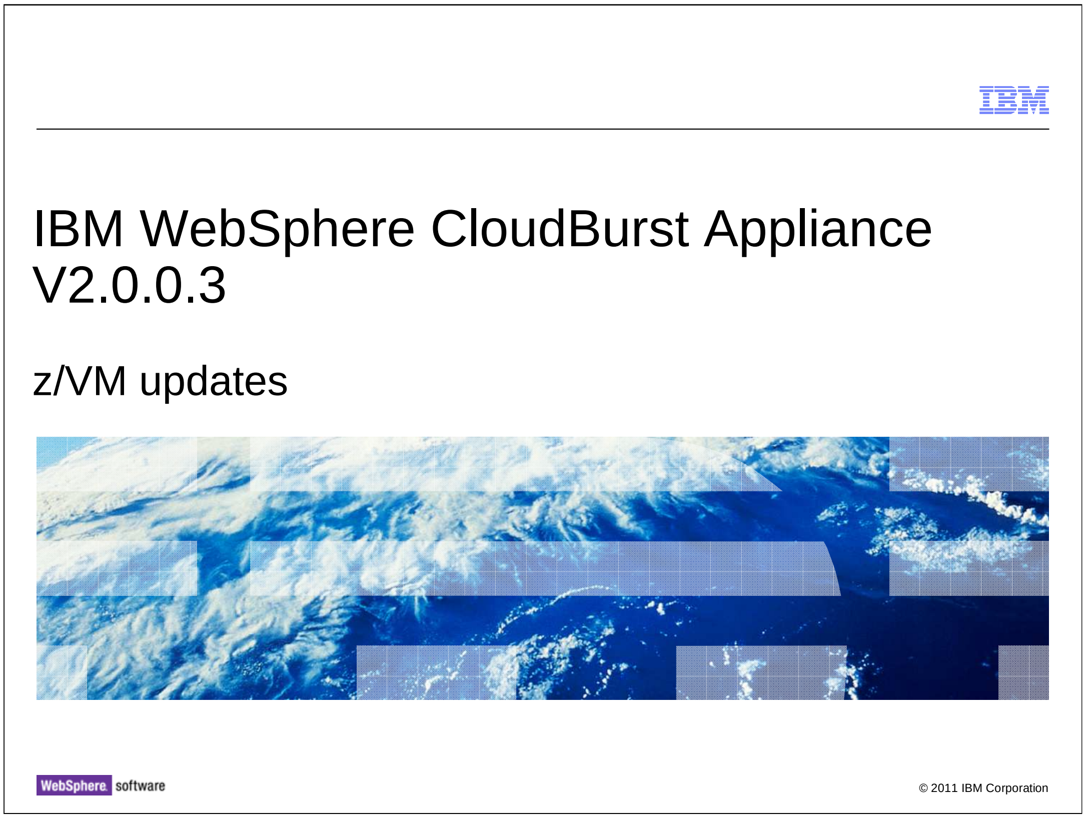

IBM WebSphere CloudBurst Appliance V2.0.0.3

z/VM updates

This presentation covers the new function available in the V2.0.0.3 fix pack for IBM WebSphere CloudBurst™ having to do with the z/VM hypervisor.

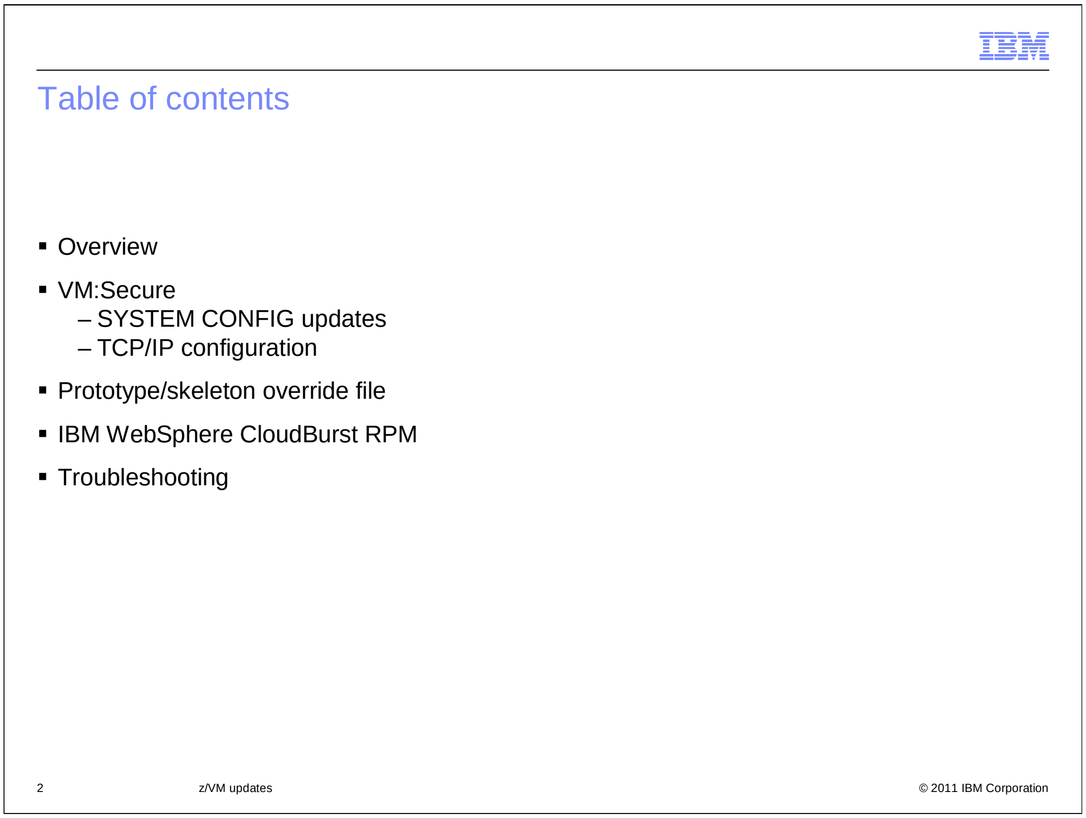

Table of contents

- Overview
- VM:Secure
  - SYSTEM CONFIG updates
  - TCP/IP configuration
- Prototype/skeleton override file
- IBM WebSphere CloudBurst RPM
- Troubleshooting

You will start out looking at an overview of the requirements necessary to configure a z/VM hypervisor with the IBM WebSphere CloudBurst appliance. The basic requirements for configuration have not really changed but you will review them in light of the VM:Secure support, which you will see in detail here. With V2.0.0.3, there is a new prototype, or skeleton, override file available for both VM:Secure and DirMaint configurations that you will briefly look at. You will review the RPM needed for the z/VM hypervisor and its interaction with the VM:Secure configuration and finally you will see some troubleshooting information that you can use if you run into problems.

Section

Overview

3 z/VM updates © 2011 IBM Corporation

This section will present the z/VM prerequisites to configure a z/VM hypervisor for IBM WebSphere CloudBurst.

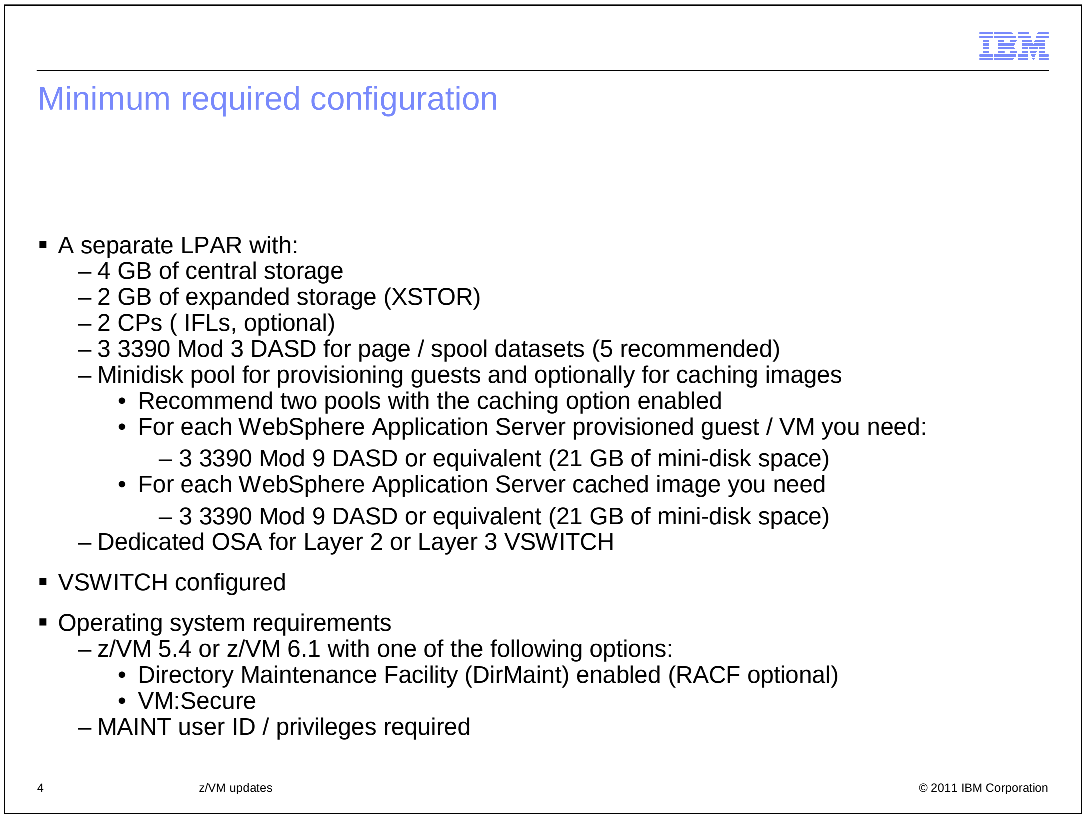

Minimum required configuration

- A separate LPAR with:
  - 4 GB of central storage
  - 2 GB of expanded storage (XSTOR)
  - 2 CPs (IFLs, optional)
  - 3 3390 Mod 3 DASD for page / spool datasets (5 recommended)
  - Minidisk pool for provisioning guests and optionally for caching images
    - Recommend two pools with the caching option enabled
    - For each WebSphere Application Server provisioned guest / VM you need:
      - 3 3390 Mod 9 DASD or equivalent (21 GB of mini-disk space)
    - For each WebSphere Application Server cached image you need
      - 3 3390 Mod 9 DASD or equivalent (21 GB of mini-disk space)
  - Dedicated OSA for Layer 2 or Layer 3 VSWITCH
- VSWITCH configured
- Operating system requirements
  - z/VM 5.4 or z/VM 6.1 with one of the following options:
    - Directory Maintenance Facility (DirMaint) enabled (RACF optional)
    - VM:Secure
  - MAINT user ID / privileges required

Shown here is the minimum required hardware configuration for the hypervisor. Note that for performance reasons, it is recommended to have five page datasets defined to the system. Also note that a minimum of one minidisk pool is required. One minidisk pool is used for provisioned guests and for a WebSphere Application Server provisioned guest, 21GB of space is needed. This space will vary depending on the image being deployed and if you have ‘Use shared minidisks’ enabled. If ‘sharing’ is enabled, each provisioned guest will link to some disks rather than having their own copies of each disk. A second minidisk pool can be used for caching images. The space needed for the cached image minidisk pool will depend on how many images you will be caching. Again, in the case of the WebSphere Application Server images, 21GB of space is required. You can use the same minidisk pool for both the provisioned guests and the cached images but two are recommended.

Starting with V2.0.0.3, you can choose either VM:Secure or DirMaint as your directory manager. The MAINT user ID privilege is needed for initial setup of the hypervisor, however once you start provisioning with IBM WebSphere CloudBurst, a Class A user ID that you create is used instead.

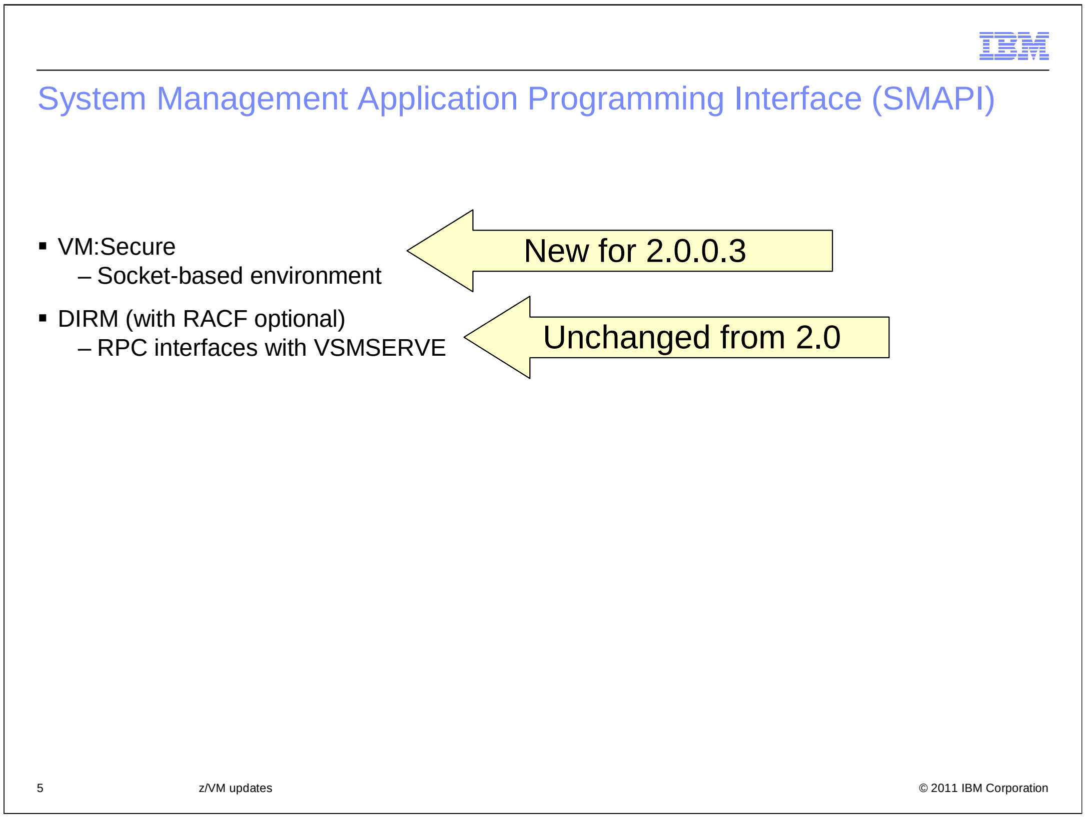

System Management Application Programming Interface (SMAPI)

- VM:Secure
  - – Socket-based environment
- DIRM (with RACF optional)
  - – RPC interfaces with VSMSERVE

New for 2.0.0.3

Unchanged from 2.0

5 z/VM updates © 2011 IBM Corporation

The System Management Application Programming Interface, or SMAPI, is used to perform system management functions for virtual images, or guests, in the IBM z/VM environment. When using VM:Secure as your directory manager, you must use the newer socket-based environment to issue the system management APIs. If you are using DirMaint, the older Remote Procedure Call, or RPC, interface is still used in V2.0.0.3 as it was in earlier releases.

This presentation will focus mainly on the configuration needed for the socket-based environment. The RPC interface configuration is highlighted in the z/VM configuration presentation found in the V2.0 Configuration section.

Section

VM:Secure

This section will present the IBM WebSphere CloudBurst configuration in the z/VM environment, using VM:Secure as your directory manager.

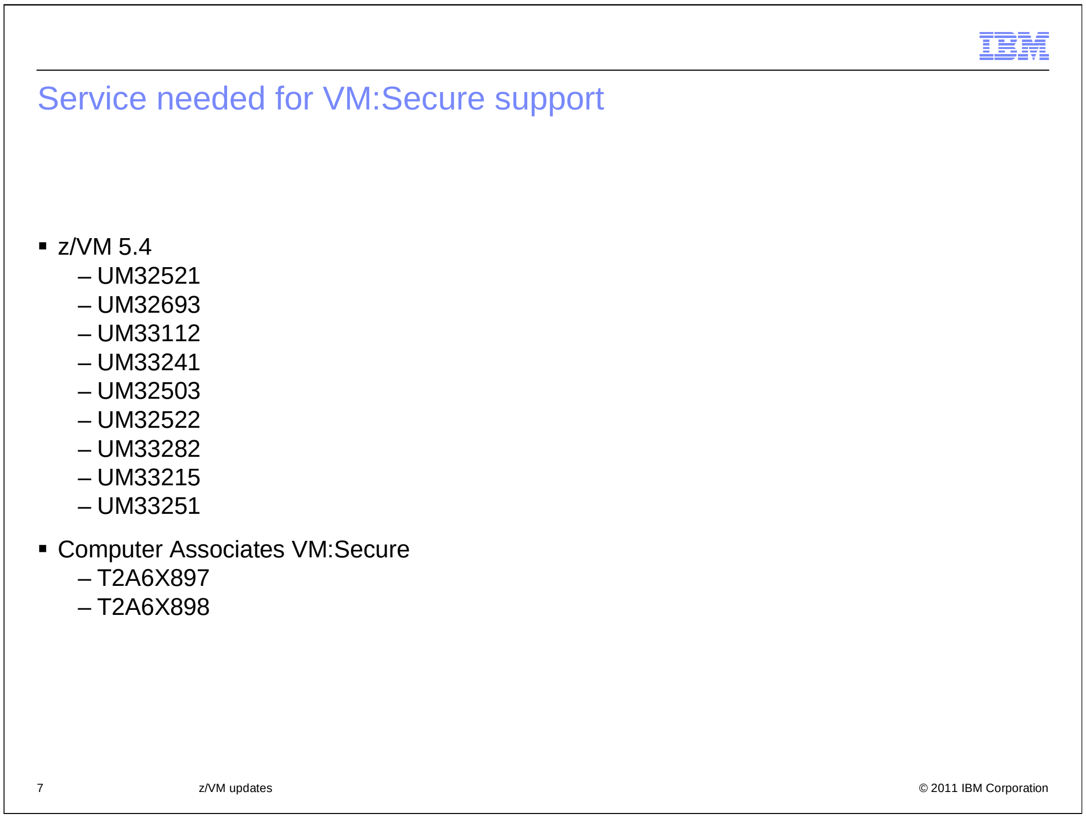

Service needed for VM:Secure support

- z/VM 5.4
  - UM32521
  - UM32693
  - UM33112
  - UM33241
  - UM32503
  - UM32522
  - UM33282
  - UM33215
  - UM33251
- Computer Associates VM:Secure
  - T2A6X897
  - T2A6X898

The first thing you need to do is put some important service on your z/VM system. This slide lists service that you must have on your z/VM 5.4 system in order for the function to work correctly. Note that there are a couple of VM:Secure fixes listed that you must also get from Computer Associates.

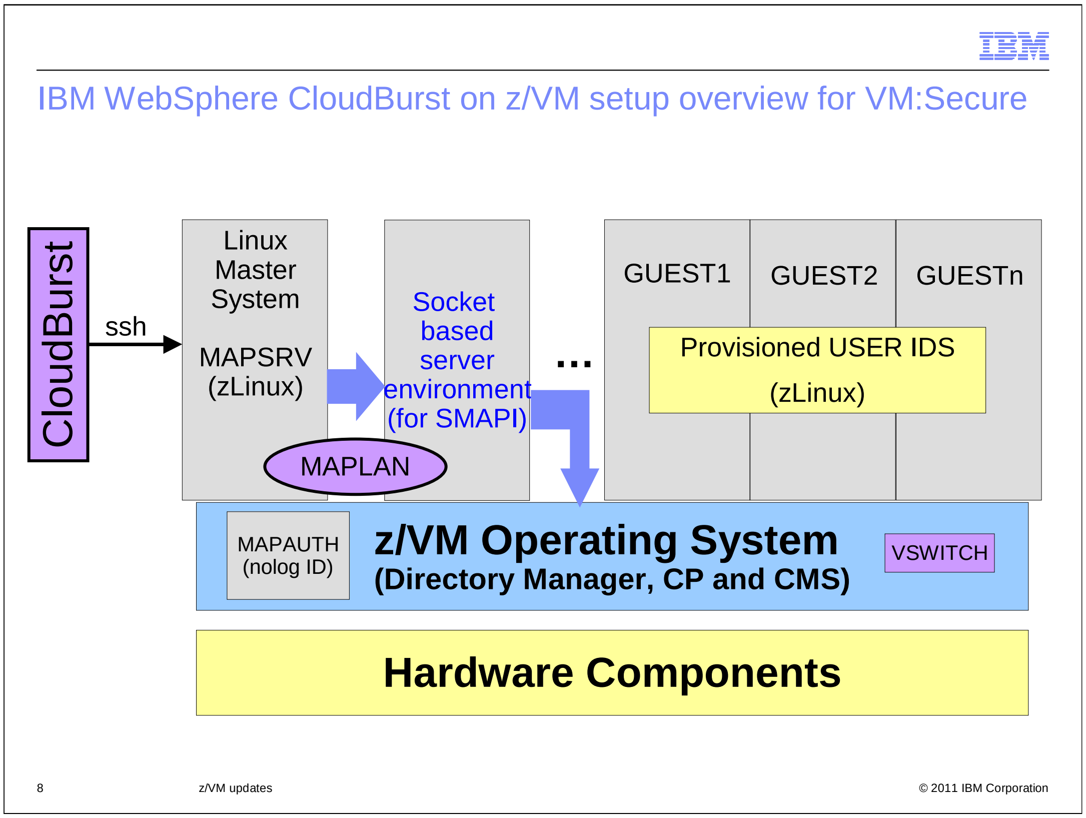

IBM WebSphere CloudBurst on z/VM setup overview for VM:Secure

CloudBurst

ssh

Linux Master System

MAPSRV (zLinux)

MAPLAN

Socket based server environment (for SMAPI)

…

GUEST1

GUEST2

GUESTn

Provisioned USER IDS

(zLinux)

MAPAUTH (nolog ID)

z/VM Operating System

(Directory Manager, CP and CMS)

VSWITCH

Hardware Components

8 z/VM updates © 2011 IBM Corporation

This slide shows an overview of some of the user IDs used in IBM WebSphere CloudBurst administration and how they interact. The big difference here for VM:Secure is the box labeled ‘Socket-based server environment (for SMAPI)’. For DirMaint in V2.0.0.3, this box is labeled as using TCP/IP by way of the VSMSERVE guest. Since this presentation is focusing on configuration for VM:Secure, you will see more information on the socket-based environment box on the next slides.

The IBM WebSphere CloudBurst appliance will communicate directly to the Linux Master System. For the purposes of this presentation, the Linux Master System is referred to as the MAPSRV user ID. An RPM is installed on the MAPSRV user ID that will take the commands coming from the appliance and interpret them into SMAPI commands. The SMAPI commands do what is necessary to provision the guests on z/VM as seen on the right side of the slide. This is just a high-level view but it should give you an idea of how the various pieces fit together. The MAPSRV user ID uses the MAPLAN TCP/IP interface shown on the slide to pass commands to the socket-based server environment. The MAPAUTH user ID shown on the slide is the one set up to issue restricted SMAPI commands. This user ID is a nolog ID and is used instead of MAINT. Finally, the VSWITCH shown on the slide is used for external communications.

