	LB3	Operating	1	1,500	z/VM
<input type="checkbox"/>	LB3	Operating	1	1,500	z/VM
<input type="checkbox"/>	LB3	Operated	1	1,500	z/VM
<input type="checkbox"/>	LND	Operated	1	1,024	z/VM
<input type="checkbox"/>	A17	Operating			z/VM
<input type="checkbox"/>	B.1.01	Operating			PowerVM
<input type="checkbox"/>	B.1.02	Operating			PowerVM
<input type="checkbox"/>	B.1.03	Operating	8	32,768	PowerVM

Max Page Size: 90 | Total: 13 | Filtered: 13 | Selected: 1

**Context Menu:**

- Image Details
- Toggle Lock
- Daily
- Recovery
- Service
- Operational Customization
- Configuration
  - Choose z/VM Virtual Servers to Manage
  - Manage Storage Resources
  - Manage Virtual Switches
  - New Virtual Server

# Enable an Existing Virtual Server to be Managed




Choose z/VM Virtual Machines to Manage - SCZP301:A12

Select or deselect the z/VM virtual machines that are to be managed by this console.

Select	Virtual Machine Name
<input type="checkbox"/>	GSKADMIN
<input type="checkbox"/>	IBMUSER
<input type="checkbox"/>	IMAP
<input type="checkbox"/>	IMAPAUTH
<input checked="" type="checkbox"/>	LBSOVS
<input checked="" type="checkbox"/>	LBSZWAS1
<input checked="" type="checkbox"/>	LBSZWAS2
<input type="checkbox"/>	LBSZWAS3
<input type="checkbox"/>	LDAPSRV
<input type="checkbox"/>	LGLOPR
<input checked="" type="checkbox"/>	LNXMNT
<input type="checkbox"/>	LPSERVE
<input type="checkbox"/>	MAINT
<input type="checkbox"/>	MIGMAINT
<input type="checkbox"/>	MONWRITE
<input type="checkbox"/>	MPROUTE
<input type="checkbox"/>	NAMESRV
<input type="checkbox"/>	NDBPMGR
<input type="checkbox"/>	NDBSRV01
<input type="checkbox"/>	NOBODY
<input type="checkbox"/>	OR1

Page 1 of 1 Total: 125 Filtered: 125 Displayed: 125 Selected: 4

OK Cancel Help



Choose z/VM Virtual Machines to Manage - SCZP301:A12

Select or deselect the z/VM virtual machines that are to be managed by this console.

Select	Virtual Machine Name
<input type="checkbox"/>	GSKADMIN
<input type="checkbox"/>	IBMUSER
<input type="checkbox"/>	IMAP
<input type="checkbox"/>	IMAPAUTH
<input checked="" type="checkbox"/>	LBSOVS
<input checked="" type="checkbox"/>	LBSZWAS1
<input checked="" type="checkbox"/>	LBSZWAS2
<input checked="" type="checkbox"/>	LBSZWAS3
<input type="checkbox"/>	LDAPSRV
<input type="checkbox"/>	LGLOPR
<input checked="" type="checkbox"/>	LNXMNT
<input type="checkbox"/>	LPSERVE
<input type="checkbox"/>	MAINT
<input type="checkbox"/>	MIGMAINT
<input type="checkbox"/>	MONWRITE
<input type="checkbox"/>	MPROUTE
<input type="checkbox"/>	NAMESRV
<input type="checkbox"/>	NDBPMGR
<input type="checkbox"/>	NDBSRV01
<input type="checkbox"/>	NOBODY