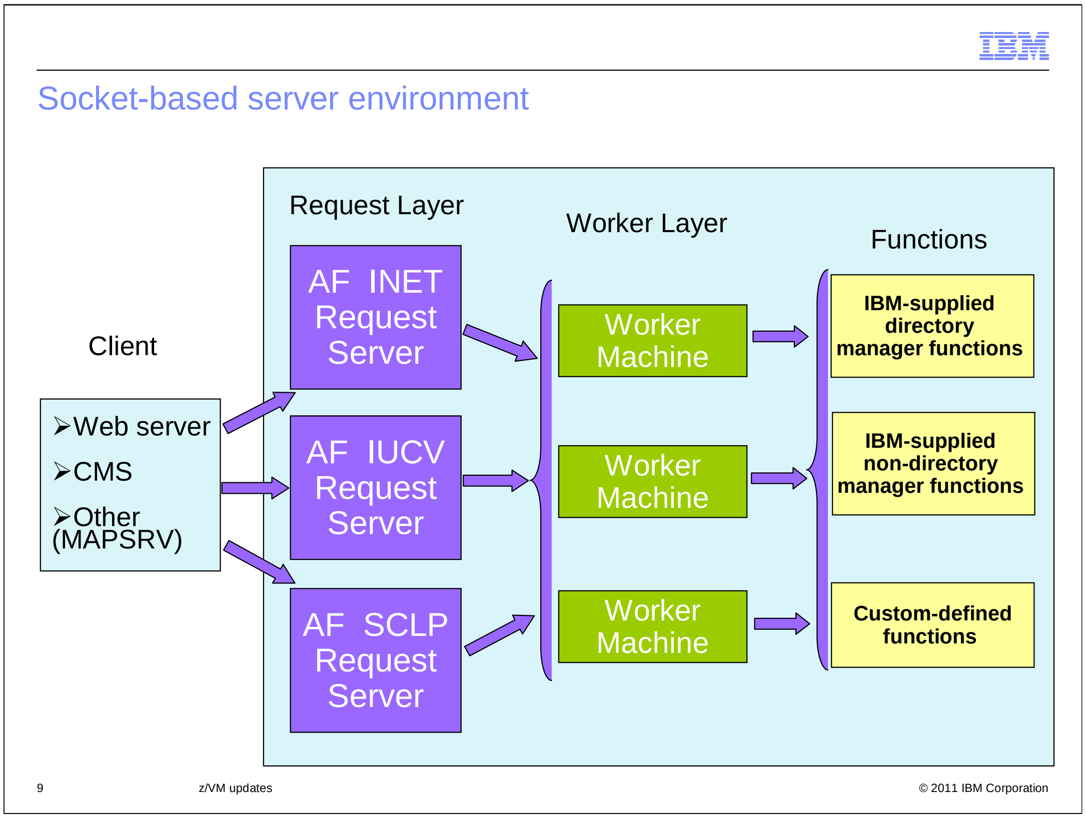

Socket-based server environment

Client

- Web server
- CMS
- Other (MAPSRV)

Request Layer

AF INET Request Server

AF IUCV Request Server

AF SCLP Request Server

Worker Layer

Worker Machine

Worker Machine

Worker Machine

Functions

IBM-supplied directory manager functions

IBM-supplied non-directory manager functions

Custom-defined functions

9 z/VM updates

The next few slides will show you the make-up of the socket-based server environment. The socket-based server environment consists of one or more request servers and two or more worker servers. The request server listens for socket connections initiated by a client program. The client program can be a web server, CMS or in the IBM WebSphere CloudBurst case, the Linux Master System, or MAPSRV user ID. The server accepts the connection, receives the data, and then calls the appropriate worker server to process the request, while the client program waits for the response.

As seem on the slide, there are three types of request servers. The **AF_INET** request server uses internet protocols for communication with the client and is the one used for IBM WebSphere CloudBurst. The **AF_IUCV** request server uses the Inter-User Communications Vehicle (IUCV) for point-to-point connections when z/VM components need to communicate with each another or with CP. Finally the **AF_SCLP** request server is used to receive and transmit Hardware Management Console events.

These servers are defined as separate virtual machines in the default z/VM installation.

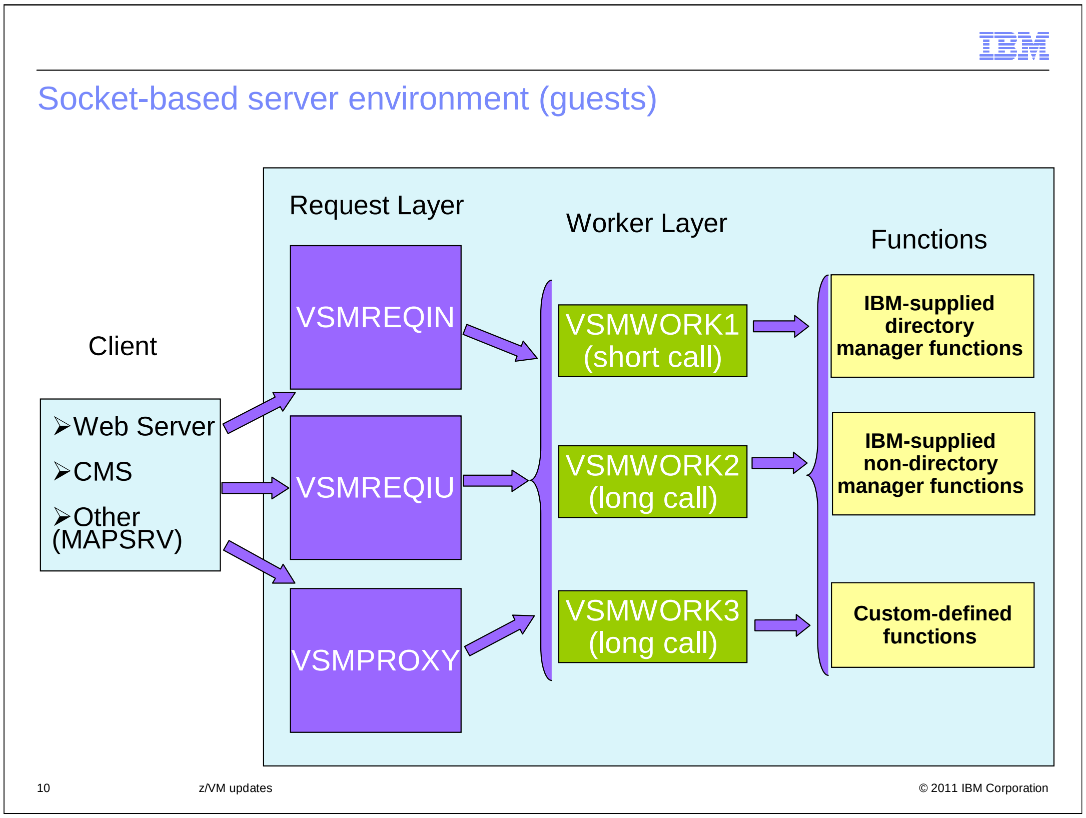

IBM

Socket-based server environment (guests)

Client

➢Web Server
➢CMS
➢Other (MAPSRV)

Request Layer

VSMREQIN

VSMREQIU

VSMPROXY

Worker Layer

VSMWORK1 (short call)

VSMWORK2 (long call)

VSMWORK3 (long call)

Functions

IBM-supplied directory manager functions

IBM-supplied non-directory manager functions

Custom-defined functions

10 z/VM updates © 2011 IBM Corporation

This slide shows the same picture as the previous slide but has the default z/VM guest names instead. By default, VSMREQIN is the AF_INET request server, VSMREQIU is the AF_IUCV request server and VSMPROXY is the AF_SCLP request server.

The worker servers, VSMWORKx by default, process the API function requests. The z/VM default installation defines three worker servers: VSMWORK1, VSMWORK2, and VSMWORK3. There are two types of API calls, a “short call" and “long call." The first worker server, VSMWORK1, is always the “short call" worker. All other worker servers are designated as “long call" workers. These workers handle API requests that require more time than the “short call" requests. When more than one “long call" server is active, a worker server that is not busy will receive the request. If all worker servers are busy, the request will be queued to one of the worker servers. Note that there must always be at least one short call worker server and at least one long call worker server. A total of three servers (one short call and two long call) is the recommended minimum however.

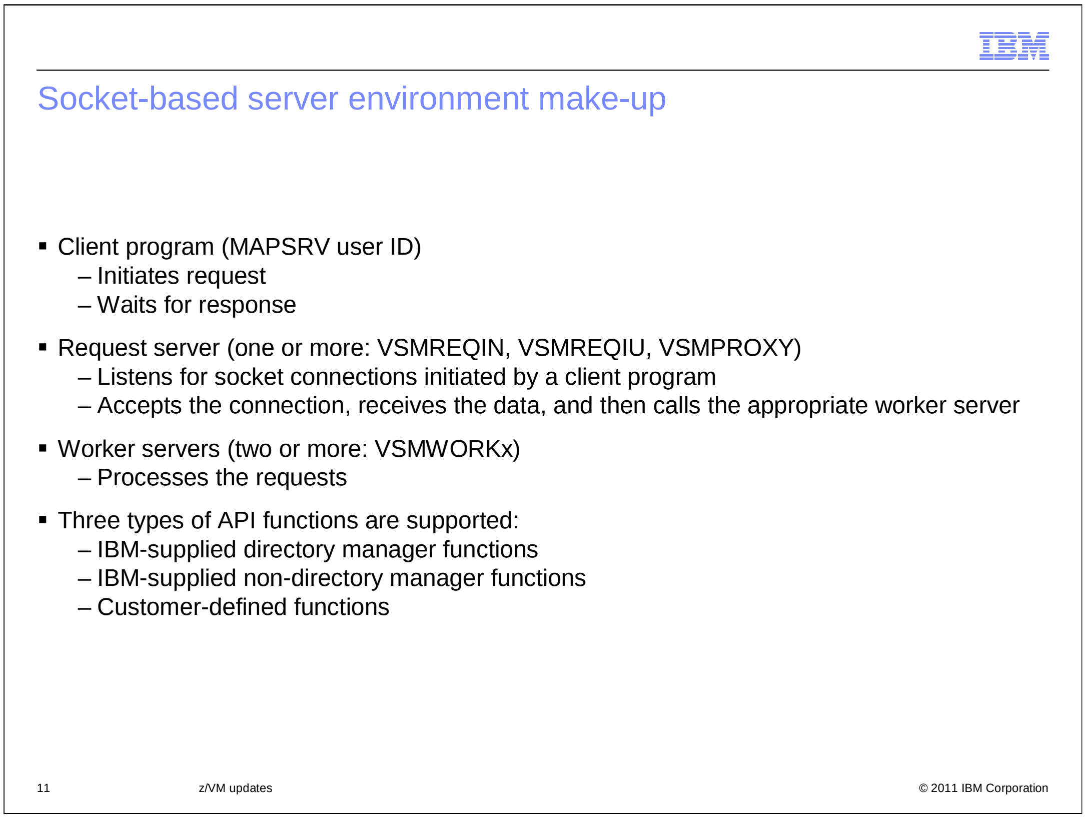

Socket-based server environment make-up

- Client program (MAPSRV user ID)
  - – Initiates request
  - – Waits for response
- Request server (one or more: VSMREQIN, VSMREQIU, VSMPROXY)
  - – Listens for socket connections initiated by a client program
  - – Accepts the connection, receives the data, and then calls the appropriate worker server
- Worker servers (two or more: VSMWORKx)
  - – Processes the requests
- Three types of API functions are supported:
  - – IBM-supplied directory manager functions
  - – IBM-supplied non-directory manager functions
  - – Customer-defined functions

Summarizing what you just saw on the diagram then, the socket-based server environment consists of a client program that initiates a request and waits for the response. In the case of IBM WebSphere CloudBurst this is the Linux Master System, or MAPSRV user ID. The request servers listen for requests from the client programs and will receive the data and call the appropriate worker server who then process the request. The request server that the IBM WebSphere CloudBurst appliance uses is the VSMREQIN guest, which process requests coming in by way of the internet protocol. There are three types of requests supported by the SMAPI that can be processed by the socket-based server environment. They are: IBM-supplied directory manager functions, IBM-supplied non-directory manager functions and Customer-defined functions

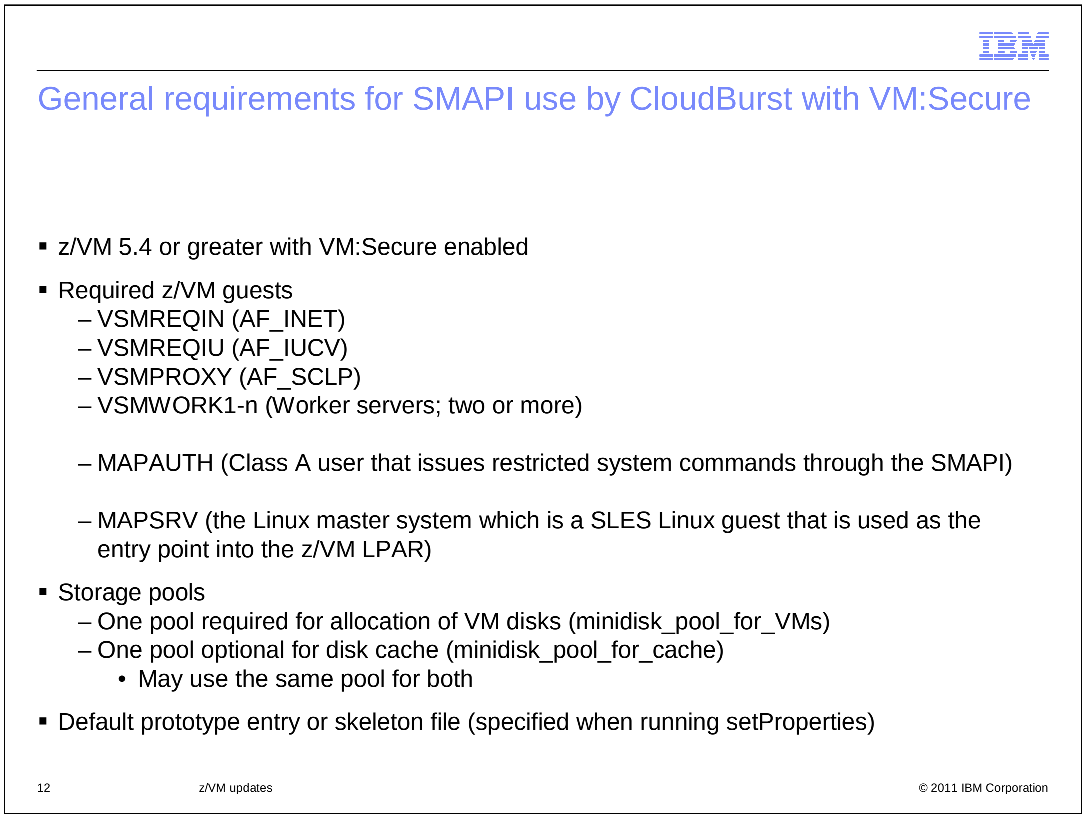

General requirements for SMAPI use by CloudBurst with VM:Secure

- z/VM 5.4 or greater with VM:Secure enabled
- Required z/VM guests
  - VSMREQIN (AF_INET)
  - VSMREQIU (AF_IUCV)
  - VSMPROXY (AF_SCLP)
  - VSMWORK1-n (Worker servers; two or more)
  - MAPAUTH (Class A user that issues restricted system commands through the SMAPI)
  - MAPSRV (the Linux master system which is a SLES Linux guest that is used as the entry point into the z/VM LPAR)
- Storage pools
  - One pool required for allocation of VM disks (minidisk_pool_for_VMs)
  - One pool optional for disk cache (minidisk_pool_for_cache)
    - May use the same pool for both
- Default prototype entry or skeleton file (specified when running setProperties)

In general, looking at the needed configuration then to use VM:Secure with IBM WebSphere CloudBurst, there are some required z/VM guests needed as listed on the slide. The list of guests includes the request and worker servers for the socket-based environment, the MAPAUTH user ID that is used to issue restricted system commands through the SMAPI and the MAPSRV user ID, which is the entry point from the appliance into the z/VM LPAR. These names may differ on your system, but these are the default names that are used in this presentation.

Another requirement, which is not specific to VM:Secure, is the definition of one or more storage pools. One pool is required for the allocation of VM disks for the guests that are deployed by IBM WebSphere CloudBurst. A second pool that can be used for caching images is optional. The variables set to define these pools are shown in parentheses. You will see how these are set later when the RPM is described. Note that you can choose to use the same pool for both the VM disks and disk cache but two pools are recommended.

Finally, you need to have a default prototype entry, or skeleton file, defined. This is used to create the directory entries for the guests that are deployed. The following slides will talk about each of these requirements in more detail.

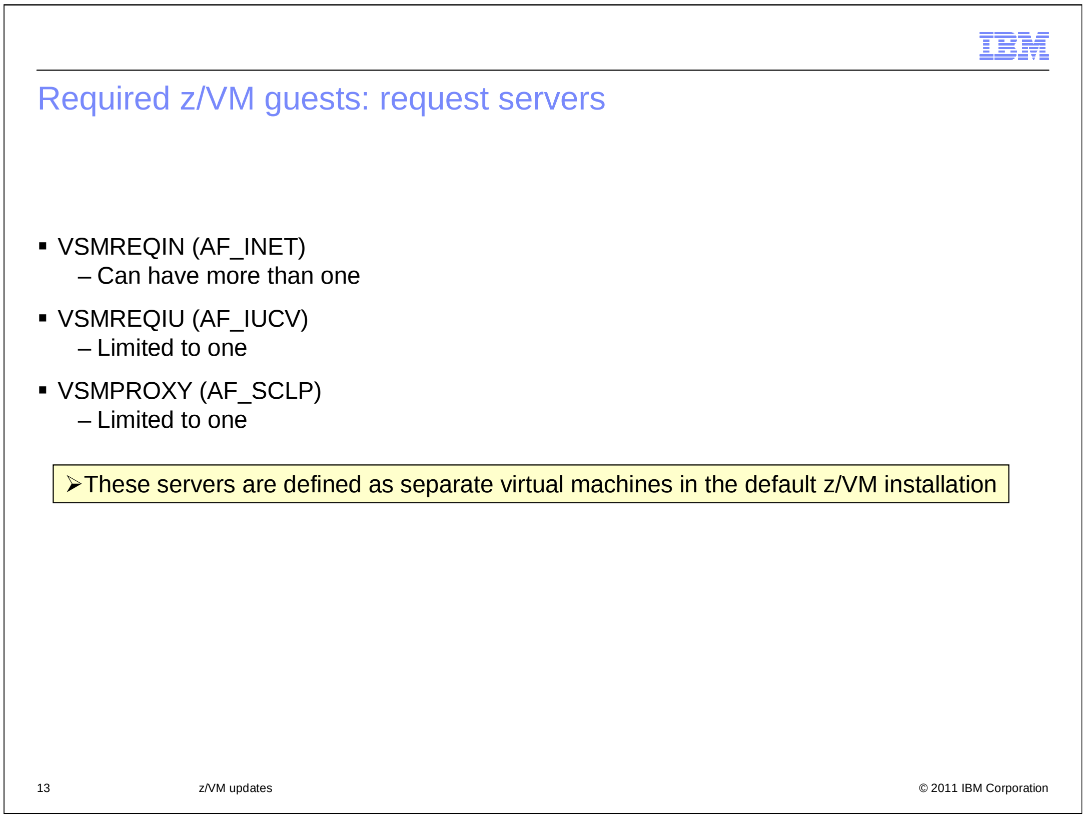

Required z/VM guests: request servers

- VSMREQIN (AF_INET)
  - Can have more than one
- VSMREQIU (AF_IUCV)
  - Limited to one
- VSMPROXY (AF_SCLP)
  - Limited to one

➢These servers are defined as separate virtual machines in the default z/VM installation

13 z/VM updates © 2011 IBM Corporation

Starting with the request servers needed by the socket-based environment, you need to define the VSMREQIN, VSMREQIU and VSMPROXY guests. These guests should already be defined for you in the default z/VM installation. Note that you can have more than one request server using the AF_INET family sockets to connect with clients but you are limited to only one AF_IUCV and one AF_SCLP request server..

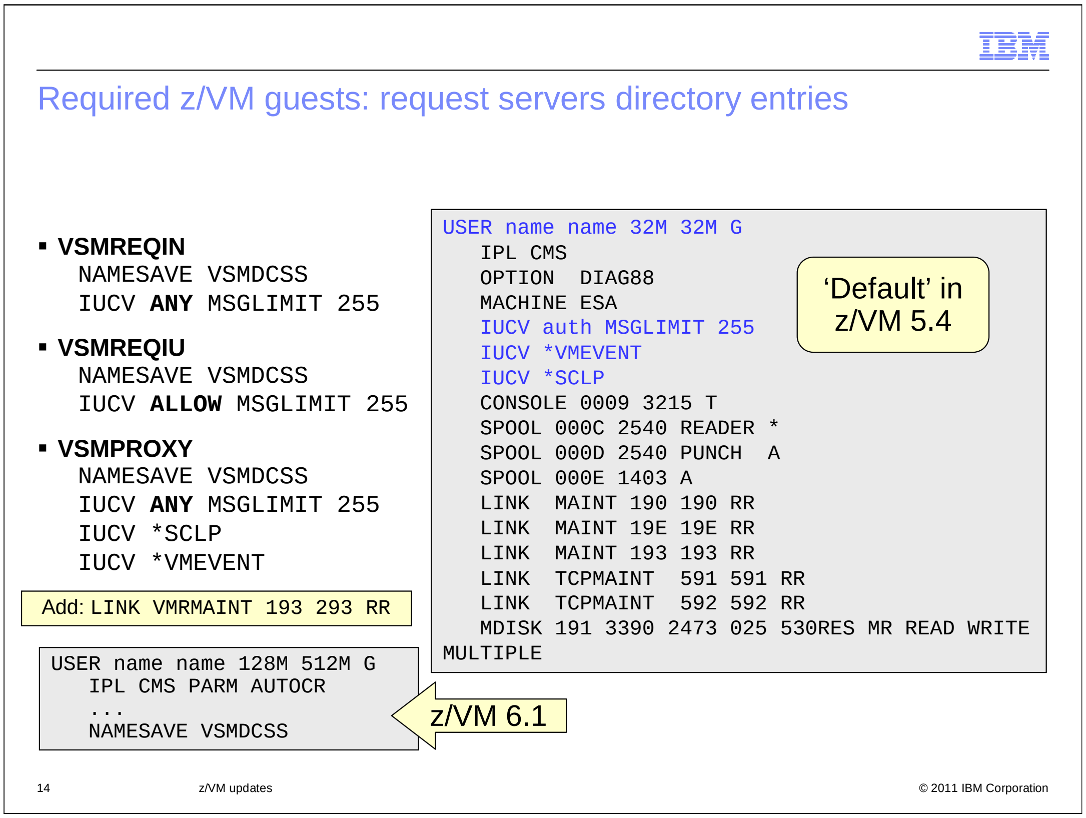

Required z/VM guests: request servers directory entries

- **VSMREQIN**

```
  NAMESAVE VSMDCSS
  IUCV ANY MSGLIMIT 255
```

- **VSMREQIU**

```
  NAMESAVE VSMDCSS
  IUCV ALLOW MSGLIMIT 255
```

- **VSMPROXY**

```
  NAMESAVE VSMDCSS
  IUCV ANY MSGLIMIT 255
  IUCV *SCLP
  IUCV *VMEVENT
```

Add: `LINK VMRMAINT 193 293 RR`

```
USER name name 128M 512M G
   IPL CMS PARM AUTOCR
   ...
   NAMESAVE VSMDCSS
```

z/VM 6.1

```
USER name name 32M 32M G
   IPL CMS
   OPTION  DIAG88
   MACHINE ESA
   IUCV auth MSGLIMIT 255
   IUCV *VMEVENT
   IUCV *SCLP
   CONSOLE 0009 3215 T
   SPOOL 000C 2540 READER *
   SPOOL 000D 2540 PUNCH  A
   SPOOL 000E 1403 A
   LINK  MAINT 190 190 RR
   LINK  MAINT 19E 19E RR
   LINK  MAINT 193 193 RR
   LINK  TCPMAINT  591 591 RR
   LINK  TCPMAINT  592 592 RR
   MDISK 191 3390 2473 025 530RES MR READ WRITE
MULTIPLE
```

‘Default’ in z/VM 5.4

14 z/VM updates © 2011 IBM Corporation

Looking at the request server directory entries, this slide shows the default directory entry in z/VM 5.4 on the right. The lines highlighted in blue are the ones that differ between the various types of request servers. The differences are shown on the left for each of the various types. In order for the request servers to access the VM:Secure code, you need to add the LINK VMRMAINT statement shown on the slide to the directory entries. In the bottom left corner, you see the differences in the default directory entries that are introduced in z/VM 6.1.