Page 1 of 1 Total: 125 Filtered: 125 Displayed: 125 Selected: 5

OK Cancel Help

# Enable an Existing Virtual Server to be Managed



**Hardware Management Console**

Virtual Server Details

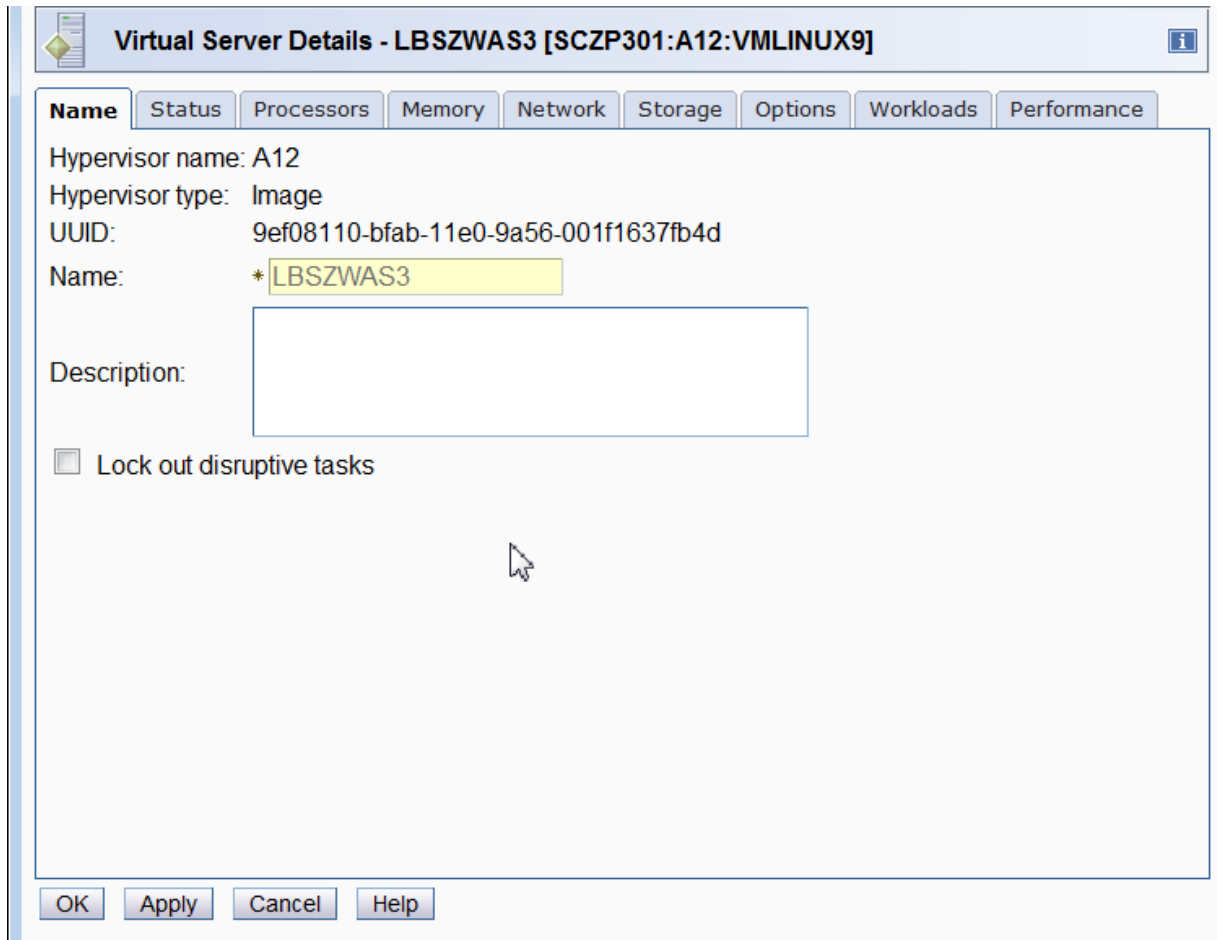
Ensemble Management > ITSO Ensemble > Members > SCZP301

Virtual Servers | Hypervisors | Blades | Topology

Select	Name	Status	Processors	Memory (MB)	Type	Auto Start
<input type="checkbox"/>	A02	Operating			z/VM	--
<input checked="" type="checkbox"/>	A12	Operating			z/VM	--
<input type="checkbox"/>	LBSOVS	Operating	1	1,500	z/VM	
<input type="checkbox"/>	LBSZWAS1	Not Activated	1	1,500	z/VM	
<input type="checkbox"/>	LBSZWAS2	Not Activated	1	1,500	z/VM	
<input type="checkbox"/>	LBSZWAS3	Not Activated	1	1,500	z/VM	
<input type="checkbox"/>	LNXMNT	Not Activated	1	1,024	z/VM	
<input type="checkbox"/>	A17	Operating			z/VM	--
<input type="checkbox"/>	B.1.01	Operating	8	32,768	PowerVM	--
<input type="checkbox"/>	B.1.02	Operating	8	32,768	PowerVM	--
<input type="checkbox"/>	B.1.03	Operating	8	32,768	PowerVM	--
<input type="checkbox"/>	B.1.04	Operating	8	32,768	PowerVM	--
<input type="checkbox"/>	B.1.05	Operating	8	32,768	PowerVM	--
<input type="checkbox"/>	B.1.10	Operating	8	32,768	PowerVM	--

Max Page Size: 90    Total: 14    Filtered: 14    Selected: 1

# Enable an Existing Virtual Server to be Managed



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

# Agenda

- 1 Why make your Linux guest part of the Ensemble
- 2 Requirements and Planning Considerations
- 3 Defining Resources for your Virtual Servers
- 4 Defining a New the Virtual Server or Managing an Existing Server
- 5 Linux Installation Considerations
- 6 Installing/Enabling the GPMP
- 7 Enabling ARM