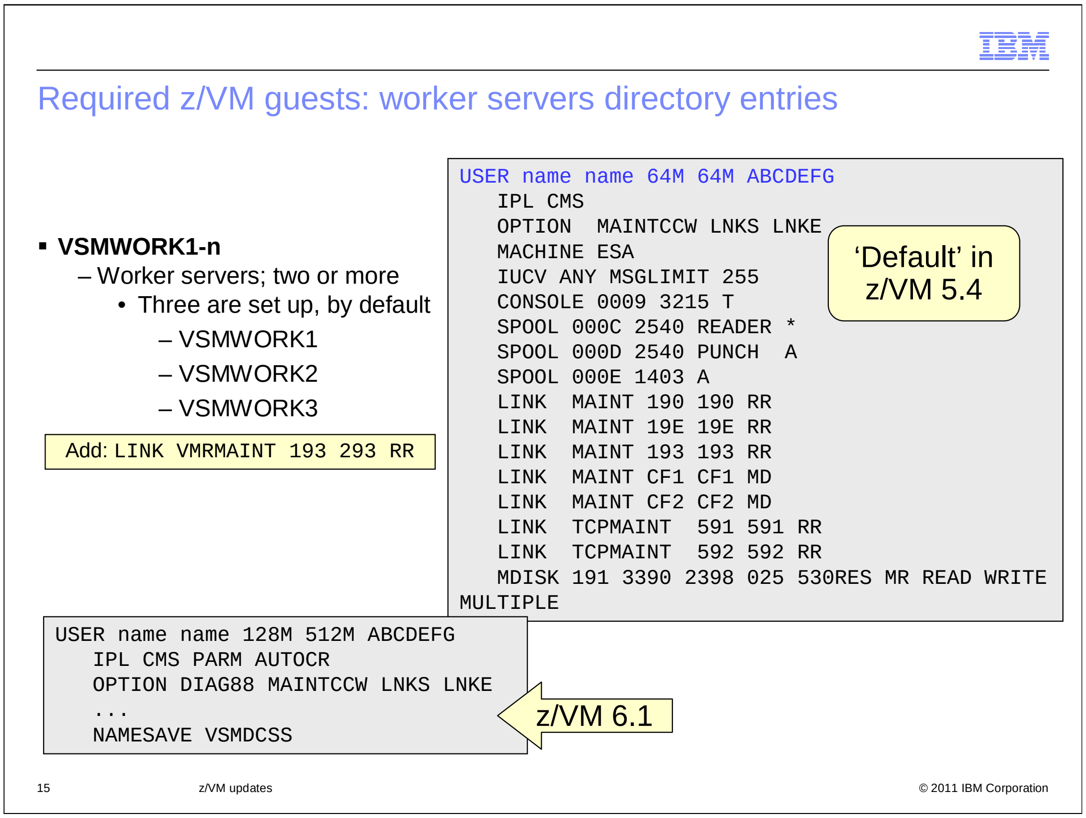

Required z/VM guests: worker servers directory entries

- **VSMWORK1-n**
  - Worker servers; two or more
    - Three are set up, by default
      - VSMWORK1
      - VSMWORK2
      - VSMWORK3

Add: `LINK VMRMAINT 193 293 RR`

```
USER name name 64M 64M ABCDEFG
   IPL CMS
   OPTION  MAINTCCW LNKS LNKE
   MACHINE ESA
   IUCV ANY MSGLIMIT 255
   CONSOLE 0009 3215 T
   SPOOL 000C 2540 READER *
   SPOOL 000D 2540 PUNCH  A
   SPOOL 000E 1403 A
   LINK  MAINT 190 190 RR
   LINK  MAINT 19E 19E RR
   LINK  MAINT 193 193 RR
   LINK  MAINT CF1 CF1 MD
   LINK  MAINT CF2 CF2 MD
   LINK  TCPMAINT  591 591 RR
   LINK  TCPMAINT  592 592 RR
   MDISK 191 3390 2398 025 530RES MR READ WRITE
MULTIPLE
```

'Default' in z/VM 5.4

```
USER name name 128M 512M ABCDEFG
   IPL CMS PARM AUTOCR
   OPTION DIAG88 MAINTCCW LNKS LNKE
   ...
   NAMESAVE VSMDCSS
```

z/VM 6.1

15 z/VM updates © 2011 IBM Corporation

This slide shows the default directory entry for the various worker servers. Again, note that you must add the LINK VMRMAINT statement to the directory entries In order for the worker servers to access the VM:Secure code.

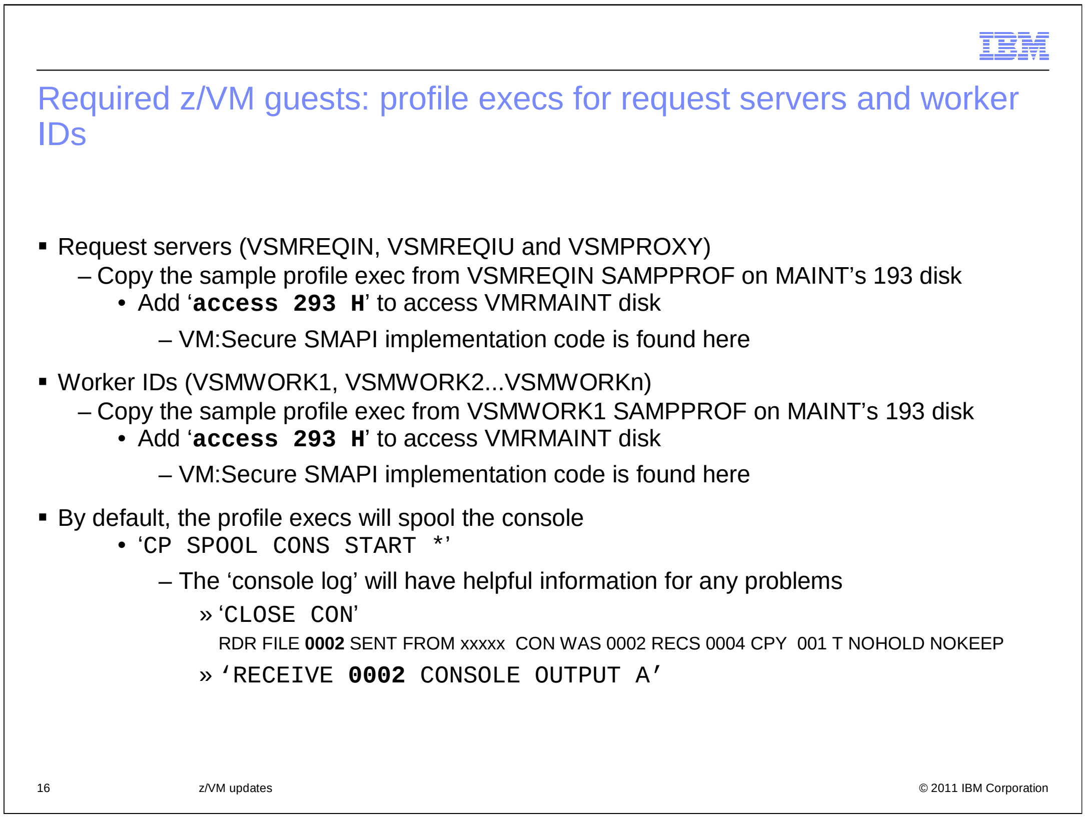

Required z/VM guests: profile execs for request servers and worker IDs

- Request servers (VSMREQIN, VSMREQIU and VSMPROXY)
  - Copy the sample profile exec from VSMREQIN SAMPPROF on MAINT’s 193 disk
    - Add ‘**`access 293 H`**’ to access VMRMAINT disk
      - VM:Secure SMAPI implementation code is found here
- Worker IDs (VSMWORK1, VSMWORK2...VSMWORKn)
  - Copy the sample profile exec from VSMWORK1 SAMPPROF on MAINT’s 193 disk
    - Add ‘**`access 293 H`**’ to access VMRMAINT disk
      - VM:Secure SMAPI implementation code is found here
- By default, the profile execs will spool the console
    - ‘`CP SPOOL CONS START *`’
      - The ‘console log’ will have helpful information for any problems
        - » ‘`CLOSE CON`’

          RDR FILE **0002** SENT FROM xxxxx CON WAS 0002 RECS 0004 CPY 001 T NOHOLD NOKEEP
        - » ‘`RECEIVE` **`0002`** `CONSOLE OUTPUT A`’

In order to logon and start the various servers required for the socket-based environment, you need to define profile execs for each guest. For the request servers, there is a sample profile exec found on MAINT’s 193 disk called VSMREQIN SAMPPROF. You should copy this to each of the request servers and add a statement to access the VMRMAINT disk that you created a link to in the directory entries.

A sample profile for the worker guests is also found on MAINT’s 193 disk called VSMWORK1 SAMPPROF. Again, you need to add a statement to access the VMRMAINT disk as shown on the slide.

By default, the profile execs will spool the console. These will often give useful information when debugging problems here.

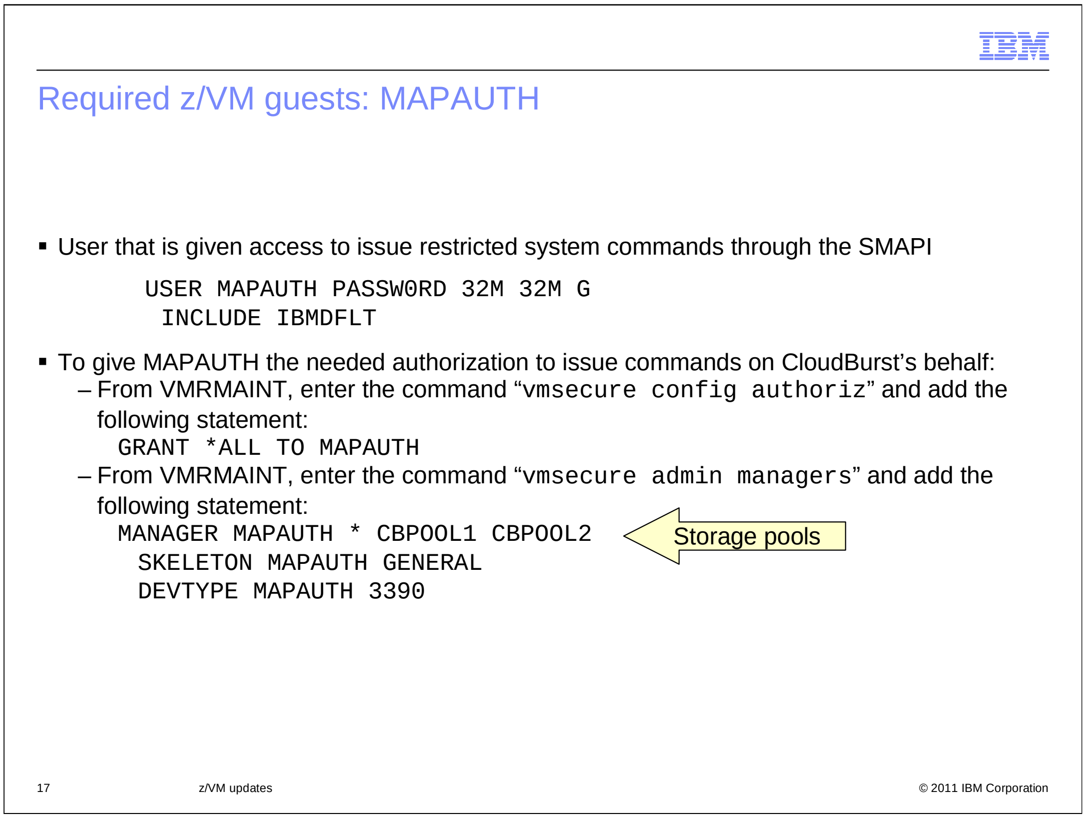

Required z/VM guests: MAPAUTH

- User that is given access to issue restricted system commands through the SMAPI

```
USER MAPAUTH PASSW0RD 32M 32M G
 INCLUDE IBMDFLT
```

- To give MAPAUTH the needed authorization to issue commands on CloudBurst’s behalf:
  - From VMRMAINT, enter the command “`vmsecure config authoriz`” and add the following statement:

    ```
    GRANT *ALL TO MAPAUTH
    ```

  - From VMRMAINT, enter the command “`vmsecure admin managers`” and add the following statement:

    ```
    MANAGER MAPAUTH * CBPOOL1 CBPOOL2
      SKELETON MAPAUTH GENERAL
      DEVTYPE MAPAUTH 3390
    ```

Storage pools

In order to issue the SMAPI commands, a user ID and password is required. This is what the MAPAUTH user ID is used for. The directory entry is very basic for MAPAUTH however you need to give the MAPAUTH guest authorization to issue commands on CloudBurst’s behalf. The authorizations shown here are all specific to the VM:Secure configuration. The first thing you need to do is give MAPAUTH authorizations to issue all commands. This is done by adding the GRANT statement shown to the ‘AUTHORIZ CONFIG’ file. This is done with the ‘vmsecure config authoriz’ command. The MAPAUTH user ID also needs to be defined to VM:Secure as a ‘manager’. You also need to give MAPAUTH access to the storage pools that you will create. To do this add the statements shown on the bottom of the slide to the VMSECURE MANAGERS file. This is done with the ‘vmsecure admin managers’ command. In the example on the slide, the storage pools are defined as CBPOOL1 and CBPOOL2.

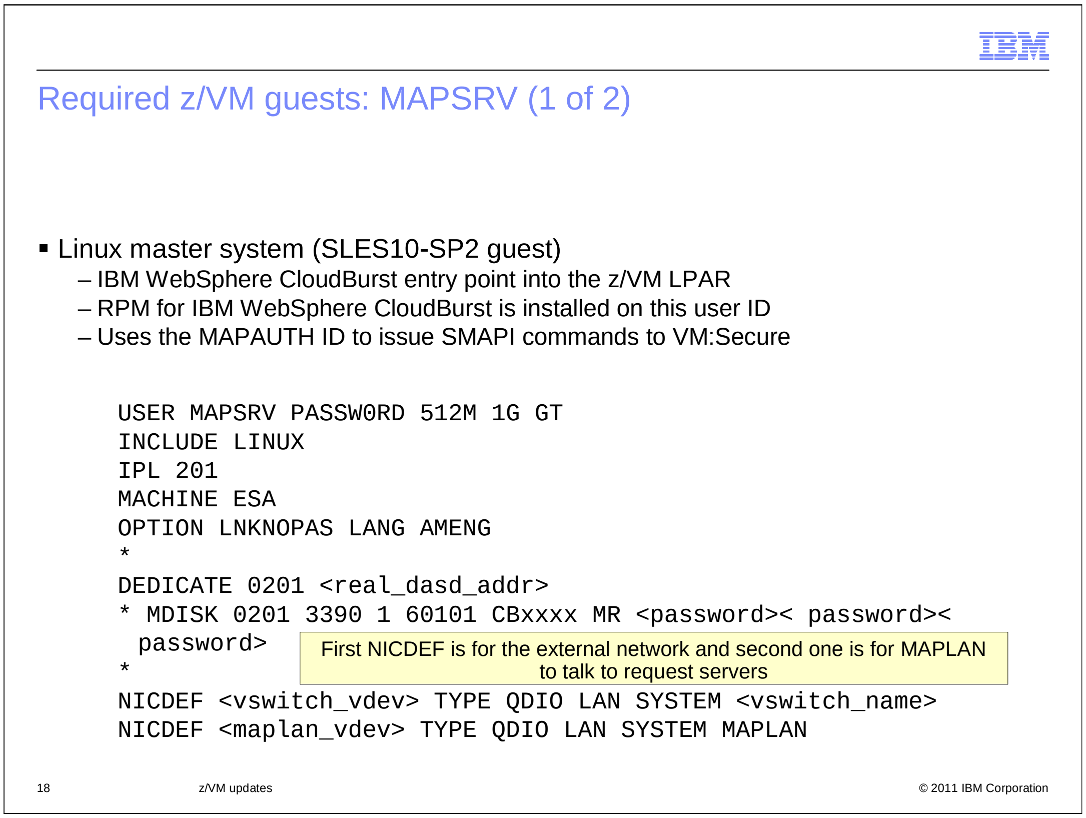

Required z/VM guests: MAPSRV (1 of 2)

- Linux master system (SLES10-SP2 guest)
  - IBM WebSphere CloudBurst entry point into the z/VM LPAR
  - RPM for IBM WebSphere CloudBurst is installed on this user ID
  - Uses the MAPAUTH ID to issue SMAPI commands to VM:Secure

```
USER MAPSRV PASSW0RD 512M 1G GT
INCLUDE LINUX
IPL 201
MACHINE ESA
OPTION LNKNOPAS LANG AMENG
*
DEDICATE 0201 <real_dasd_addr>
* MDISK 0201 3390 1 60101 CBxxxx MR <password>< password><
  password>
*
NICDEF <vswitch_vdev> TYPE QDIO LAN SYSTEM <vswitch_name>
NICDEF <maplan_vdev> TYPE QDIO LAN SYSTEM MAPLAN
```

First NICDEF is for the external network and second one is for MAPLAN to talk to request servers

18 z/VM updates © 2011 IBM Corporation

The Linux master system is a Novell SUSE Linux Enterprise Server (SLES10-SP2) guest that is used as the entry point into the z/VM logical partition (LPAR). While it can be SLES 10 or higher, increasing the size of the disks when performing an extend and capture of an image requires the SLES10-SP2 level. The RPM for IBM WebSphere CloudBurst is installed on this system and using the MAPAUTH ID the MAPSRV guest is allowed to issue SMAPI commands to VM:Secure.

This slide shows an example of a MAPSRV directory entry. The example shows a dedicated volume being used for the Linux operating system. You can also use a minidisk here, which is shown in the commented out MDISK statement. In either case, the volume is mounted as virtual DASD address 201 which is used to IPL the guest. There are two NICDEF statements defined here. The first one is used for the external network and the second one is used to for MAPLAN to talk to the request server, VSMREQIN.

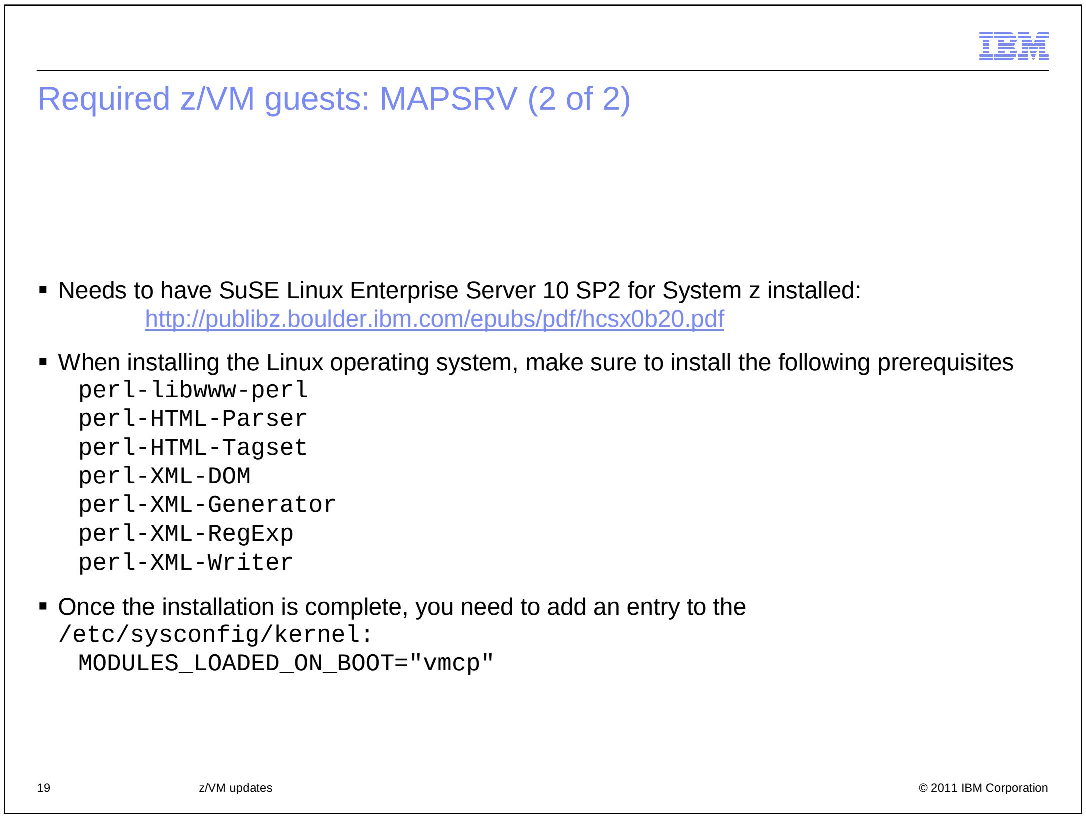

Required z/VM guests: MAPSRV (2 of 2)

- Needs to have SuSE Linux Enterprise Server 10 SP2 for System z installed:
  http://publibz.boulder.ibm.com/epubs/pdf/hcsx0b20.pdf
- When installing the Linux operating system, make sure to install the following prerequisites

```
perl-libwww-perl
perl-HTML-Parser
perl-HTML-Tagset
perl-XML-DOM
perl-XML-Generator
perl-XML-RegExp
perl-XML-Writer
```

- Once the installation is complete, you need to add an entry to the `/etc/sysconfig/kernel:`

```
MODULES_LOADED_ON_BOOT="vmcp"
```

19 z/VM updates © 2011 IBM Corporation

The MAPSRV guest needs to have SuSE Linux Enterprise Server 10 SP2 for System z installed on it. You can reference the “Getting Started with Linux on System z” document if you need help installing it to the MAPSRV guest. The packages that are required on the MAPSRV guest are listed on the slide.

The **vmcp** module/command allows z/VM CP commands to be issued from Linux. You need to specify that it be loaded on boot. Add the *MODULES_LOADED_ON_BOOT* statement to your /etc/sysconfig/kernel file after the installation completes to accomplish this.

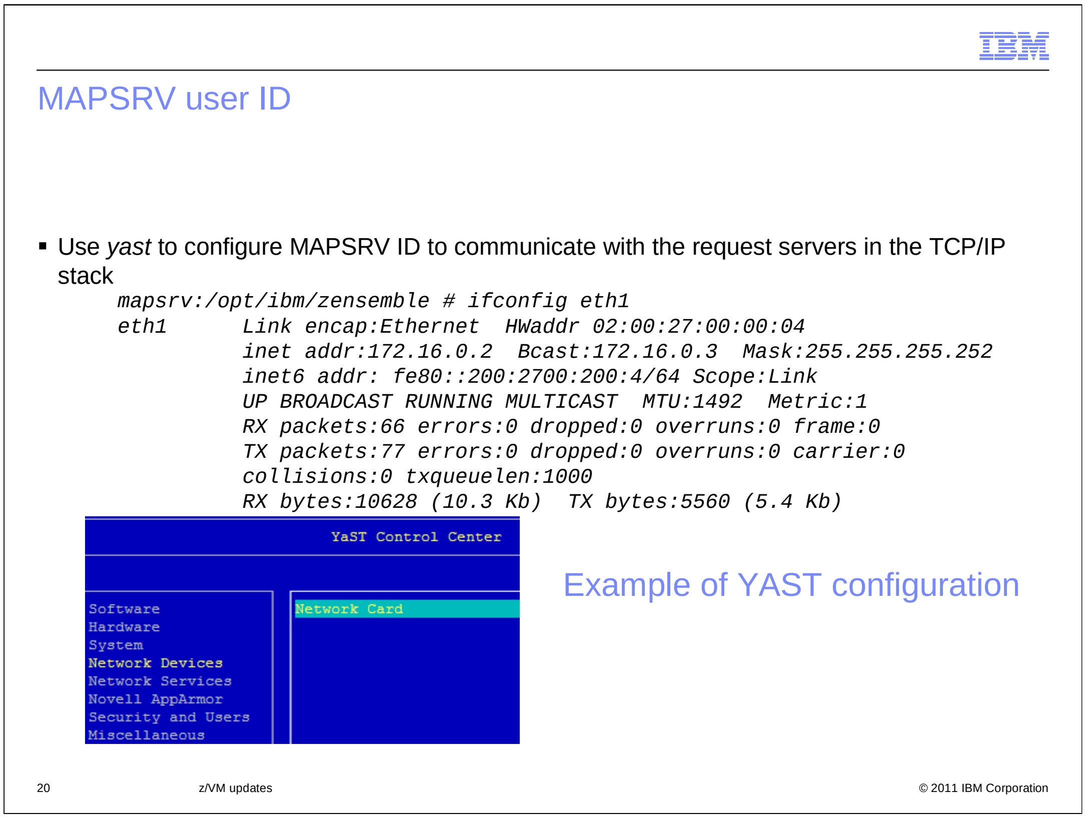

MAPSRV user ID

- Use *yast* to configure MAPSRV ID to communicate with the request servers in the TCP/IP stack

```
mapsrv:/opt/ibm/zensemble # ifconfig eth1
eth1      Link encap:Ethernet  HWaddr 02:00:27:00:00:04
          inet addr:172.16.0.2  Bcast:172.16.0.3  Mask:255.255.255.252
          inet6 addr: fe80::200:2700:200:4/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1492  Metric:1
          RX packets:66 errors:0 dropped:0 overruns:0 frame:0
          TX packets:77 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:10628 (10.3 Kb)  TX bytes:5560 (5.4 Kb)
```

Example of YAST configuration

20 z/VM updates

Finally, you need to configure your MAPSRV guest to communicate with the VSMREQIN request server by way of the TCP/IP stack. You can use ‘yast' to configure this as shown on the slide. Once configured, the resulting ifconfig will show ETH1 setup correctly to talk to VSMREQIN as shown on the slide. The “**z/VM System Setup”** document has more detailed instructions on how to configure this if needed.

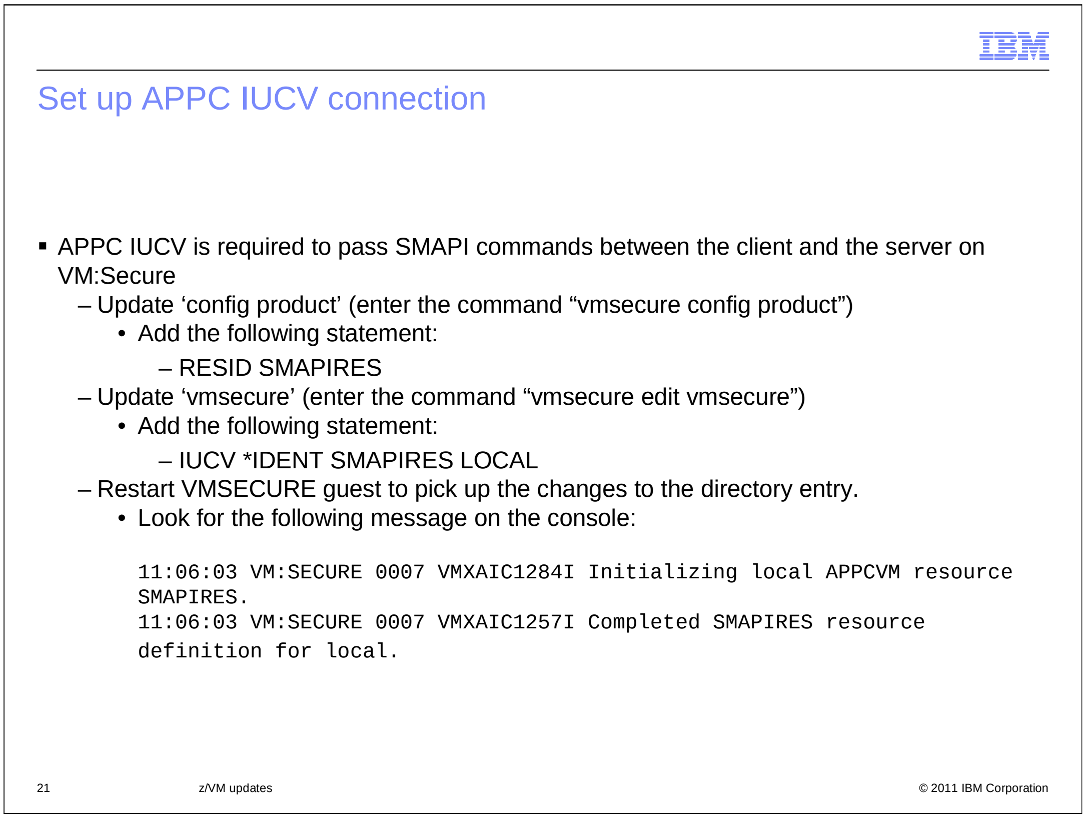

Set up APPC IUCV connection

- APPC IUCV is required to pass SMAPI commands between the client and the server on VM:Secure
  - – Update ‘config product’ (enter the command “vmsecure config product”)
    - • Add the following statement:
      - – RESID SMAPIRES
  - – Update ‘vmsecure’ (enter the command “vmsecure edit vmsecure”)
    - • Add the following statement:
      - – IUCV *IDENT SMAPIRES LOCAL
  - – Restart VMSECURE guest to pick up the changes to the directory entry.
    - • Look for the following message on the console:

```
11:06:03 VM:SECURE 0007 VMXAIC1284I Initializing local APPCVM resource
SMAPIRES.
11:06:03 VM:SECURE 0007 VMXAIC1257I Completed SMAPIRES resource
definition for local.
```

21 z/VM updates

For VM:Secure, APPC IUCV is required to pass SMAPI commands between the client and the server. This slide shows a few VM:Secure commands needed to accomplish this.

Update the PRODUCT CONFIG file with the RESID statement shown and the VMSECURE directory entry with the IUCV statement shown. The VMSECURE guest needs to be restarted to pick up these changes. Once these changes are made, you can check the VMSECURE console log for the messages shown at the bottom of the slide to be sure the configuration is correct.

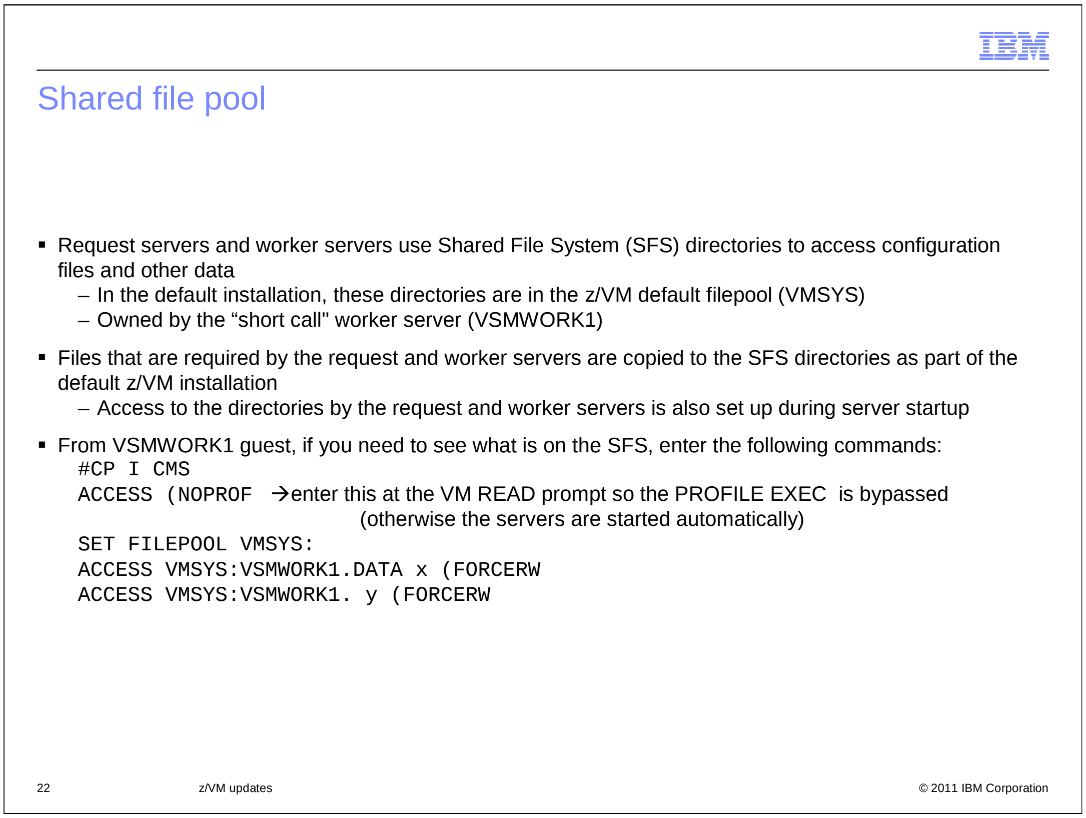

Shared file pool

- Request servers and worker servers use Shared File System (SFS) directories to access configuration files and other data
  - In the default installation, these directories are in the z/VM default filepool (VMSYS)
  - Owned by the “short call" worker server (VSMWORK1)
- Files that are required by the request and worker servers are copied to the SFS directories as part of the default z/VM installation
  - Access to the directories by the request and worker servers is also set up during server startup
- From VSMWORK1 guest, if you need to see what is on the SFS, enter the following commands:

```
#CP I CMS
ACCESS (NOPROF  →enter this at the VM READ prompt so the PROFILE EXEC  is bypassed
                      (otherwise the servers are started automatically)
SET FILEPOOL VMSYS:
ACCESS VMSYS:VSMWORK1.DATA x (FORCERW
ACCESS VMSYS:VSMWORK1. y (FORCERW
```

22 z/VM updates © 2011 IBM Corporation

The configuration files for the request and worker servers are kept in a shared file pool. Both the request and worker servers can access data here as well. In the default installation, these directories are in the z/VM default filepool, VMSYS. The default directories are owned by the “short call" worker server, VSMWORK1. Files that are required by the request and worker servers are copied to the SFS directories as part of the default z/VM installation. Looking in the PROFILE EXECs for the servers, you see that access to the directories by the request and worker servers is set up during server startup.

If you need to see or change the configuration files for the servers, commands to accomplish that are at the bottom of the slide. If you want to access the shared file pool directories from the actual server guests, you first need to issue an ‘ACCESS (NOPROF’ command. This bypasses running the PROFILE EXEC so the servers are not started. The ACCESS commands then allow you to assign a disk, represented as ‘x’ and ‘y’ in the example to the shared file pool directories. You can then look at the files using regular CMS commands such as FILELIST.

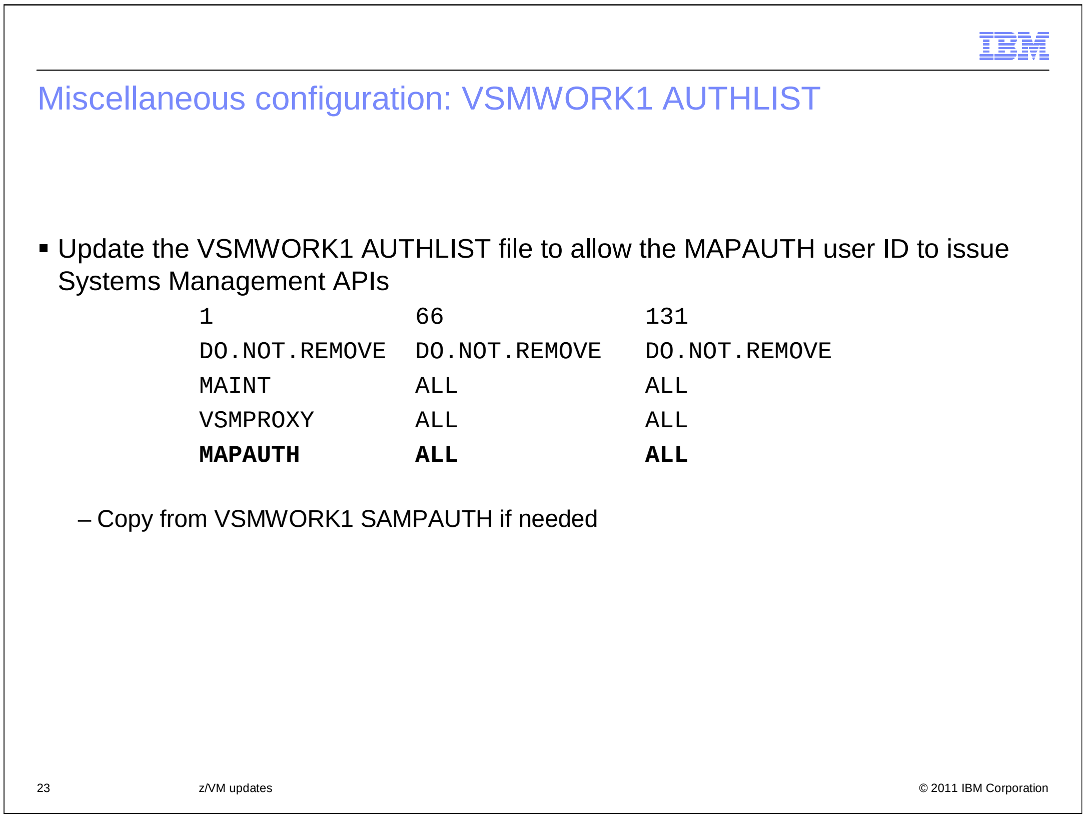

Miscellaneous configuration: VSMWORK1 AUTHLIST

- Update the VSMWORK1 AUTHLIST file to allow the MAPAUTH user ID to issue Systems Management APIs

```
1               66                131
DO.NOT.REMOVE   DO.NOT.REMOVE     DO.NOT.REMOVE
MAINT           ALL               ALL
VSMPROXY        ALL               ALL
MAPAUTH         ALL               ALL
```

  – Copy from VSMWORK1 SAMPAUTH if needed

In order for the MAPAUTH guest to issue SMAPI commands, he needs to be added to the VSMWORK1 AUTHLIST file. That is shown here. Be sure to copy one of the existing lines in the file as the file is very dependent on the position of the keywords. The numbers represent the columns that the parameters are expected in. The third parameter must start in column 131. A sample is found on MAINT’s 193 disk as VSMWORK1 SAMPAUTH.

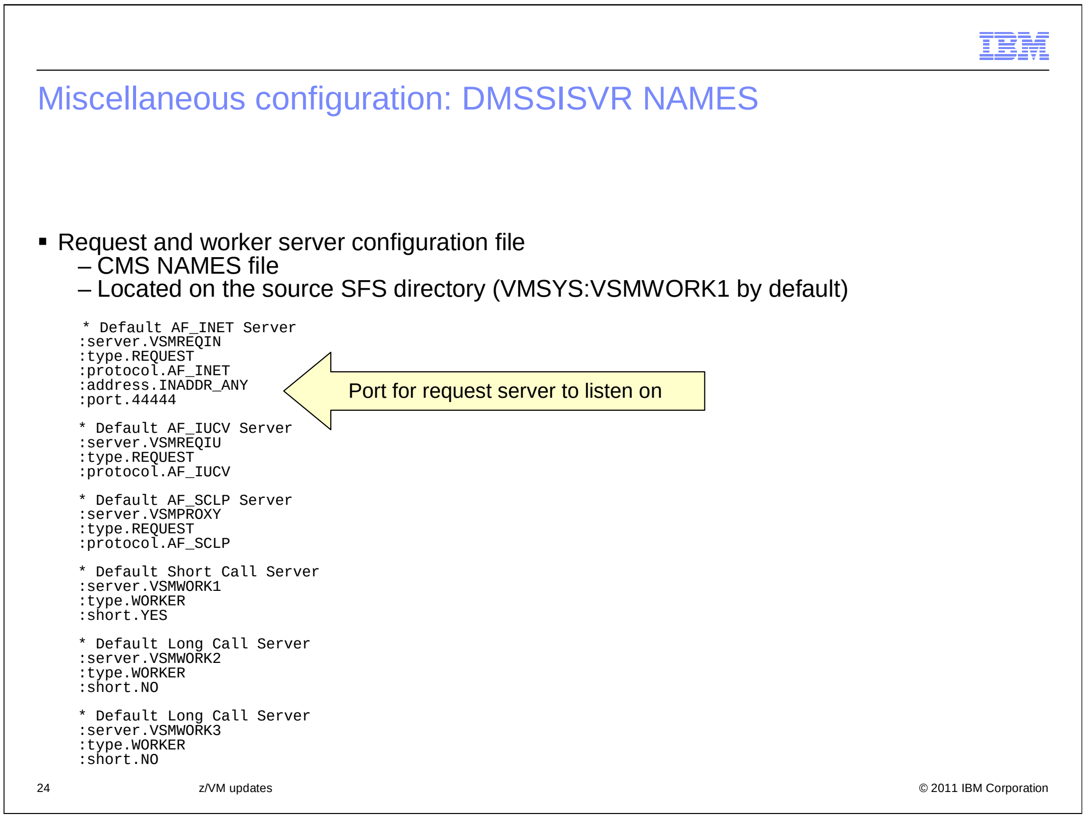

Miscellaneous configuration: DMSSISVR NAMES

- Request and worker server configuration file
  - CMS NAMES file
  - Located on the source SFS directory (VMSYS:VSMWORK1 by default)

```
 * Default AF_INET Server
:server.VSMREQIN
:type.REQUEST
:protocol.AF_INET
:address.INADDR_ANY
:port.44444

* Default AF_IUCV Server
:server.VSMREQIU
:type.REQUEST
:protocol.AF_IUCV

* Default AF_SCLP Server
:server.VSMPROXY
:type.REQUEST
:protocol.AF_SCLP

* Default Short Call Server
:server.VSMWORK1
:type.WORKER
:short.YES

* Default Long Call Server
:server.VSMWORK2
:type.WORKER
:short.NO

* Default Long Call Server
:server.VSMWORK3
:type.WORKER
:short.NO
```

Port for request server to listen on

24 z/VM updates © 2011 IBM Corporation

Looking at some of the configuration files next, this slide shows the DMSSISVR NAMES file, or the server names file. It is a CMS NAMES file that determines how each specific request and worker server will function in the overall server environment. It is located on the source SFS directory, VMSYS:VSMWORK1, by default. A sample is found on MAINT's 193 disk. Most of the sample can remain unchanged, however, one field you should note is the port number for the VSMREQIN server. The default is '44444' so if your system uses a different port, this value needs to be changed.

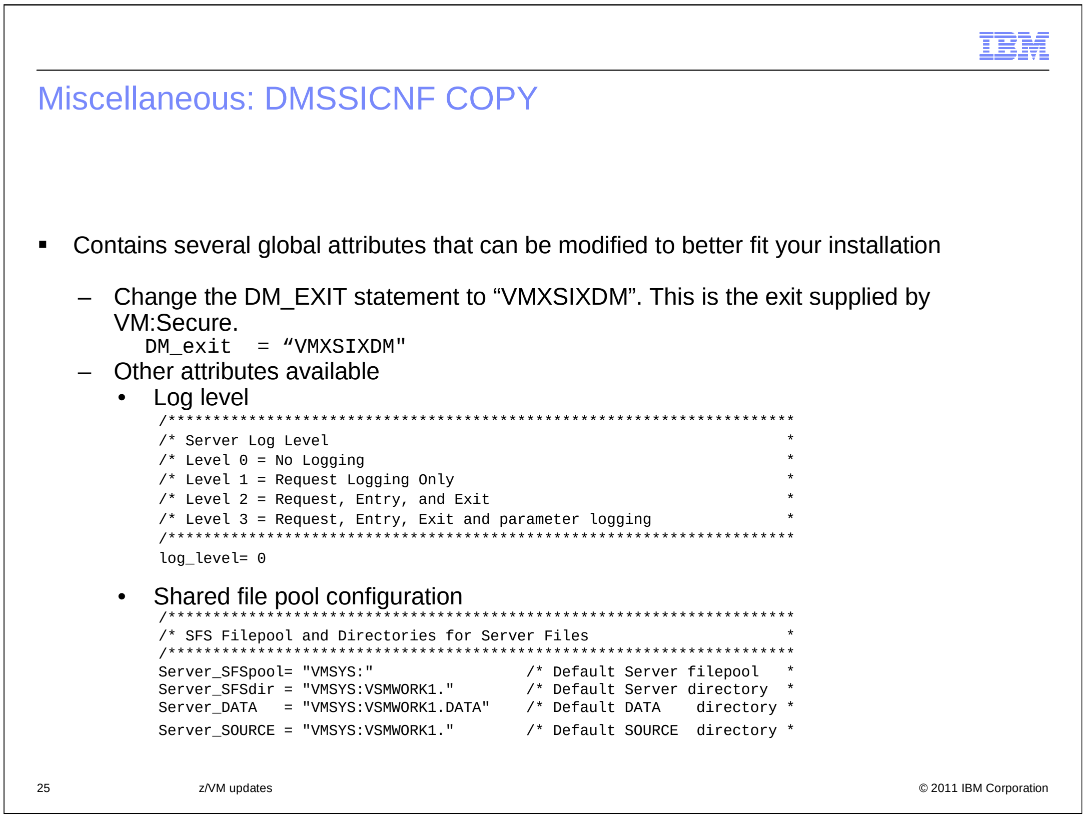

Miscellaneous: DMSSICNF COPY

- Contains several global attributes that can be modified to better fit your installation
  - Change the DM_EXIT statement to “VMXSIXDM”. This is the exit supplied by VM:Secure.