# Installer Boot With OSX Interface

- The SLES 11 SP1 (and older) installer system will not find an OSX interface
- Install with OSDSIM interface and add or convert after kernel upgrade
- (An interface with OSDSIM would have type QDIO, but connect to an IEDN network)

```
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 >
```

# Installer Boot with OSD Simulated Interface


- If the Network Adapter is redefined as OSD instead of OSX, the OSA devices are discovered
- The SLES 11 SP2 installer system has all the required support built in and a simulated interface is not required

```
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 >
```

## 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
- LLADDR is added on SLES11 SP1, SLES11 SP2 does not add the LLADDR.

```
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 lbxzas1:/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 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 lbxzas1:/etc/sysconfig/network #
```

# Ensemble Virtual Networking – Virtual NICs

# Ensemble Virtual Networking – Virtual NICs



OSD NIC+ IEDN  
Switch = OSDSIM

Native OSX NIC

OSM NIC for  
INMN

OSX NIC without  
Vswitch  
(Trunked)

Virtual Server Details - LBSZWAS3 [SCZP301:A12:VMLINUX9]

General | Status | Processors | Memory | **Network** | Storage | Options | Workloads | Performance

MAC Prefix: 02:3c:90:00:00:24

Network Adapters:

Select	Virtual Device	Device Count	Type	Switch	Port Mode	Network
<input checked="" type="radio"/>	600	3	OSD	IEDN250B	Access	LBS VLAN 250
<input type="radio"/>	700	3	OSX	IEDN251	Access	LBS VLAN 251
<input type="radio"/>	800	3	RMC			Resource Monitoring and Control connection.
<input type="radio"/>	900	3	OSX			LBS VLAN 252
Total: 4						

Add Edit Remove

[Manage Virtual Networks](#)

Network Options:

Replicate VLAN IDs to Network Adapters

OK Apply Cancel Help



# Ensemble Virtual Networking – Virtual NICs



Loopback results omitted

```
lhxzwas3:~ # ip addr
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc pfifo_fast state UNKNOWN qlen 1000
   link/ether 02:3c:90:00:00:35 brd ff:ff:ff:ff:ff:ff
   inet 172.27.250.11/24 brd 172.27.250.255 scope global eth0
   inet6 fe80::3c:90ff:fe00:35/64 scope link
       valid_lft forever preferred_lft forever

3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc pfifo_fast state UNKNOWN qlen 1000
   link/ether 02:3c:90:00:00:36 brd ff:ff:ff:ff:ff:ff
   inet 172.27.251.11/24 brd 172.27.251.255 scope global eth1
   inet6 fe80::3c:90ff:fe00:36/64 scope link
       valid_lft forever preferred_lft forever

4: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc pfifo_fast state UNKNOWN qlen 1000
   link/ether 02:3c:90:00:00:37 brd ff:ff:ff:ff:ff:ff
   inet 169.254.49.126/16 brd 169.254.255.255 scope global eth3
   inet6 fe80::3c:90ff:fe00:37/64 scope link
       valid_lft forever preferred_lft forever

5: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc pfifo_fast state UNKNOWN qlen 1000
   link/ether 02:a0:30:00:00:14 brd ff:ff:ff:ff:ff:ff
   inet6 fe80::a0:30ff:fe00:14/64 scope link
       valid_lft forever preferred_lft forever

6: vlan252@eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc noqueue state UP
   link/ether 02:a0:30:00:00:14 brd ff:ff:ff:ff:ff:ff
   inet 172.27.252.11/24 brd 172.27.252.255 scope global vlan252
   inet6 fe80::a0:30ff:fe00:14/64 scope link
       valid_lft forever preferred_lft forever
```

OSDSIM NIC  
VLAN 250

OSX NIC  
VLAN 251

OSM NIC (IPV6)

OSX Trunk

OSX VLAN  
Subinterface  
VLAN 252



2012

## Ensemble Virtual Networking – Virtual NICs

```
lhxzwas3:/etc/sysconfig/network # ls
config  if-down.d  ifcfg-eth0  ifcfg-eth2  ifcfg-lo
        ifcfg.template  providers  routes.YaST2save
dhcp    if-up.d    ifcfg-eth1  ifcfg-eth3  ifcfg-vlan252
        ifroute-lo      routes      scripts
```

- Every permanent network interface needs an “ifcfg” script.
- Normally these would be created for you by Yast, if you use it.
- The VLAN sub interface can not be created by Yast and must be setup manually

# Virtual NICs – VLAN Subinterface



```
lbxzwas3:/etc/sysconfig/network # cat ifcfg-vlan252
BOOTPROTO='static'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='172.27.252.11/24'
MTU=''
NAME='IBM Ethernet controller (0.0.0900)'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
VLAN='yes'
ETHERDEVICE='eth2'
USERCONTROL='no'
```

Note: The actual VLAN ID # is only coded in the file/script name

# Virtual NICs – VLAN Subinterface



```
lbxzas3:/etc/sysconfig/network # cat ifcfg-eth2
#BOOTPROTO='static'
#BROADCAST=''
#ETHTOOL_OPTIONS=''
#IPADDR='/24'
#MTU=''
NAME='IBM Ethernet controller (0.0.0900)'
#NETWORK=''
#REMOTE_IPADDR=''
STARTMODE='auto'
#USERCONTROL='no'
```

This definition is referenced by the VLAN sub interface. Everything is commented out except for the STARTMODE and NAME.

# Virtual NICs - OSDSIM Interface



```
lbxzwas3:/etc/sysconfig/network # cat ifcfg-eth0
BOOTPROTO='static'
IPADDR='172.27.250.11/24'
BROADCAST=''
STARTMODE='auto'
NAME='OSA Express Network card (0.0.0
ETHTOOL_OPTIONS=''
INTERFACETYPE='qeth'
MTU=''
NETWORK=''
REMOTE_IPADDR=''
USERCONTROL='no'
```

This definition is for an interface using OSDSIM (an OSD NIC tied to an IEDN vswitch). No VLAN sub-interface is required

This looks like a “regular” OSD OSA interface.

## Virtual NICs – OSX Interface

```
lboxwas3:/etc/sysconfig/network # cat ifcfg-eth1
BOOTPROTO='static'
BROADCAST=''
ETHTOOL_OPTIONS=''
IPADDR='172.27.251.11/24'
MTU=''
NAME='IBM Ethernet controller (0.0.0700)'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

This interface is using a “Native” OSX device. However from this script you see know difference. Kernel support is required for a OSX interface connection.



# Virtual NICs – INMN IPV6



```
lboxwas3:/etc/sysconfig/network # cat ifcfg-eth3
```

```
BOOTPROTO='autoip'
```

```
BROADCAST=''
```

```
ETHTOOL_OPTIONS=''
```

```
INTERFACETYPE='qeth'
```

```
IPADDR=''
```

```
LLADDR=''
```

```
MTU=''
```

```
NAME='OSA Express Network card (0.0.0800)'
```

```
NETMASK=''
```

```
NETWORK=''
```

```
REMOTE_IPADDR=''
```

```
STARTMODE='auto'
```

```
USERCONTROL='no'
```

“autoip” is required  
for the linklocal  
address  
acquisition

IPV6 support is  
built in.

## Ensemble Virtual Networking – Virtual NICs -