    ```
    DM_exit  = “VMXSIXDM"
    ```

  - Other attributes available
    - Log level

      ```
      /**********************************************************************
      /* Server Log Level                                                   *
      /* Level 0 = No Logging                                               *
      /* Level 1 = Request Logging Only                                     *
      /* Level 2 = Request, Entry, and Exit                                 *
      /* Level 3 = Request, Entry, Exit and parameter logging               *
      /**********************************************************************
      log_level= 0
      ```

    - Shared file pool configuration

      ```
      /**********************************************************************
      /* SFS Filepool and Directories for Server Files                      *
      /**********************************************************************
      Server_SFSpool= "VMSYS:"                /* Default Server filepool   *
      Server_SFSdir = "VMSYS:VSMWORK1."       /* Default Server directory  *
      Server_DATA   = "VMSYS:VSMWORK1.DATA"   /* Default DATA    directory *
      Server_SOURCE = "VMSYS:VSMWORK1."       /* Default SOURCE  directory *
      ```

The next configuration file is the DMSSICNF COPY file. This is an important file for the VM:Secure configuration. This is where you need to point to the VM:Secure directory manager exit instead of the DirMaint directory manager exit. In order to use VM:Secure, the DM_exit parameter needs to be set to ‘VMXSIXDM’. This file also has some other important configuration attributes that you may want to change. The first one is ‘log_level’. This determines how much logging takes place for the server environment. The various levels are shown on the slide but it is important to note that the default is ‘no logging’. The other parameters that you may want to change involve the shared file pool. If you want to use a shared file pool other than the default VMSYS, you need to specify that in the DMSSICNF COPY file as shown on the slide.

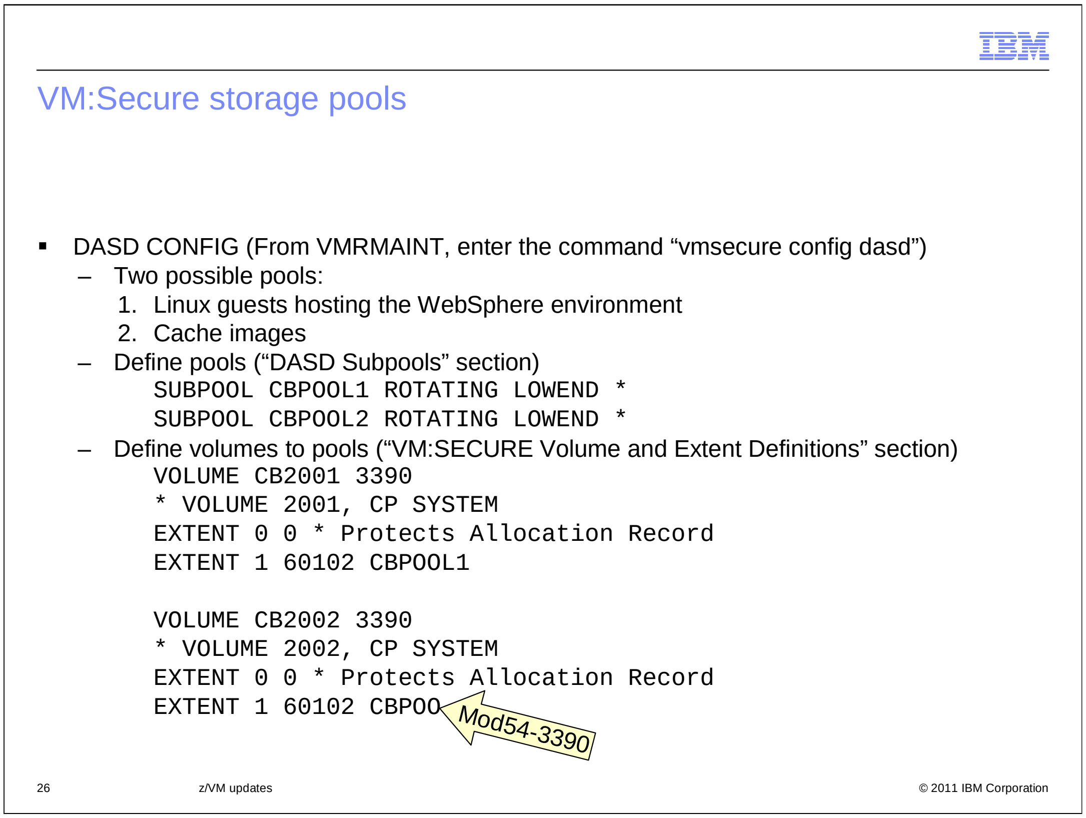

VM:Secure storage pools

- DASD CONFIG (From VMRMAINT, enter the command “vmsecure config dasd”)
  - Two possible pools:
    1. Linux guests hosting the WebSphere environment
    2. Cache images
  - Define pools (“DASD Subpools” section)

```
SUBPOOL CBPOOL1 ROTATING LOWEND *
SUBPOOL CBPOOL2 ROTATING LOWEND *
```

  - Define volumes to pools (“VM:SECURE Volume and Extent Definitions” section)

```
VOLUME CB2001 3390
* VOLUME 2001, CP SYSTEM
EXTENT 0 0 * Protects Allocation Record
EXTENT 1 60102 CBPOOL1

VOLUME CB2002 3390
* VOLUME 2002, CP SYSTEM
EXTENT 0 0 * Protects Allocation Record
EXTENT 1 60102 CBPOO
```

Mod54-3390

Recall that there are two possible storage pools that can be used by IBM WebSphere CloudBurst. The first one is required and is used to define the disks needed for the guests that IBM WebSphere CloudBurst deploys. The second one is optional and is used to cache images, if specified. The same pool can be used for both functions but separate pools are recommended. In order to define these pools to VM:Secure, you need to add the SUBPOOL statements to the “DASD Subpools” section. You also need to define the volumes that the pools will use in the ‘VM:Secure Volume and Extent Definitions’ section. In the example shown here the CB2001 volume is used for the CBPOOL1 storage pool and the CB2002 volume is used for the CBPOOL2 storage pool. Both volumes are of type MOD54-3390.

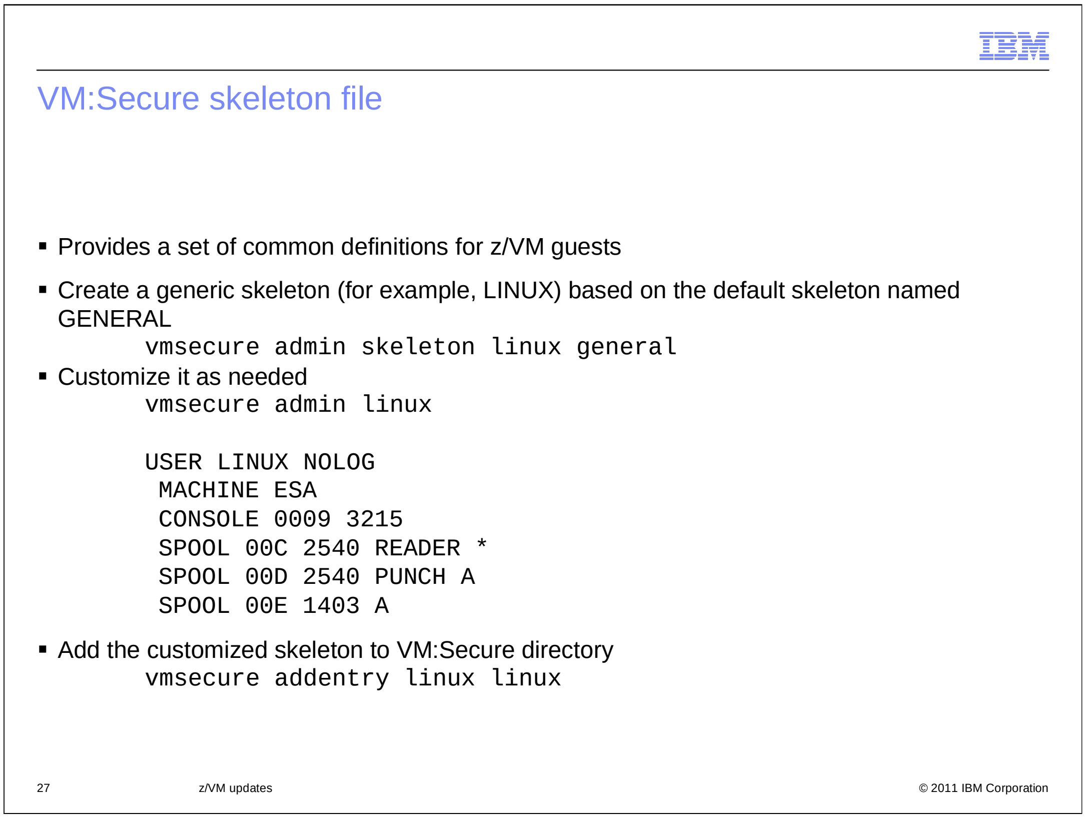

VM:Secure skeleton file

- Provides a set of common definitions for z/VM guests
- Create a generic skeleton (for example, LINUX) based on the default skeleton named GENERAL

```
      vmsecure admin skeleton linux general
```

- Customize it as needed

```
      vmsecure admin linux

      USER LINUX NOLOG
       MACHINE ESA
       CONSOLE 0009 3215
       SPOOL 00C 2540 READER *
       SPOOL 00D 2540 PUNCH A
       SPOOL 00E 1403 A
```

- Add the customized skeleton to VM:Secure directory

```
      vmsecure addentry linux linux
```

VM:Secure has the concept of a skeleton file, which is equivalent to the prototype directory entries created for the DirMaint directory manager. Skeleton files are used to define common definitions for all provisioned guests. You are asked for its name when configuring the IBM WebSphere CloudBurst for deployments. The first statement shown on the slide creates a skeleton named LINUX based on the skeleton named SKELETON. Once you have the ‘LINUX’ skeleton customized for your environment, you can use the VM:Secure addentry command to add it to the VM:Secure directory. You will see later how to point IBM WebSphere CloudBurst to it for its use.

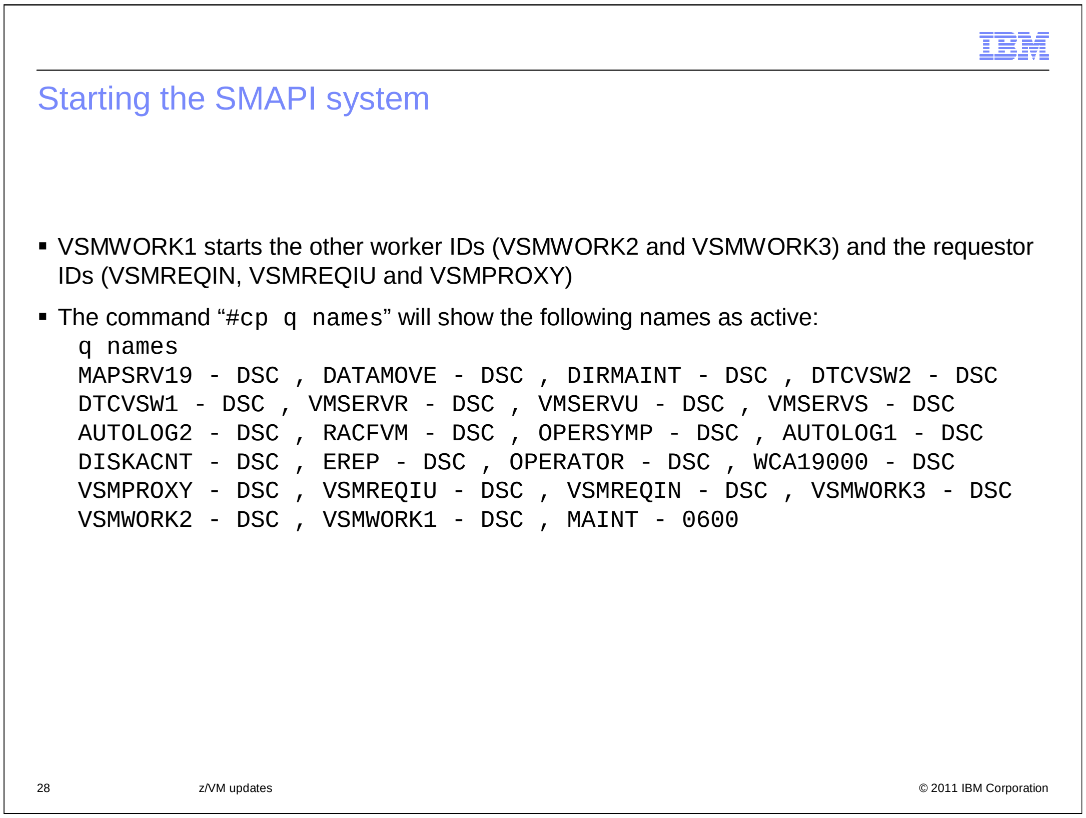

Starting the SMAPI system

- VSMWORK1 starts the other worker IDs (VSMWORK2 and VSMWORK3) and the requestor IDs (VSMREQIN, VSMREQIU and VSMPROXY)
- The command “`#cp q names`” will show the following names as active:

```
q names
MAPSRV19 - DSC , DATAMOVE - DSC , DIRMAINT - DSC , DTCVSW2 - DSC
DTCVSW1 - DSC , VMSERVR - DSC , VMSERVU - DSC , VMSERVS - DSC
AUTOLOG2 - DSC , RACFVM - DSC , OPERSYMP - DSC , AUTOLOG1 - DSC
DISKACNT - DSC , EREP - DSC , OPERATOR - DSC , WCA19000 - DSC
VSMPROXY - DSC , VSMREQIU - DSC , VSMREQIN - DSC , VSMWORK3 - DSC
VSMWORK2 - DSC , VSMWORK1 - DSC , MAINT - 0600
```

Once all the guests and configuration files are set up and ready to go, you need to actually start the socket-based server environment for SMAPI. This is done by starting the short call worker server, VSMWORK1. VSMWORK1 will automatically start the other worker IDs and requester IDs. Once VSMWORK1 is started, you can issue the ‘q names’ command to verify all the guests are active. You should see VSMWORK2, VSMWORK3, VSMREQIN, VSMREQIU and VSMPROXY. You will want to add the VSMWORK1 to AUTOLOG so that it is automatically started when the system is IPLed.

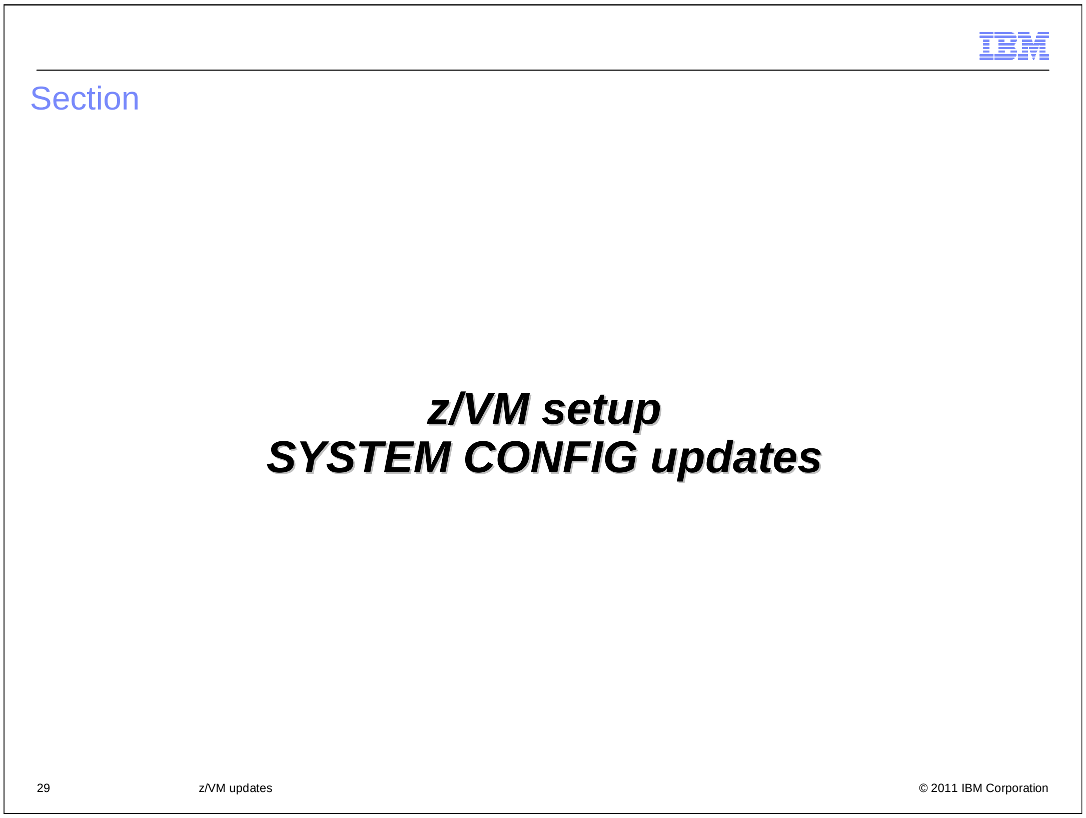

Section

z/VM setup SYSTEM CONFIG updates

This section will present SYSTEM CONFIG updates that are required on your system. These updates are required for both the VM:Secure and the DirMaint directory manager configurations.

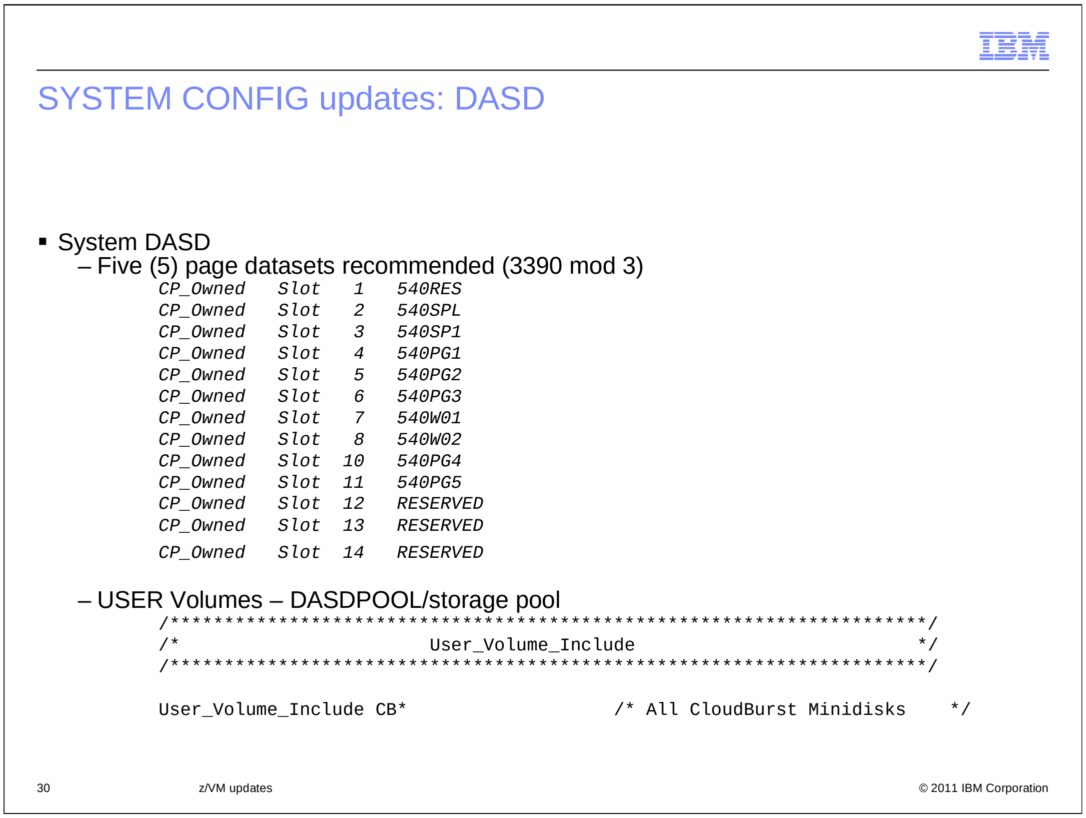

SYSTEM CONFIG updates: DASD

- System DASD
  - Five (5) page datasets recommended (3390 mod 3)

```
CP_Owned   Slot   1   540RES
CP_Owned   Slot   2   540SPL
CP_Owned   Slot   3   540SP1
CP_Owned   Slot   4   540PG1
CP_Owned   Slot   5   540PG2
CP_Owned   Slot   6   540PG3
CP_Owned   Slot   7   540W01
CP_Owned   Slot   8   540W02
CP_Owned   Slot  10   540PG4
CP_Owned   Slot  11   540PG5
CP_Owned   Slot  12   RESERVED
CP_Owned   Slot  13   RESERVED
CP_Owned   Slot  14   RESERVED
```

  - USER Volumes – DASDPOOL/storage pool

```
/***********************************************************************/
/*                       User_Volume_Include                           */
/***********************************************************************/

User_Volume_Include CB*                  /* All CloudBurst Minidisks    */
```

You want to make sure that there are enough page and spool volumes defined for the LPAR that will host IBM WebSphere CloudBurst. In order to improve system performance for z/VM, you should define five page datasets to the system. 3390 Mod 3 DASD is the recommended type of DASD to use for both SPOOL and PAGE datasets. You use the command CPFMTXA to format the disks as type PAGE and make sure you provide a unique label for each PAGE dataset, for instance 530PG1-530PG5, as shown above. You should also add additional reserved slots so that page volumes can be dynamically added to the system as needed. It is important to note that you must not re-arrange the volumes in the CP_Owned list. Moving SPOOL volumes will result in deleted z/VM segments and other catastrophic errors. When changing the configuration, always append new volumes to the end of the list.

User volumes are required to define minidisks to the z/VM system. This is required for IBM WebSphere CloudBurst since all target disks for provisioning Linux guests are obtained from a DASDPOOL or storage pool. The DASDPOOL, or storage pool, is a group of minidisks managed by the directory manager and the z/VM system. As a shortcut, you can format all the minidisks with the same first two characters. As an example, you can use CBxxxx, where xxxx is the DASD address. This will allow you to use a wildcard when defining the User_Volume_List statement in the SYSTEM CONFIG.