```
lbxzw3:/etc/sysconfig/network # ip link show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state
   UNKNOWN
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc
   pfifo_fast state UNKNOWN qlen 1000
   link/ether 02:3c:90:00:00:35 brd ff:ff:ff:ff:ff:ff
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc
   pfifo_fast state UNKNOWN qlen 1000
   link/ether 02:3c:90:00:00:36 brd ff:ff:ff:ff:ff:ff
4: eth3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc
   pfifo_fast state UNKNOWN qlen 1000
   link/ether 02:3c:90:00:00:37 brd ff:ff:ff:ff:ff:ff
5: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492 qdisc
   pfifo_fast state UNKNOWN qlen 1000
   link/ether 02:a0:30:00:00:14 brd ff:ff:ff:ff:ff:ff
6: vlan252@eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1492
   qdisc noqueue state UP
   link/ether 02:a0:30:00:00:14 brd ff:ff:ff:ff:ff:ff
```

## UDEV and OSX Interfaces

- OSX Interfaces (like OSD) require UDEV entries
- Older levels of distros do not create the udev entry even if you have a kernel update that supports OSX
- Format of the UDEV entry is the same as OSD
- Copy and update an existing OSD UDEV entry if needed
- Don't forget OSX interfaces are layer2, ensure layer2=1 in the udev entry
- SLES 11 SP2 will automatically create the required UDEV entries

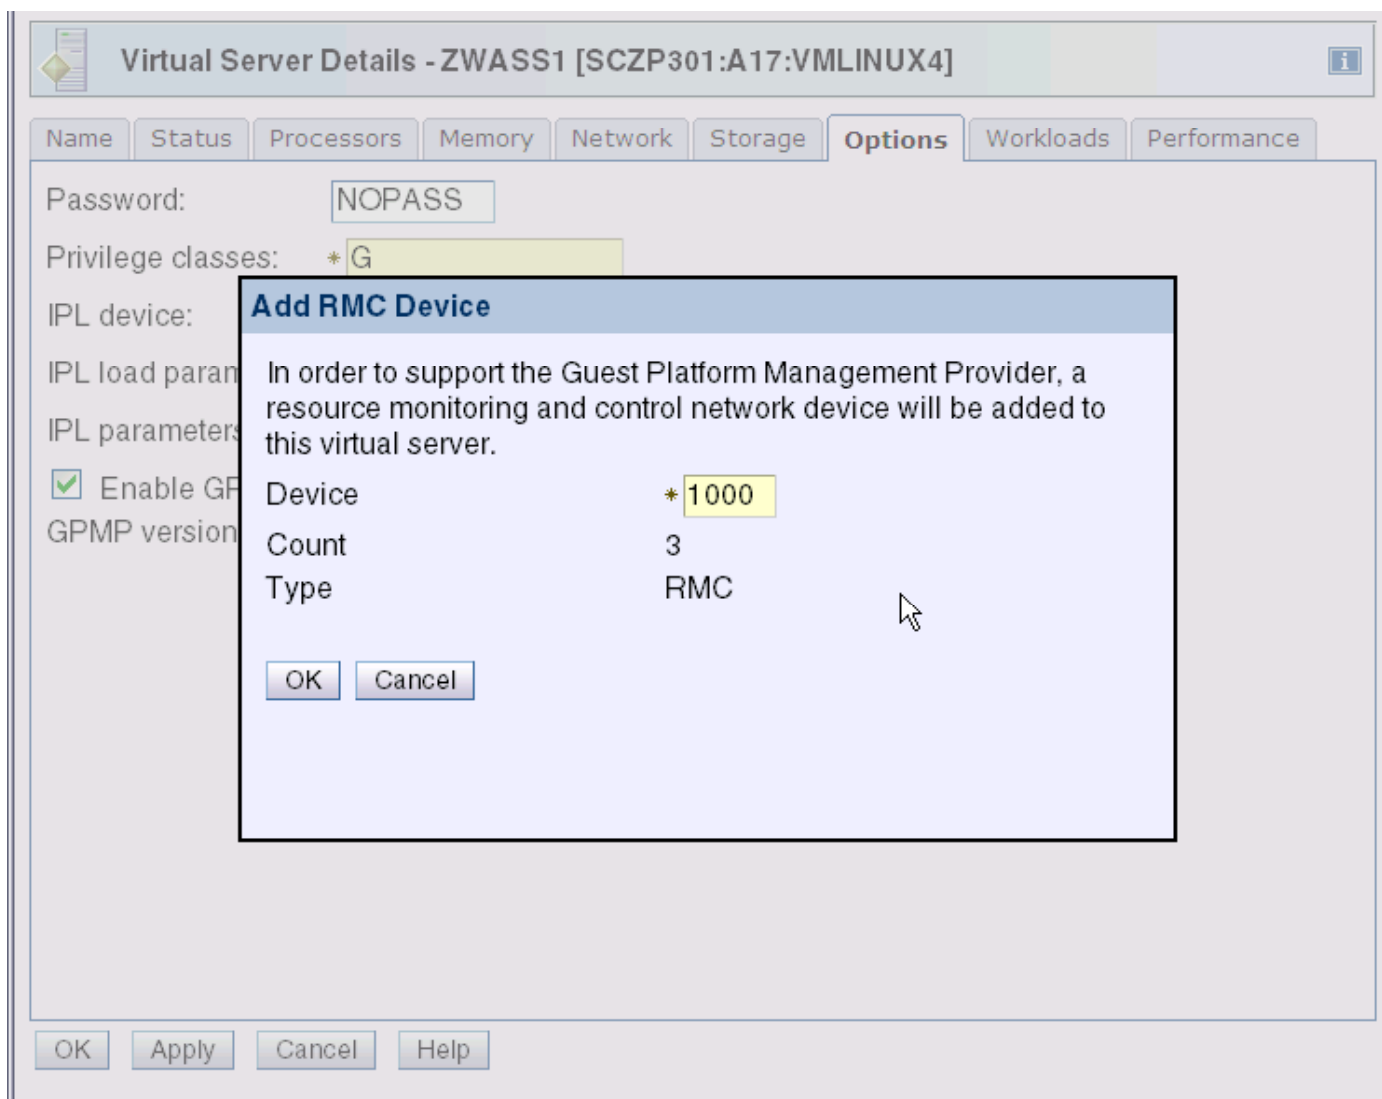
# Agenda

- 1 Why make your Linux guest part of the Ensemble
- 2 Requirements and Planning Considerations
- 3 Defining Resources for your Virtual Servers
- 4 Defining a New the Virtual Server or Managing an Existing Server
- 5 Linux Installation Considerations
- 6 Installing/Enabling the GPMP
- 7 Enabling ARM

# The Guest Platform Management Provider

- Provides detailed performance data from the guest operating system to the Unified Resource Manager
  - Is provided via the firmware/microcode update stream
- Ensemble performance data flows across the INMN from your guest, but never your application data
- Optional, but is required if you want to feed Application Response Measurement data (ARM) from middleware to the zManager and is also required for “workload management” by the Unified Resource Manager
- Lets take a look at the steps to enable it

# Enabling INMN on Linux for System z



Virtual Server Details - ZWASS1 [SCZP301:A17:VMLINUX4]

Name Status Processors Memory Network Storage **Options** Workloads Performance

Password: NOPASS

Privilege classes: \*G

IPL device:

IPL load param

IPL parameters

Enable GP

GPMP version

**Add RMC Device**

In order to support the Guest Platform Management Provider, a resource monitoring and control network device will be added to this virtual server.

Device	*1000
Count	3
Type	RMC

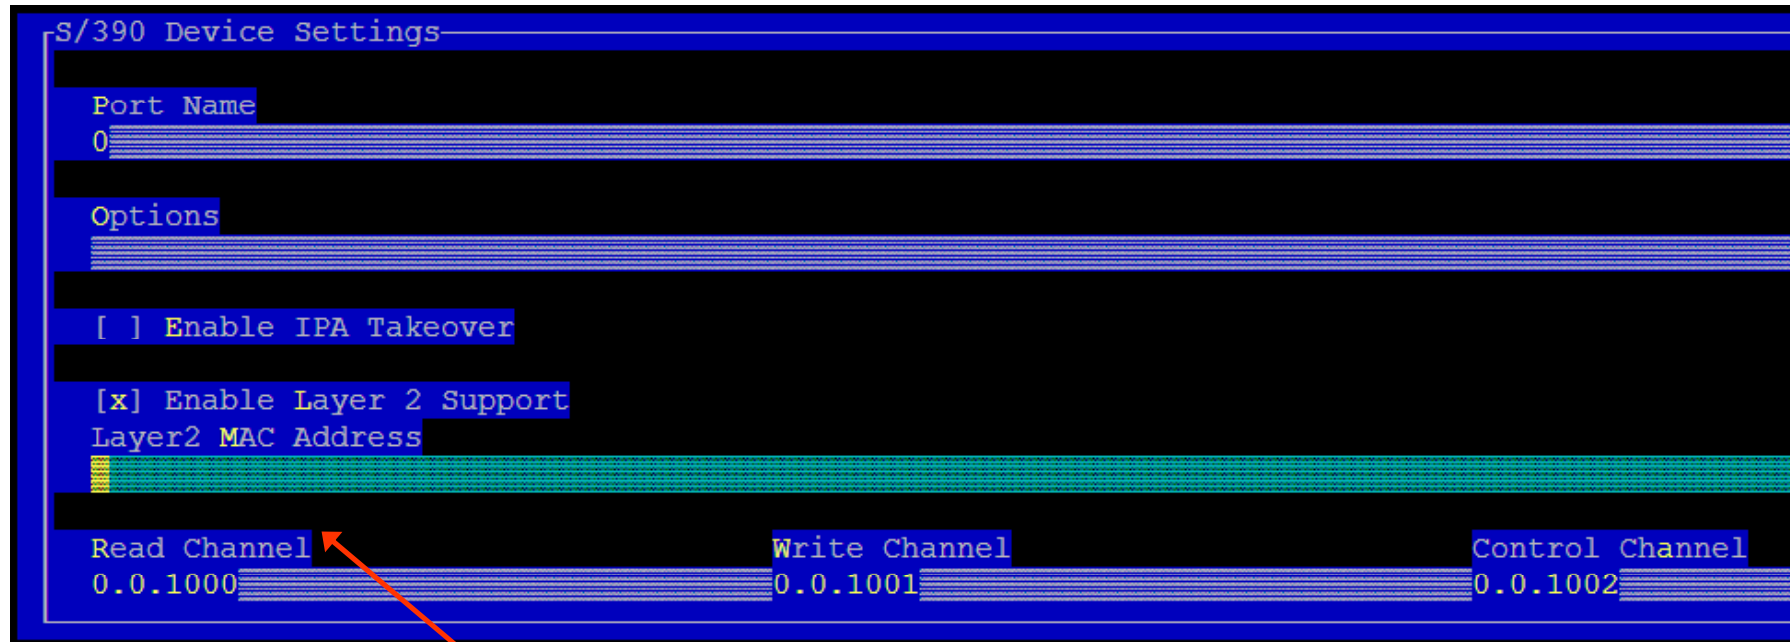
OK Cancel

OK Apply Cancel Help



## Enabling INMN on Linux – Yast Example