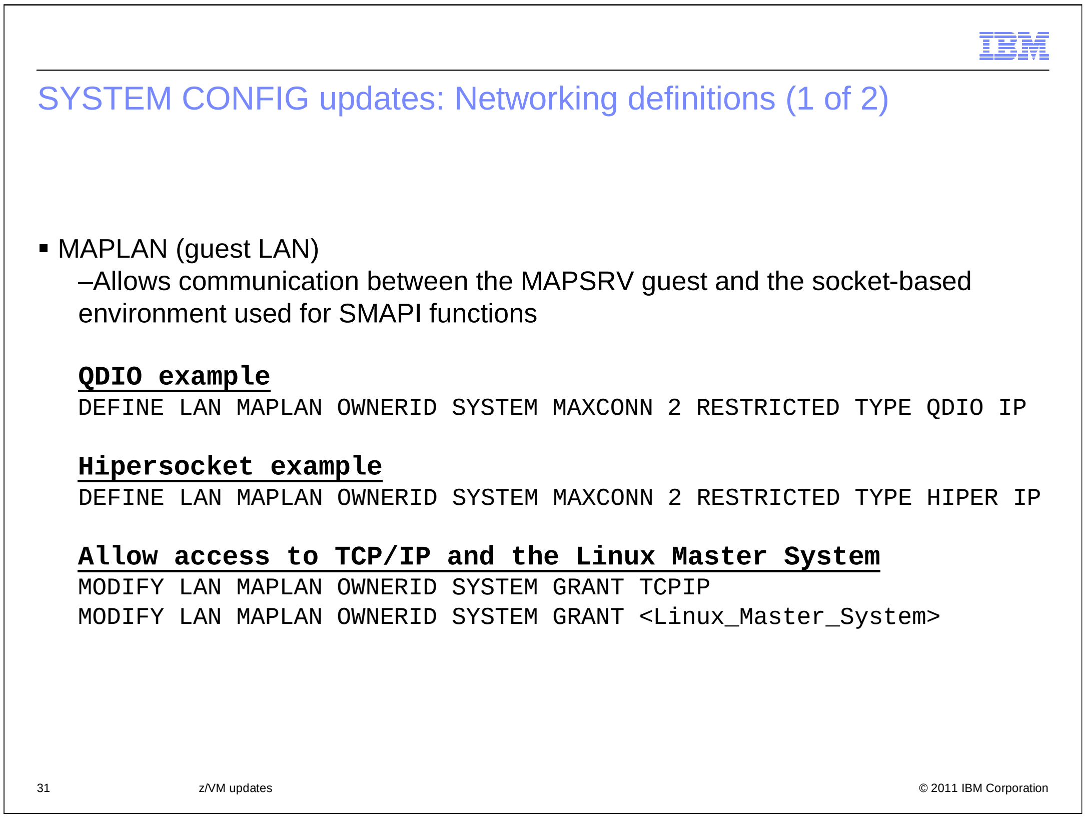

SYSTEM CONFIG updates: Networking definitions (1 of 2)

- MAPLAN (guest LAN)
  - –Allows communication between the MAPSRV guest and the socket-based environment used for SMAPI functions

QDIO example

```
DEFINE LAN MAPLAN OWNERID SYSTEM MAXCONN 2 RESTRICTED TYPE QDIO IP
```

Hipersocket example

```
DEFINE LAN MAPLAN OWNERID SYSTEM MAXCONN 2 RESTRICTED TYPE HIPER IP
```

Allow access to TCP/IP and the Linux Master System

```
MODIFY LAN MAPLAN OWNERID SYSTEM GRANT TCPIP
MODIFY LAN MAPLAN OWNERID SYSTEM GRANT <Linux_Master_System>
```

31 z/VM updates

There are a minimum of two network definitions required in the SYSTEM CONFIG: MAPLAN and VSWITCH. The MAPLAN is used to allow communication between the MAPSRV guest and the socket-based environment for SMAPI functions. Two examples are shown. The MAXCONN statement limits this to two connections.

Note that both TCP/IP and the Linux Master system both need to be granted access to the MAPLAN. This is shown on the bottom of the slide.

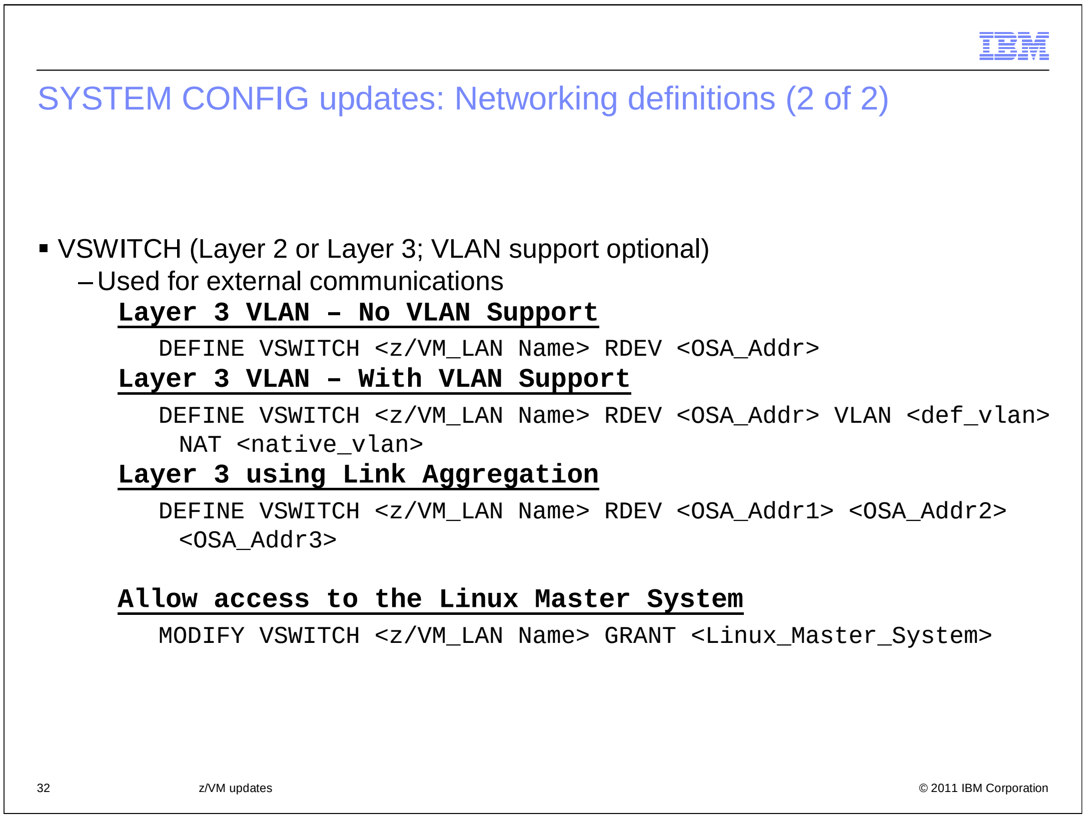

SYSTEM CONFIG updates: Networking definitions (2 of 2)

- VSWITCH (Layer 2 or Layer 3; VLAN support optional)
  - Used for external communications

Layer 3 VLAN – No VLAN Support

```
DEFINE VSWITCH <z/VM_LAN Name> RDEV <OSA_Addr>
```

Layer 3 VLAN – With VLAN Support

```
DEFINE VSWITCH <z/VM_LAN Name> RDEV <OSA_Addr> VLAN <def_vlan>
  NAT <native_vlan>
```

Layer 3 using Link Aggregation

```
DEFINE VSWITCH <z/VM_LAN Name> RDEV <OSA_Addr1> <OSA_Addr2>
  <OSA_Addr3>
```

Allow access to the Linux Master System

```
MODIFY VSWITCH <z/VM_LAN Name> GRANT <Linux_Master_System>
```

32 z/VM updates © 2011 IBM Corporation

The VSWITCH, shown here, is used for external communications. The VSWITCH can be Layer 2 or Layer 3 and VLAN support is optional. VSWITCH virtualizes a single OSA adapter or multiple OSAs can be aggregated using LACP to provide redundancy, failover and additional bandwidth, if required.

Network definitions for provisioned systems must be in the main SYSTEM CONFIG file. Using imbed files for network definitions is not supported.

By not coding CON CONTROLLER, default owners DTCVSW1 and DTCVSW2 are used.
These guests should be added to the AUTOLOG PROFILE EXEC.
The Linux Master System needs to be given access to this VSWITCH.

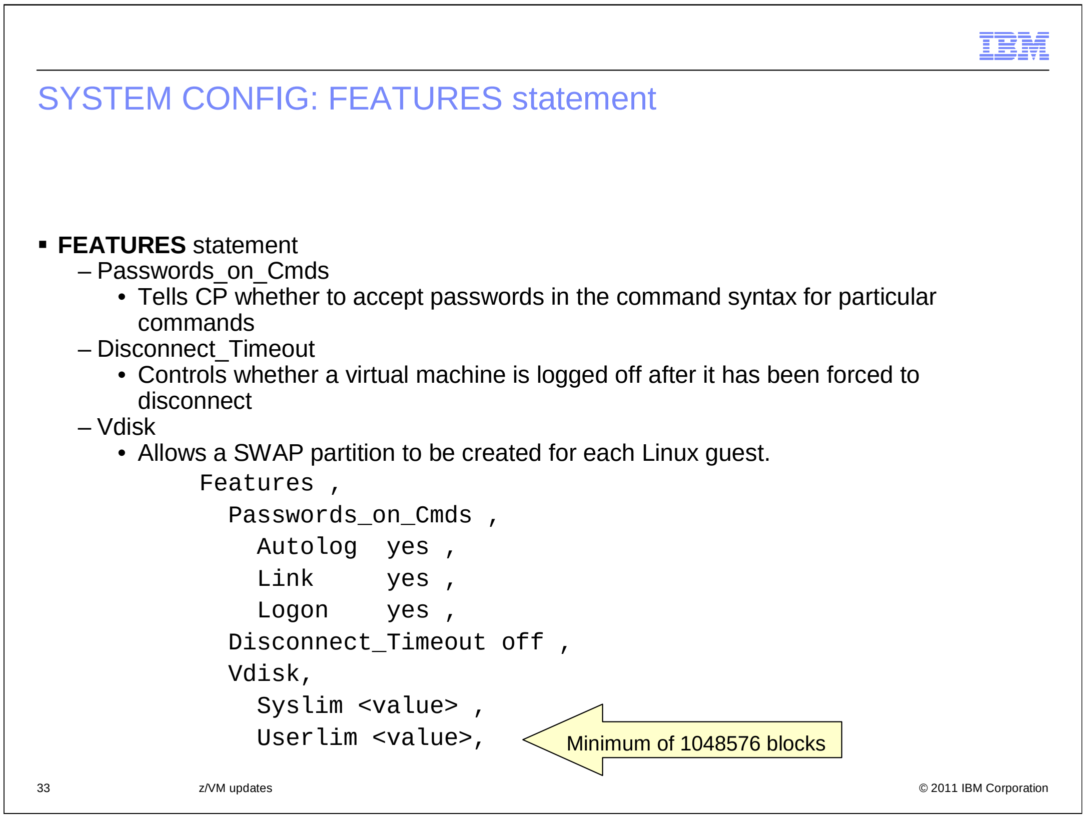

SYSTEM CONFIG: FEATURES statement

- **FEATURES** statement
  - Passwords_on_Cmds
    - Tells CP whether to accept passwords in the command syntax for particular commands
  - Disconnect_Timeout
    - Controls whether a virtual machine is logged off after it has been forced to disconnect
  - Vdisk
    - Allows a SWAP partition to be created for each Linux guest.

```
Features ,
  Passwords_on_Cmds ,
    Autolog  yes ,
    Link     yes ,
    Logon    yes ,
  Disconnect_Timeout off ,
  Vdisk,
    Syslim <value> ,
    Userlim <value>,
```

Minimum of 1048576 blocks

33 z/VM updates © 2011 IBM Corporation

The FEATURES statement in SYSTEM CONFIG allows you to modify attributes associated with the running system at IPL time. IBM WebSphere CloudBurst requires that you allow Passwords_on Cmds. This feature tells CP whether to accept passwords in the command syntax (in clear text) when users issue the CP AUTOLOG, LINK, or LOGON commands.

The Disconnect_Timeout feature controls whether and when a virtual machine is logged off after it has been forced to disconnect. You will turn this feature off, so that any virtual machine that has been forced to disconnect is not logged off.

For Linux guests deployed by IBM WebSphere CloudBurst, VDISK is used for the swap disk, requiring 512 MB per guest, as a performance enhancement. VDISK space comes from the z/VM dynamic paging area, main memory, and is allocated as needed. Be sure to account for this use of VDISK when you are planning the memory capacity for the z/VM system. Using VDISK puts additional stress on PAGE datasets, so plan the Syslim value according to the amount of PAGE space you have allocated on your system. The Userlim value must be a minimum of 1048576 blocks to allow the SWAP partition to be created for each Linux guest.

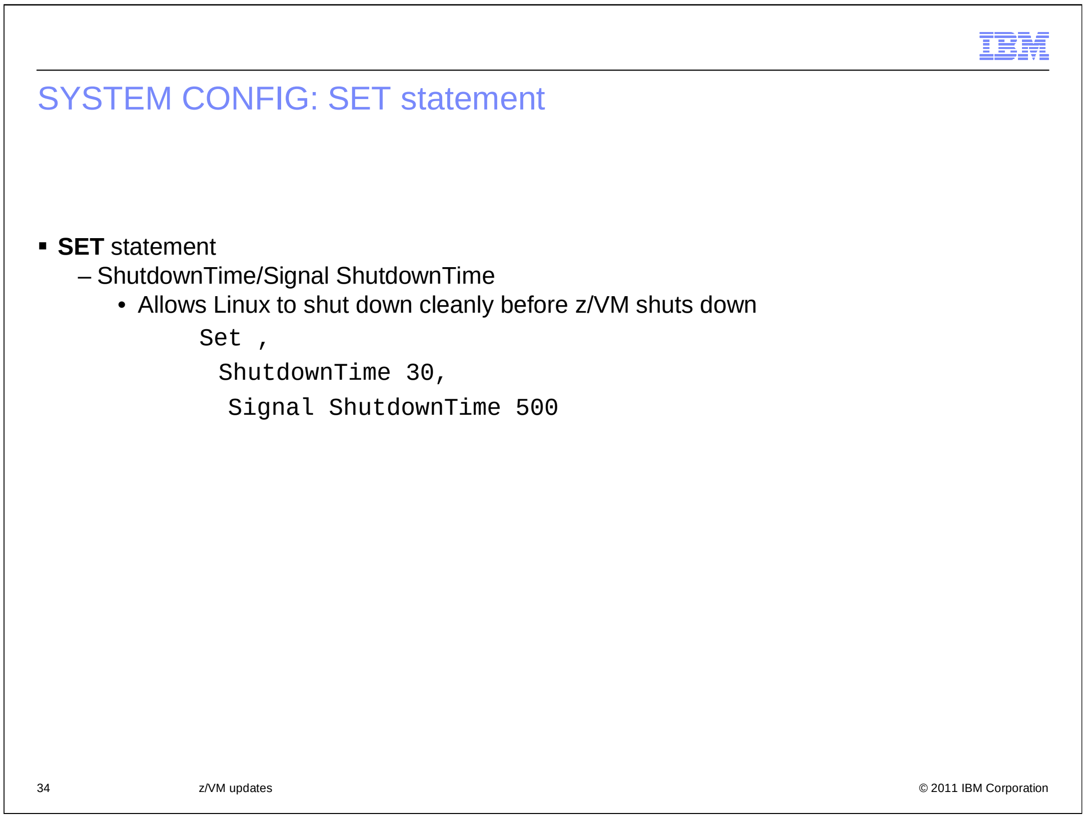

SYSTEM CONFIG: SET statement

- **SET** statement
  - ShutdownTime/Signal ShutdownTime
    - Allows Linux to shut down cleanly before z/VM shuts down

```
Set ,
  ShutdownTime 30,
   Signal ShutdownTime 500
```

The ShutdownTime and Signal ShutdownTime system configuration values enable a virtual machine to register with CP to receive a shutdown signal when z/VM is shutting down. CP waits to shut itself down until the time interval (in seconds) is exceeded, or all of the virtual machines enabled for the signal shutdown have reported a successful shutdown. Linux distributions support this function, which allows Linux to shut down cleanly before z/VM shuts down.

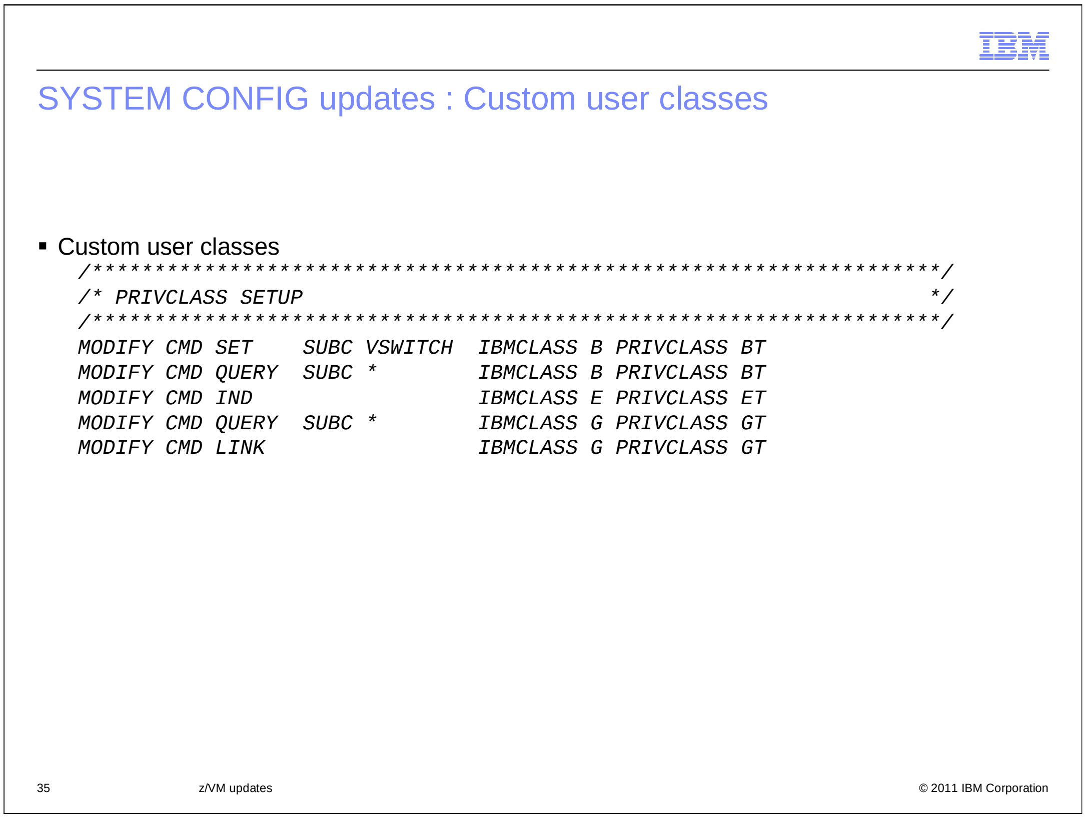

SYSTEM CONFIG updates : Custom user classes

- Custom user classes

```
/*********************************************************************/
/* PRIVCLASS SETUP                                                   */
/*********************************************************************/
MODIFY CMD SET    SUBC VSWITCH  IBMCLASS B PRIVCLASS BT
MODIFY CMD QUERY  SUBC *        IBMCLASS B PRIVCLASS BT
MODIFY CMD IND                  IBMCLASS E PRIVCLASS ET
MODIFY CMD QUERY  SUBC *        IBMCLASS G PRIVCLASS GT
MODIFY CMD LINK                 IBMCLASS G PRIVCLASS GT
```

The Modify commands shown on the slide need to be added to SYSTEM CONFIG to allow the Linux Master system (MAPSRV) to manage and update the virtual network devices.

Once all the changes have been implemented in your SYSTEM CONFIG, you will need to IPL the z/VM LPAR to enable the changes.

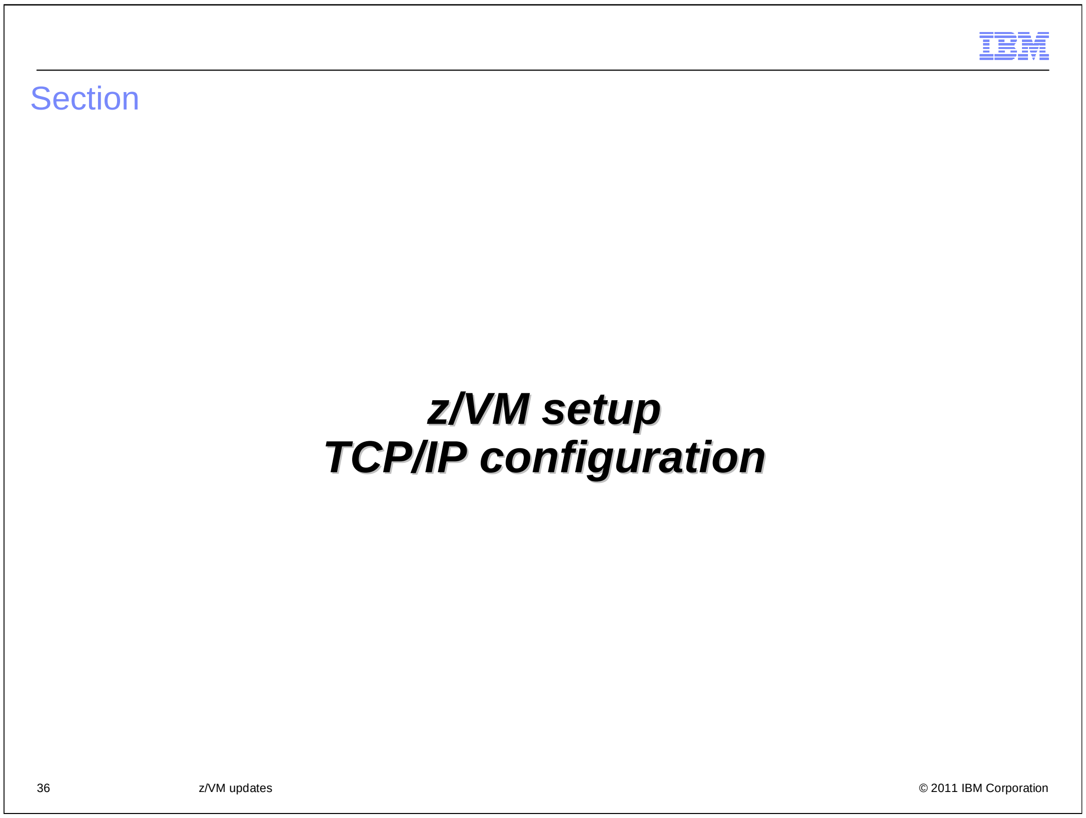

Section

z/VM setup
TCP/IP configuration

This section will present the TCP/IP configuration changes that are required, Again these updates are required for both the VM:Secure and the DirMaint directory manager configurations.