```
S/390 Device Settings
Port Name
0
Options
[ ] Enable IPA Takeover
[x] Enable Layer 2 Support
Layer2 MAC Address
Read Channel      Write Channel      Control Channel
0.0.1000          0.0.1001          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

## Enabling INMN on Linux for System z

```
YaST2 - lan @ wass1
Network Card Setup
General—Address—Hardware—
Device Type _____ Configuration Name _____
QETH _____ â eth3 _____
( ) No IP Address (for Bonding Devices)
(x) Dynamic Address _____
Zeroconf _____ â DHCP version 6 only _____
( ) Statically assigned IP Address _____
IP Address _____ Subnet Mask _____ Hostname _____
Additional Addresses _____
Alias Name | IP Address | Netmask _____
```

- Utilize “Zeroconf” to dynamically assign the IP address! DON’T use DHCP!

## Validating INMN on Linux for System z

```
eth3      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)
```

- Expect to see a 169.254.x.x address assigned to you INMN interface
- Remember it is copper Ethernet so no jumbo frames like the IEDN

## Validating INMN on Linux for System z

SCZP301:A17 Details - SCZP301:A17 i

Instance Information    Acceptable Status    **Hypervisor Information**

Description:

Virtual Server shutdown timeout (seconds):

Management Guest IPv6 Address:

Apply    Change Options...    Cancel    Help

## 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
```

- Pinging 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

```
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

```
chown ibmgpmp:ibmgpmp /opt/ibm/gpmp/CollectFFDC.sh /opt/ibm/gpmp/armsad /opt/ibm/gpmp/gpmpmain /opt/ibm/gpmp/gpmpsad /opt/ibm/gpmp/gpmpshm /opt/ibm/gpmp/post-install-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 /opt/ibm/gpmp/gpmpmain /opt/ibm/gpmp/gpmpsad /opt/ibm/gpmp/gpmpshm /opt/ibm/gpmp/java /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

```
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

```
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 -e
wass1:~ #
```

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

# Agenda

- 1 Why make your Linux guest part of the Ensemble
- 2 Requirements and Planning Considerations
- 3 Defining Resources for your Virtual Servers
- 4 Defining a New the Virtual Server or Managing an Existing Server
- 5 Linux Installation Considerations
- 6 Installing/Enabling the GPMP
- 7 Enabling ARM

## What is ARM and why do I want to use it?

- Application Response Measurement (ARM) – An open standard for monitoring and measuring performance
- Allows for tracking the HOPs of an application request across different middleware components.
- The API can be leveraged from C and Java programs
- By enabling the ARM API, you can generate a HOPs report in the Unified Resource Manager to better understand application performance and identify bottlenecks.

# Enabling ARM – GPMP Post Install Configuration

```
wasgl:/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

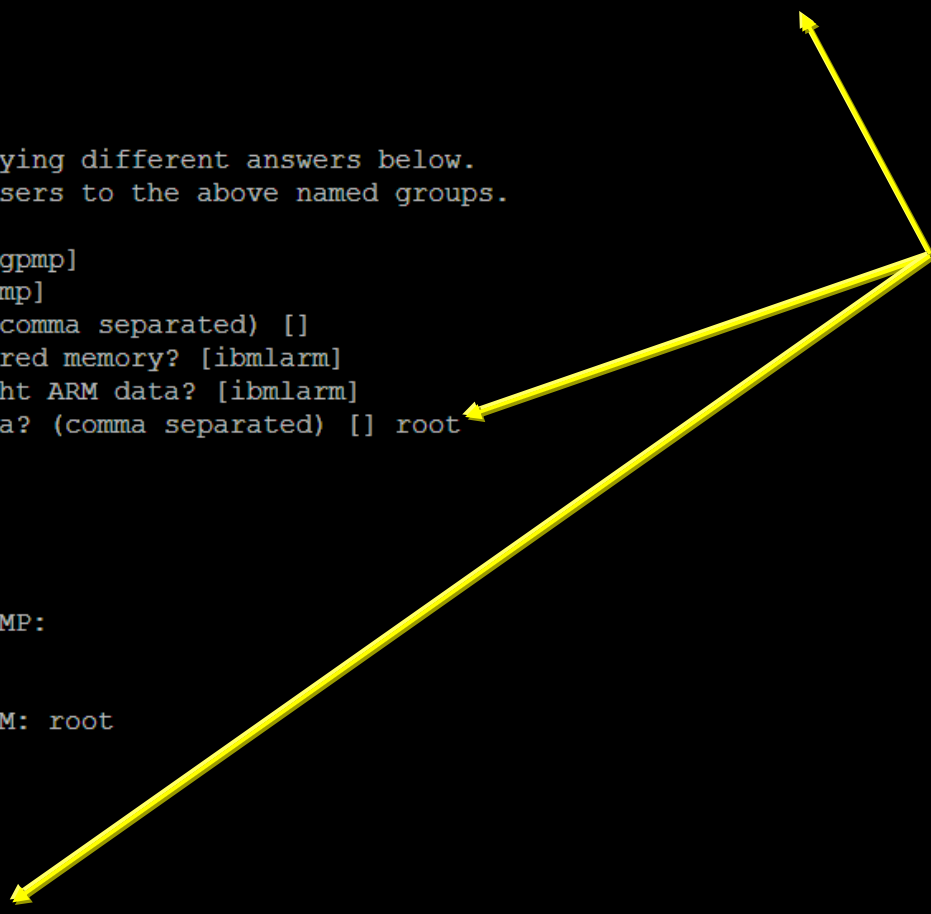
About to change GPMP configuration:

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, WAS on Linux


**Application servers**

[Application servers](#) > [server1](#) > [Process definition](#) > [Java Virtual Machine](#) > **Custom properties**

Use this page to specify an arbitrary name and value pair. The value that is specified for the name and value pair is a string that can set internal system configuration properties.

⊕ Preferences

New Delete



Select	Name	Value	Description
You can administer the following resources:			
<input type="checkbox"/>	<a href="#">com.ibm.security.jgss.debug</a>	off	
<input type="checkbox"/>	<a href="#">com.ibm.security.krb5.Krb5Debug</a>	off	
<input type="checkbox"/>	<a href="#">com.ibm.websphere.pmi.reqmetrics.PassCorrelatorToDB</a>	true	
<input type="checkbox"/>	<a href="#">ws.ext.dirs</a>	/opt/ibm/gpmp/java	
Total 4			



Add two custom properties to WAS JVM

# Enabling ARM, WAS on Linux

## Request Metrics

Request metrics tracks each individual transaction in WebSphere Application Server, recording the response time of the major components such as time in the Web server or in the Enterprise JavaBeans (EJB) container. Use this page to enable request metrics, select the components that are instrumented by request metrics, set trace levels, enable standard logs, enable Application Response Measurement (ARM), specify the type of ARM agent, and specify the ARM transaction factory implementation class name.

Configuration

**General Properties** **Additional Properties**

Prepare Servers for Request metrics collection

**Components to be instrumented**

None  
 All  
 Custom

AsyncBeans  
 EJB  
 JCA  
 JDBC  
 JMS  
 JNDI  
 Portlet  
 SIB  
 Servlet  
 Servlet Filter  
 WebServices

\* Trace level  
 Hops

**Request Metrics Destination**

Standard Logs  
 Application Response Measurement(ARM) agent

Agent Type  
 ARM40

ARM transaction factory implementation class name  
 com.ibm.wlm.arm40SDK.tri

# Enabling ARM on Linux

```
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 to add /usr/lib64 to WAS\_LIBPATH

# Enabling ARM, WAS on Linux



- Start the WebSphere server

```
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/lssarm -a
FEW6046I APPL: WebSphere:APPLICATION_SERVER
wasg1:/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin # █
```

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

## References

- Via IBM Resource Link:
  - IBM zEnterprise Ensemble Performance Management Guide *GC27-2607*
  - IBM zEnterprise Ensemble Planning and Configuration Guide *GC27-2608*
  - HMC Web Services API *SC27-2616*
- Red Book
  - IBM zEnterprise Unified Resource Manager *SG24-7921*
- [www.vm.ibm.com](http://www.vm.ibm.com)
  - *Required service* <http://www.vm.ibm.com/service/vmrequrm.html>
  - *SC24-6178-03 z/VM CP Planning and Administration Guide*
  - *SC24-6175-03 z/VM CP Commands and Utilities Reference*
  - *SC24-6234-03 z/VM Systems Management API*
  - *SC24-6174-03 z/VM V6R2 Connectivity*



# Session Evaluations

- Enabling Linux on System z for Ensemble management
- Session 11937
- [www.SHARE.org/AnaheimEval](http://www.SHARE.org/AnaheimEval)







**Richard G. Young**  
*Senior Certified I.T. Specialist*  
*IBM STG Lab Services*  
*Virtualization & Linux on*  
*zEnterprise Team Lead*

*777 East Wisconsin Ave*  
*Milwaukee, WI 53202*

*Tel 262 893 8662*  
*Email: ryoung1@us.ibm.com*