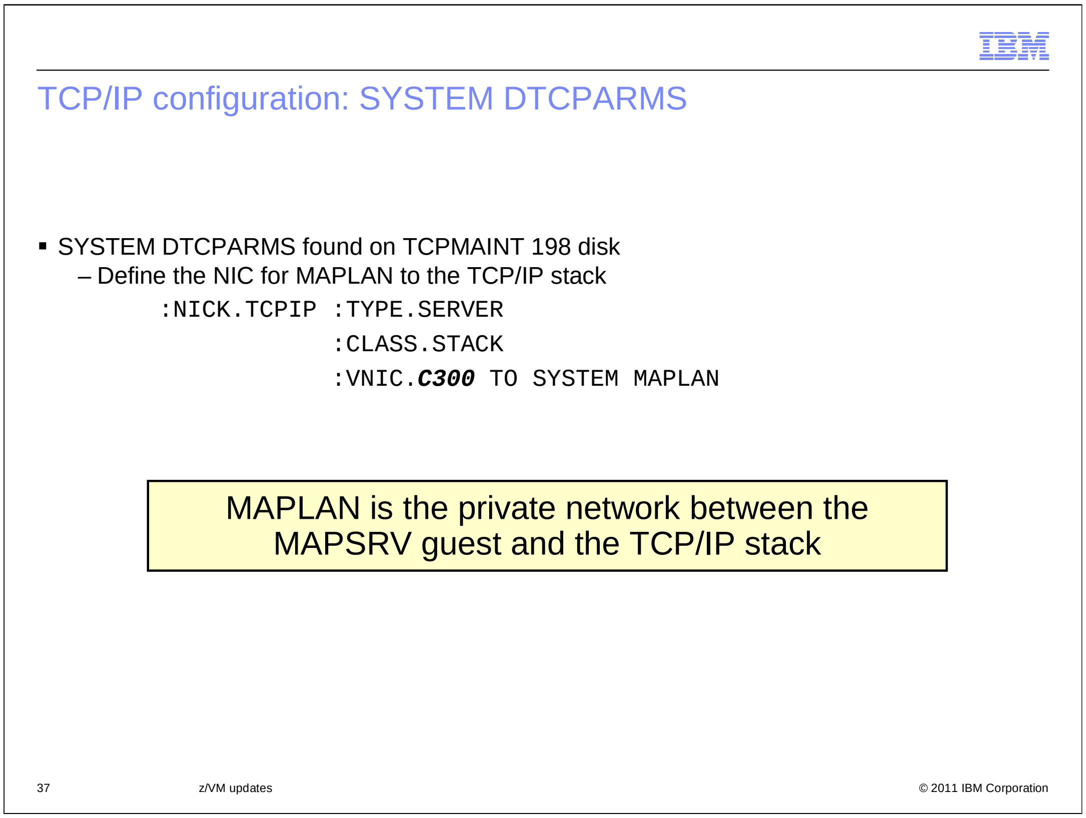

TCP/IP configuration: SYSTEM DTCPARMS

- SYSTEM DTCPARMS found on TCPMAINT 198 disk
  - Define the NIC for MAPLAN to the TCP/IP stack

```
:NICK.TCPIP :TYPE.SERVER
            :CLASS.STACK
            :VNIC.C300 TO SYSTEM MAPLAN
```

MAPLAN is the private network between the MAPSRV guest and the TCP/IP stack

The first step is to define the network interface card (NIC) for MAPLAN to the TCP/IP stack. This is the private network between the MAPSRV guest and the TCP/IP stack. In the example shown above, you will define NIC C300 to the system. This is done in the SYSTEM DTCPARMS file found on the TCPMAINT 198 disk.

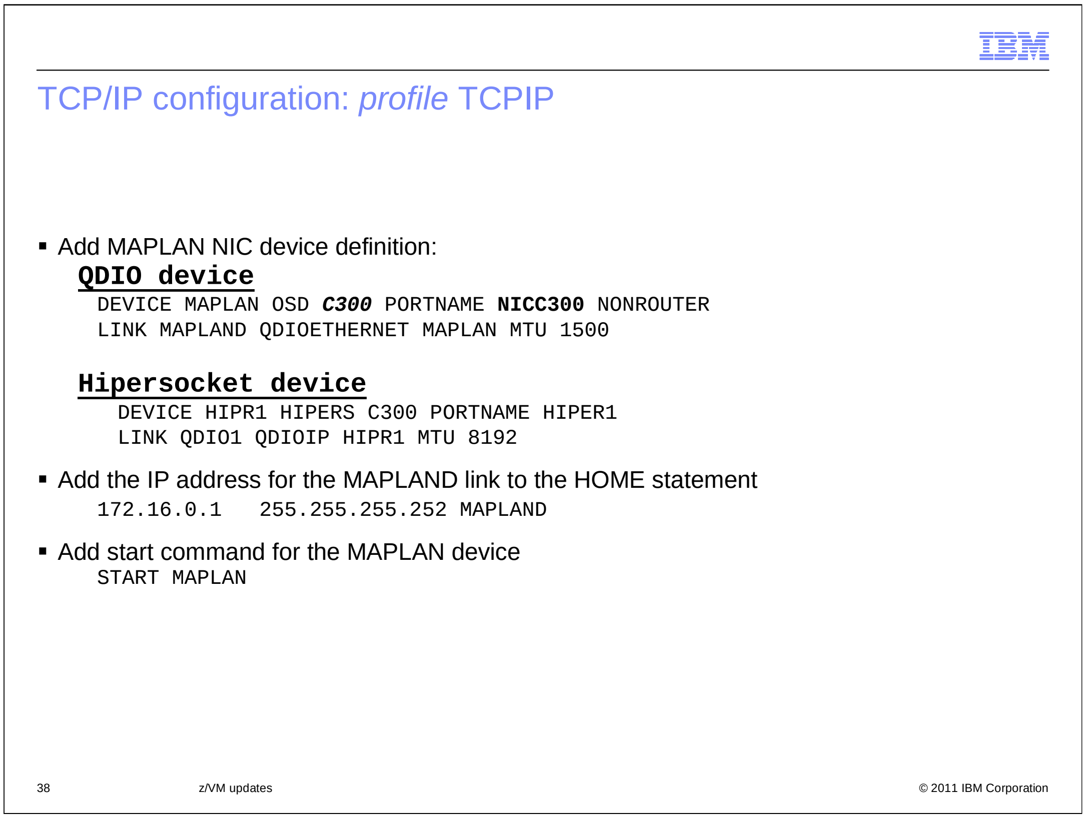

TCP/IP configuration: *profile* TCPIP

- Add MAPLAN NIC device definition:

QDIO device

```
DEVICE MAPLAN OSD C300 PORTNAME NICC300 NONROUTER
LINK MAPLAND QDIOETHERNET MAPLAN MTU 1500
```

Hipersocket device

```
DEVICE HIPR1 HIPERS C300 PORTNAME HIPER1
LINK QDIO1 QDIOIP HIPR1 MTU 8192
```

- Add the IP address for the MAPLAND link to the HOME statement

```
172.16.0.1   255.255.255.252 MAPLAND
```

- Add start command for the MAPLAN device

```
START MAPLAN
```

38 z/VM updates © 2011 IBM Corporation

Now that the MAPLAN NIC is defined to the guest, you must add the device to the TCP/IP profile. On the previous slide, you defined device C300 as the NIC for accessing MAPLAN. Now you will use this address when defining the DEVICE statement. The PORTNAME value is arbitrary; portnames NICC300 and HIPER1 are used in the examples on the slide.

Next you must add the IP address for the MAPLAND link to the HOME statement as seen on the slide. Finally, at the bottom of the TCP/IP profile, there is a section that will start the devices when the TCP/IP guest is started. Make sure you add the start command for the MAPLAN device.

Once the required changes have been made to the TCP/IP profile, save the changes and restart the TCP/IP stack.

Section

Prototype/skeleton override file

This next section presents a function that is new for IBM WebSphere CloudBurst V2.0.0.3 when used with a z/VM hypervisor. It allows you to define a prototype/skeleton override file so that you can use different prototype or skeleton files depending on the image part being provisioned.

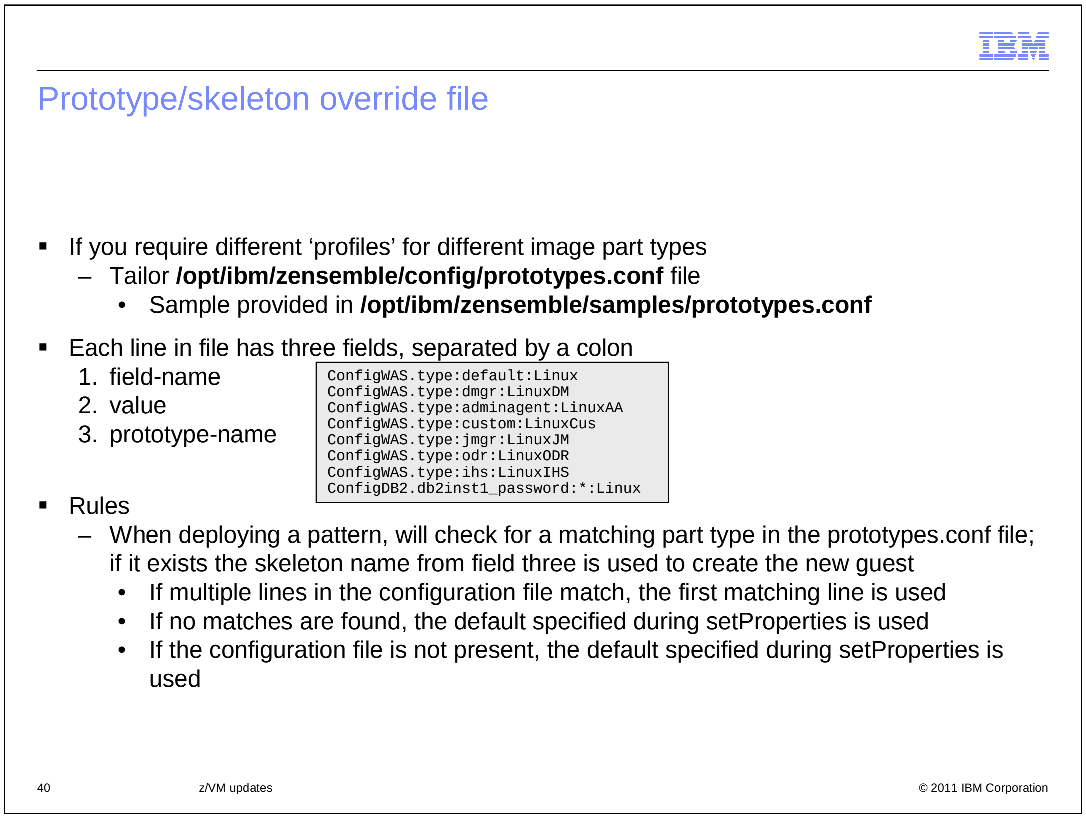

Prototype/skeleton override file

- If you require different ‘profiles’ for different image part types
  - Tailor **/opt/ibm/zensemble/config/prototypes.conf** file
    - Sample provided in **/opt/ibm/zensemble/samples/prototypes.conf**
- Each line in file has three fields, separated by a colon
  1. field-name
  2. value
  3. prototype-name

```
ConfigWAS.type:default:Linux
ConfigWAS.type:dmgr:LinuxDM
ConfigWAS.type:adminagent:LinuxAA
ConfigWAS.type:custom:LinuxCus
ConfigWAS.type:jmgr:LinuxJM
ConfigWAS.type:odr:LinuxODR
ConfigWAS.type:ihs:LinuxIHS
ConfigDB2.db2inst1_password:*:Linux
```

- Rules
  - When deploying a pattern, will check for a matching part type in the prototypes.conf file; if it exists the skeleton name from field three is used to create the new guest
    - If multiple lines in the configuration file match, the first matching line is used
    - If no matches are found, the default specified during setProperties is used
    - If the configuration file is not present, the default specified during setProperties is used

40 z/VM updates

In order to use different prototypes or skeletons for different image part types, you need to tailor a prototypes.conf file in the /opt/ibm/zensemble/config directory. There is a sample provided in the /opt/ibm/zensemble/samples directory. Each line in the file has three fields, separated by a colon: field-name, value and prototype-name. In the example, you can see that when deploying a dmgr part for the WebSphere Application Server image, you want to use the LinuxDM prototype. When deploying a pattern, IBM WebSphere CloudBurst will first check for a match for the pattern type in the prototypes.conf file. If a match is found, it will use the specified prototype. If no matches are found, the default specified during setProperties is used. It will also use the default specified during setProperties if the prototypes.conf file does not exist. If more than one match is found, it will use the first matching line. You will see information on setProperties command in the next section that talks about the RPM for IBM WebSphere CloudBurst.

Section

z/VM setup
IBM Websphere CloudBurst RPM

41 z/VM updates © 2011 IBM Corporation

This section will discuss the installation and use of the RPM zensemble file to allow for the provision and management of a Linux virtual machine on a z/VM system.

IBM WebSphere CloudBurst RPM (1 of 3)

- Install the *zensemble RPM* file that is delivered with the firmware
  - Go to **Cloud > Cloud Groups** and click your defined Cloud
  - A field, **Required z/VM agent:** is shown with clickable links titled
    - **Download agent** – for VM:Secure installations
    - **Download agent (RPC)** – for DIRM installations

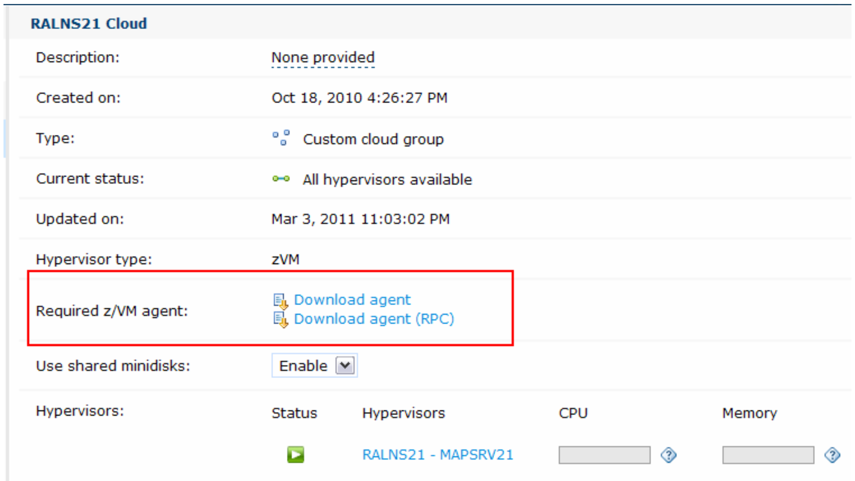

RALNS21 Cloud

Description:	None provided
Created on:	Oct 18, 2010 4:26:27 PM
Type:	Custom cloud group
Current status:	All hypervisors available
Updated on:	Mar 3, 2011 11:03:02 PM
Hypervisor type:	zVM
Required z/VM agent:	Download agent Download agent (RPC)
Use shared minidisks:	Enable

Hypervisors:

Status	Hypervisors	CPU	Memory
▶	RALNS21 - MAPSRV21		

The *zensemble* RPM file needs to be downloaded using the IBM WebSphere CloudBurst appliance administration console. Select Cloud -> Cloud Groups and then your defined cloud. You will see a field called “Required z/VM agent” with two clickable links titled ‘Download agent’ and ‘Download agent (RPC)’. For VM:Secure, select the ‘Download agent’ link. ‘Download agent (RPC)’ is used with configurations using DirMaint.

IBM WebSphere CloudBurst RPM (2 of 3)

- Save and upload the file to your MAPSRV id

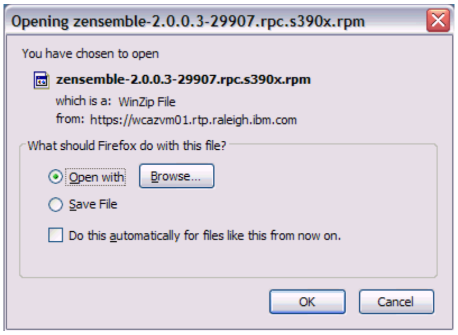

Opening zensemble-2.0.0.3-29907.rpc.s390x.rpm

You have chosen to open

zensemble-2.0.0.3-29907.rpc.s390x.rpm

which is a: WinZip File

from: https://wcazvm01.rtp.raleigh.ibm.com

What should Firefox do with this file?

● Open with

○ Save File

□ Do this automatically for files like this from now on.

- Install RPM

```
rpm –ivh zensemble-2.0.0.3-29907.rpc.s390x.rpm
```

  – Delete RPM

```
      rpm -e zensemble-2.0.0.2-26833
```

Clicking on the hyperlink will prompt you to either open the RPM or save it. Save the RPM file to your workstation, upload it to your MAPSRV ID and install it. The hypervisor will ensure that you are using the RPM that matches the firmware on the appliance. To install the RPM use the ‘rpm’ command shown on the slide. Note that you need to delete any old ones before installing a newer one.

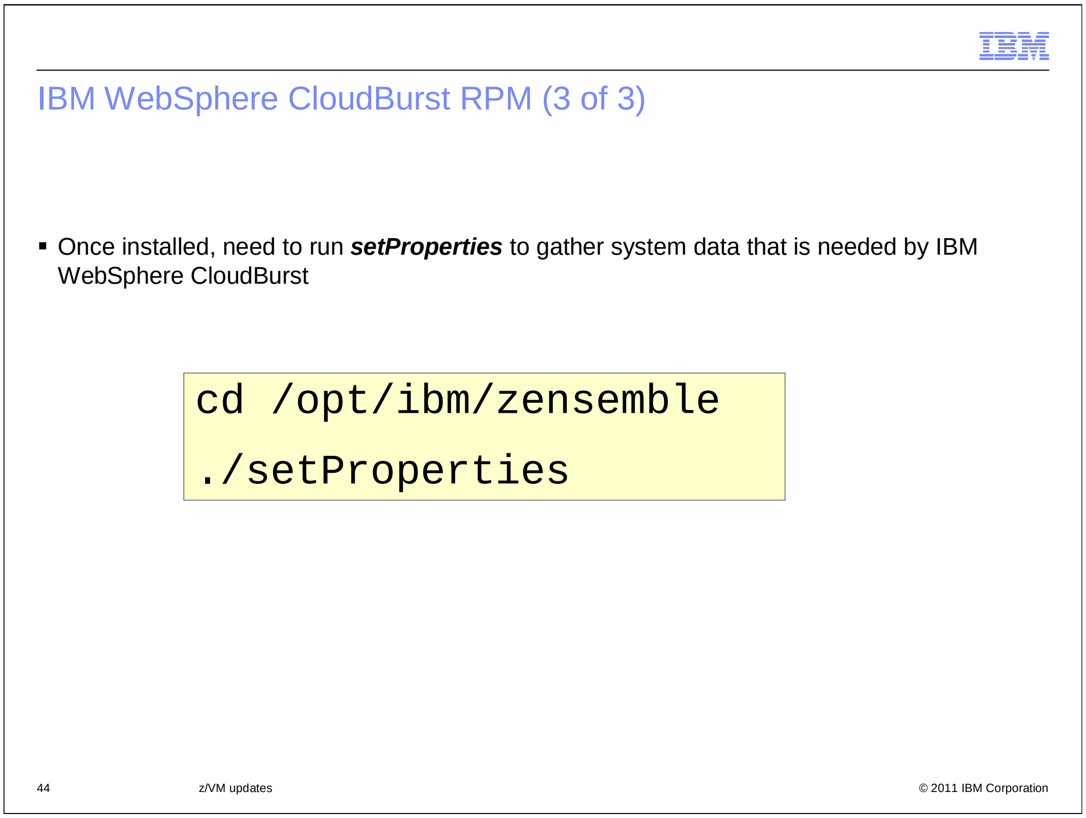

IBM WebSphere CloudBurst RPM (3 of 3)

- Once installed, need to run ***setProperties*** to gather system data that is needed by IBM WebSphere CloudBurst

```
cd /opt/ibm/zensemble
./setProperties
```

Once installed, you will notice one of the files that is installed is the setProperties script that needs to be run in order to gather system data for the IBM WebSphere CloudBurst appliance. The next few slides will show the information that you are prompted for when running the setProperties command.

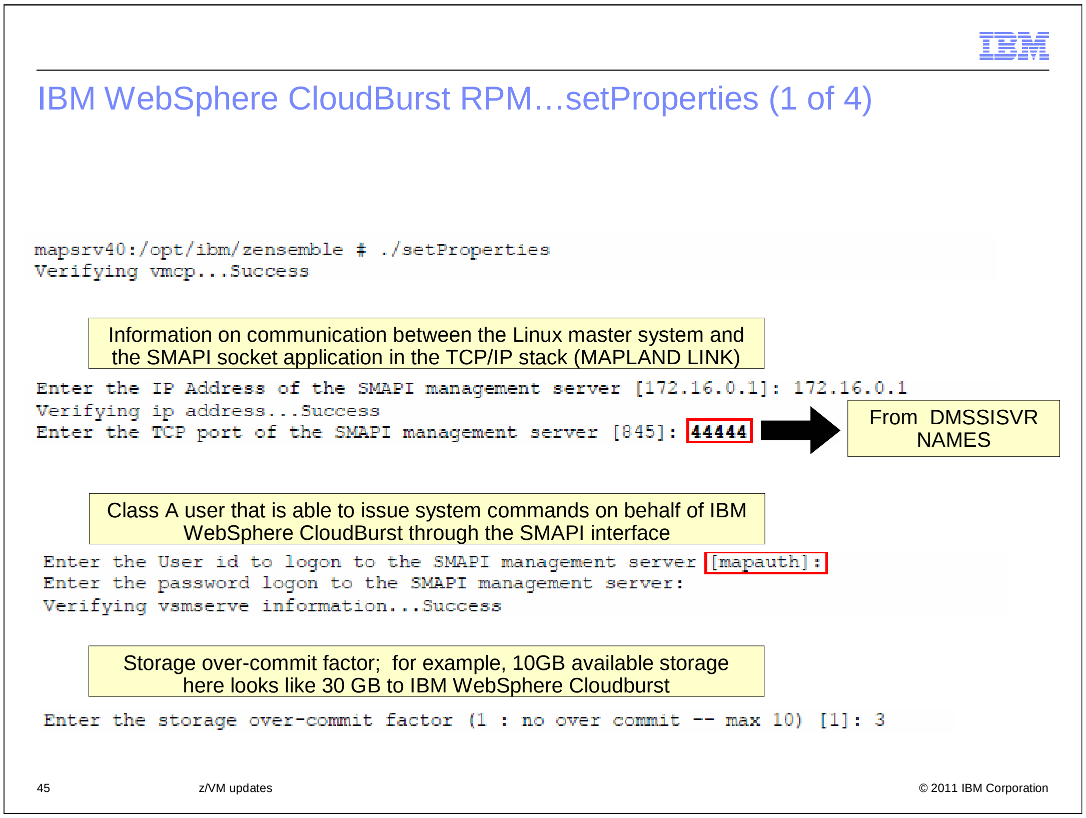

IBM WebSphere CloudBurst RPM…setProperties (1 of 4)

```
mapsrv40:/opt/ibm/zensemble # ./setProperties
Verifying vmcp...Success
```

Information on communication between the Linux master system and the SMAPI socket application in the TCP/IP stack (MAPLAND LINK)

```
Enter the IP Address of the SMAPI management server [172.16.0.1]: 172.16.0.1
Verifying ip address...Success
Enter the TCP port of the SMAPI management server [845]: 44444
```

From DMSSISVR NAMES

Class A user that is able to issue system commands on behalf of IBM WebSphere CloudBurst through the SMAPI interface

```
Enter the User id to logon to the SMAPI management server [mapauth]:
Enter the password logon to the SMAPI management server:
Verifying vsmserve information...Success
```

Storage over-commit factor; for example, 10GB available storage here looks like 30 GB to IBM WebSphere Cloudburst

```
Enter the storage over-commit factor (1 : no over commit -- max 10) [1]: 3
```

At the end of a successful installation of the zensemble rpm file, a reminder is posted to run /opt/ibm/zensemble/setProperties script to gather system data/information for use by IBM WebSphere CloudBurst. This script will gather z/VM system information that is used by IBM WebSphere CloudBurst for deployment of Linux VMs. It will query or discover various pieces of information. The first piece of information needed by setProperties is the IP address and port of the SMAPI management server. The port information comes from the port you specified in the DMSSISVR NAMES file. The next piece of information needed is the name of the Class A user that you defined for use by IBM WebSphere CloudBurst. The default is MAPAUTH. It is the user ID that is used to issue the SMAPI commands used to provision guests from IBM WebSphere CloudBurst. You will also need to know MAPAUTH's logon password. For the over-commit factor, if physical memory is 8 GB, and over-commit ratio is 2, IBM WebSphere CloudBurst Appliance will assume 16 GB is available for VM deployment. The IBM WebSphere CloudBurst Appliance monitors the VM Working Set size and estimates usage of memory for placement.

IBM WebSphere CloudBurst RPM…setProperties (2 of 4)

Information on VSWITCH to be used for external communications

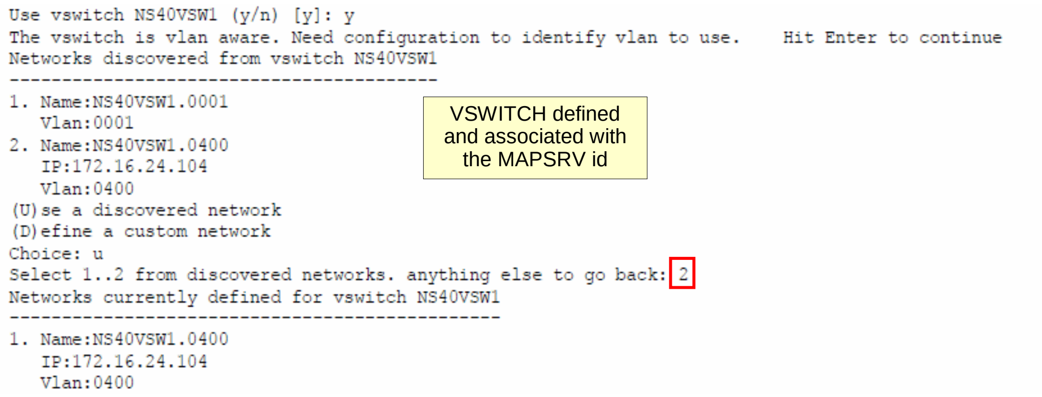

```
Use vswitch NS40VSW1 (y/n) [y]: y
The vswitch is vlan aware. Need configuration to identify vlan to use.    Hit Enter to continue
Networks discovered from vswitch NS40VSW1
-----------------------------------------
1. Name:NS40VSW1.0001
   Vlan:0001
2. Name:NS40VSW1.0400
   IP:172.16.24.104
   Vlan:0400
(U)se a discovered network
(D)efine a custom network
Choice: u
Select 1..2 from discovered networks. anything else to go back: 2
Networks currently defined for vswitch NS40VSW1
-----------------------------------------------
1. Name:NS40VSW1.0400
   IP:172.16.24.104
   Vlan:0400
```

VSWITCH defined and associated with the MAPSRV id

If you defined a VSWITCH for use by the MAPSRV user ID, the VSWITCH ID can be discovered and then selected by number. This is shown here on the slide.

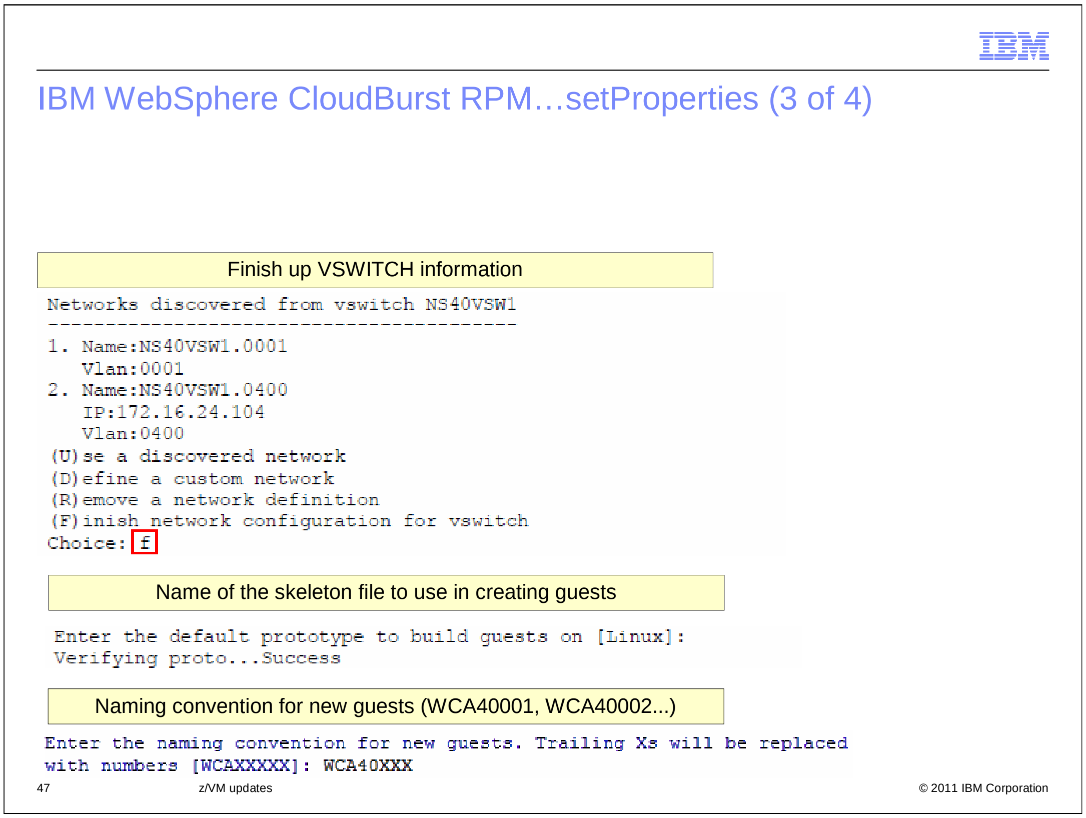

IBM WebSphere CloudBurst RPM…setProperties (3 of 4)

Finish up VSWITCH information

```
Networks discovered from vswitch NS40VSW1
-----------------------------------------
1. Name:NS40VSW1.0001
   Vlan:0001
2. Name:NS40VSW1.0400
   IP:172.16.24.104
   Vlan:0400
(U)se a discovered network
(D)efine a custom network
(R)emove a network definition
(F)inish network configuration for vswitch
Choice: f
```

Name of the skeleton file to use in creating guests

```
Enter the default prototype to build guests on [Linux]:
Verifying proto...Success
```

Naming convention for new guests (WCA40001, WCA40002...)

```
Enter the naming convention for new guests. Trailing Xs will be replaced
with numbers [WCAXXXXX]: WCA40XXX
```

47 z/VM updates © 2011 IBM Corporation

After finishing with the VSWITCH configuration information, you are asked to provide a default prototype or skeleton file to use to build the Linux guests that get provisioned. This prototype is used if the **prototypes.conf** file does not exist or a matching image part type is not found there. You are also asked to provide a naming convention for new guests. Any X’s are replaced with numbers. In the example shown here, the first guest is defined as WCA40000. The second one is then defined as WCA40001 and so on.

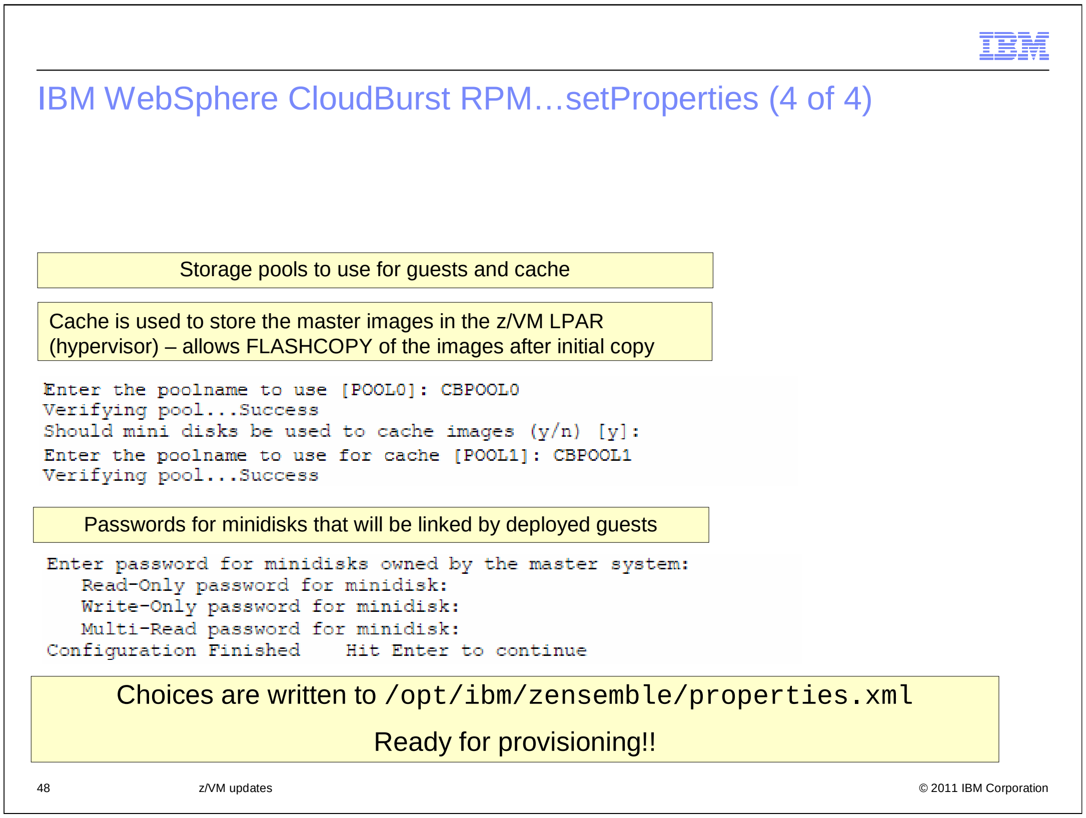

IBM WebSphere CloudBurst RPM…setProperties (4 of 4)

Storage pools to use for guests and cache

Cache is used to store the master images in the z/VM LPAR (hypervisor) – allows FLASHCOPY of the images after initial copy

```
Enter the poolname to use [POOL0]: CBPOOL0
Verifying pool...Success
Should mini disks be used to cache images (y/n) [y]:
Enter the poolname to use for cache [POOL1]: CBPOOL1
Verifying pool...Success
```

Passwords for minidisks that will be linked by deployed guests

```
Enter password for minidisks owned by the master system:
   Read-Only password for minidisk:
   Write-Only password for minidisk:
   Multi-Read password for minidisk:
Configuration Finished    Hit Enter to continue
```

Choices are written to `/opt/ibm/zensemble/properties.xml`

Ready for provisioning!!

48 z/VM updates © 2011 IBM Corporation

You have seen the configuration of two possible storage pools. You need to let IBM WebSphere CloudBurst know about the storage pools defined so that they can be used for the provisioning of guests. The first poolname specified here is the one that is required and is used for the guest disks. The second one is optional and is used to cache the images. You can use the same pool for both but in the example, CBPOOL0 and CBPOOL1 are being used.

The last piece of information you must supply when running the setProperties script are the passwords for the minidisks that are linked by the deployed guests. When you are finished, the values you specified are written to /opt/ibm/zensemble/properties.xml.

Section

Troubleshooting

Finally, you will briefly look at where to find information if things go wrong.

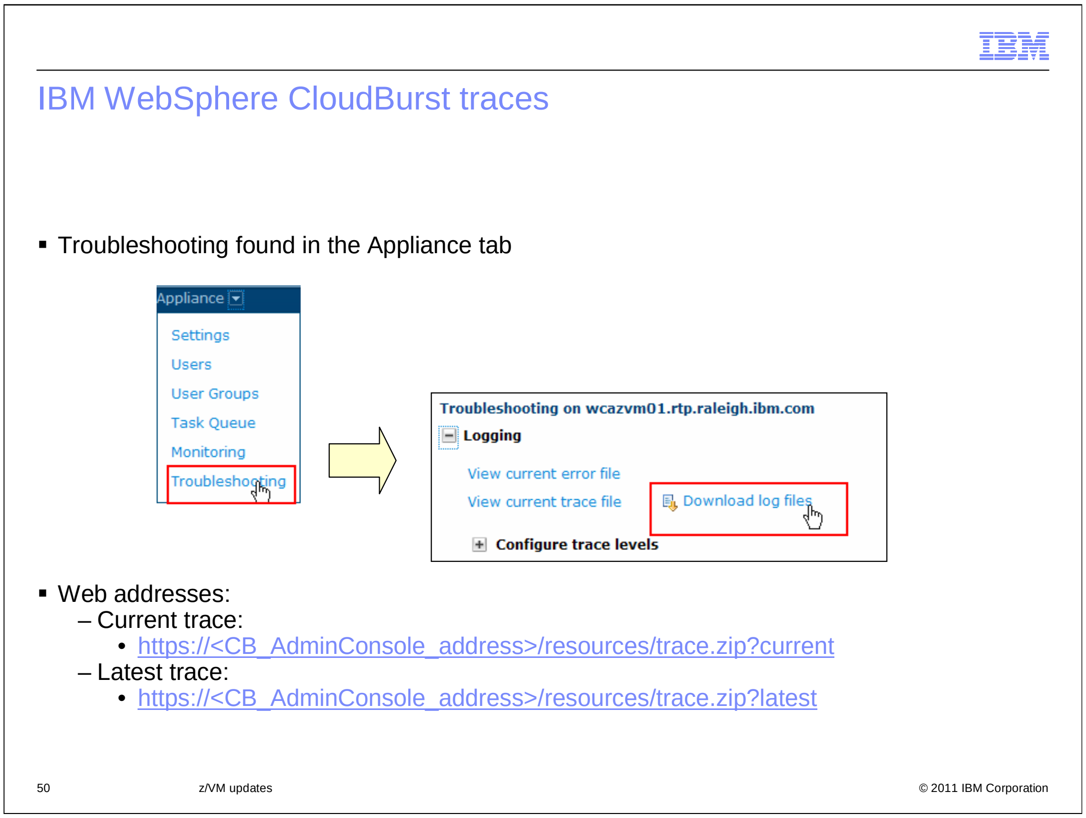

IBM WebSphere CloudBurst traces

- Troubleshooting found in the Appliance tab

- Web addresses:
  - Current trace:
    - https://<CB_AdminConsole_address>/resources/trace.zip?current
  - Latest trace:
    - https://<CB_AdminConsole_address>/resources/trace.zip?latest

50 z/VM updates

There are logs available for you to look at or supply to the service team from the IBM WebSphere CloudBurst appliance under the Appliance tab as shown on the slide. This log can be quite large with lots of information. It is also available using the URL shown on the slide under ‘Current trace’. The current trace contains multiple trace files that are maintained on a rolling basis. To get just the latest logs, and thus a much smaller file, use the URL shown under ‘Latest trace’ instead.

Sample current trace, partial listing

https://<CB_AdminConsole_address>/resources/trace.zip?current

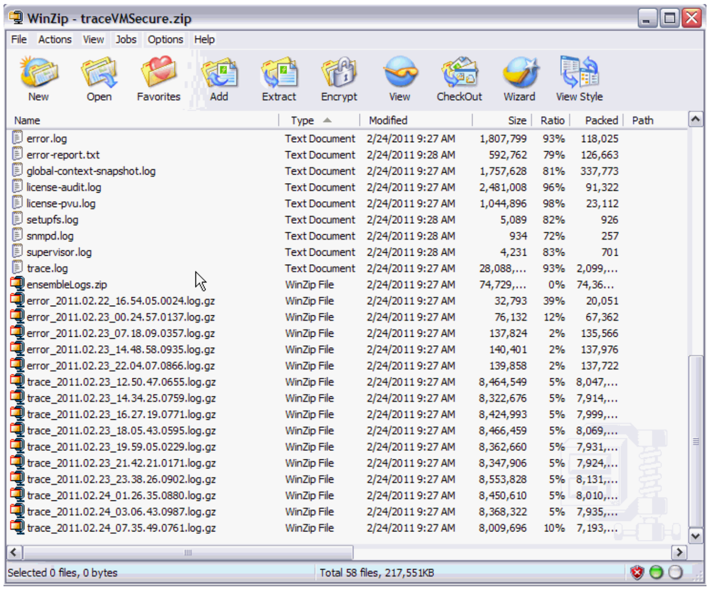

Name	Type	Modified	Size	Ratio	Packed	Path
error.log	Text Document	2/24/2011 9:27 AM	1,807,799	93%	118,025	
error-report.txt	Text Document	2/24/2011 9:28 AM	592,762	79%	126,663	
global-context-snapshot.log	Text Document	2/24/2011 9:27 AM	1,757,628	81%	337,773	
license-audit.log	Text Document	2/24/2011 9:27 AM	2,481,008	96%	91,322	
license-pvu.log	Text Document	2/24/2011 9:27 AM	1,044,896	98%	23,112	
setupfs.log	Text Document	2/24/2011 9:28 AM	5,089	82%	926	
snmpd.log	Text Document	2/24/2011 9:28 AM	934	72%	257	
supervisor.log	Text Document	2/24/2011 9:28 AM	4,231	83%	701	
trace.log	Text Document	2/24/2011 9:27 AM	28,088,...	93%	2,099,...	
ensembleLogs.zip	WinZip File	2/24/2011 9:27 AM	74,729,...	0%	74,36...	
error_2011.02.22_16.54.05.0024.log.gz	WinZip File	2/24/2011 9:27 AM	32,793	39%	20,051	
error_2011.02.23_00.24.57.0137.log.gz	WinZip File	2/24/2011 9:27 AM	76,132	12%	67,362	
error_2011.02.23_07.18.09.0357.log.gz	WinZip File	2/24/2011 9:27 AM	137,824	2%	135,566	
error_2011.02.23_14.48.58.0935.log.gz	WinZip File	2/24/2011 9:27 AM	140,401	2%	137,976	
error_2011.02.23_22.04.07.0866.log.gz	WinZip File	2/24/2011 9:27 AM	139,858	2%	137,722	
trace_2011.02.23_12.50.47.0655.log.gz	WinZip File	2/24/2011 9:27 AM	8,464,549	5%	8,047,...	
trace_2011.02.23_14.34.25.0759.log.gz	WinZip File	2/24/2011 9:27 AM	8,322,676	5%	7,914,...	
trace_2011.02.23_16.27.19.0771.log.gz	WinZip File	2/24/2011 9:27 AM	8,424,993	5%	7,999,...	
trace_2011.02.23_18.05.43.0595.log.gz	WinZip File	2/24/2011 9:27 AM	8,466,459	5%	8,069,...	
trace_2011.02.23_19.59.05.0229.log.gz	WinZip File	2/24/2011 9:27 AM	8,362,660	5%	7,931,...	
trace_2011.02.23_21.42.21.0171.log.gz	WinZip File	2/24/2011 9:27 AM	8,347,906	5%	7,924,...	
trace_2011.02.23_23.38.26.0902.log.gz	WinZip File	2/24/2011 9:27 AM	8,553,828	5%	8,131,...	
trace_2011.02.24_01.26.35.0880.log.gz	WinZip File	2/24/2011 9:27 AM	8,450,610	5%	8,010,...	
trace_2011.02.24_03.06.43.0987.log.gz	WinZip File	2/24/2011 9:27 AM	8,368,322	5%	7,935,...	
trace_2011.02.24_07.35.49.0761.log.gz	WinZip File	2/24/2011 9:27 AM	8,009,696	10%	7,193,...	

Selected 0 files, 0 bytes

Total 58 files, 217,551KB

This slide shows the contents of a trace file that has been downloaded from the administrative console. This is also what you see when the you ask for the 'current trace'. The bottom of the listing is shown here so you can see the list of 'older' logs that have been zipped up. The 'latest' log includes just the 'error.log' and trace.log' and does not include all the older, zipped up ones, making it much smaller. You should also notice here the ensembles.log. This has the z/VM-specific information in it. You will see that in more detail on the next slide.

Sample current trace, ensembleLogs.zip

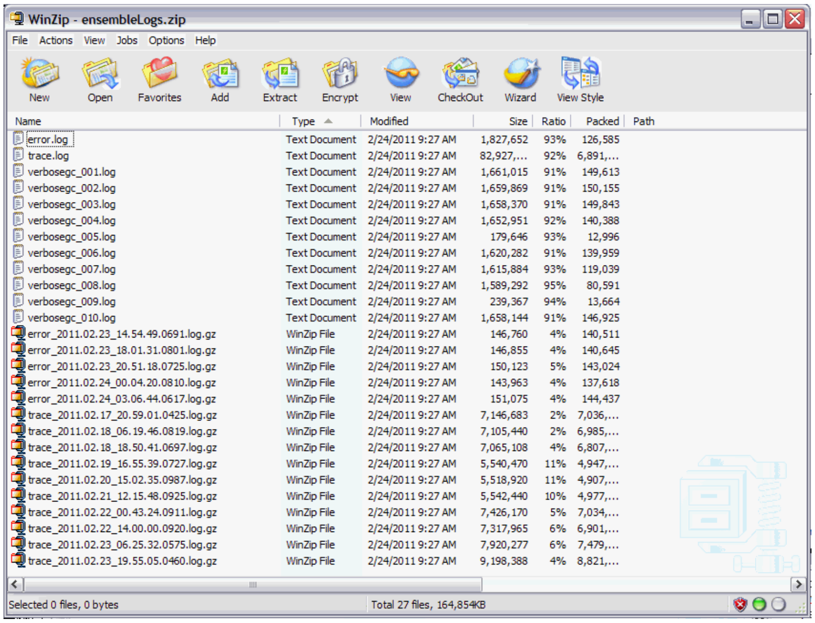

Name	Type	Modified	Size	Ratio	Packed	Path
error.log	Text Document	2/24/2011 9:27 AM	1,827,652	93%	126,585	
trace.log	Text Document	2/24/2011 9:27 AM	82,927,...	92%	6,891,...	
verbosegc_001.log	Text Document	2/24/2011 9:27 AM	1,661,015	91%	149,613	
verbosegc_002.log	Text Document	2/24/2011 9:27 AM	1,659,869	91%	150,155	
verbosegc_003.log	Text Document	2/24/2011 9:27 AM	1,658,370	91%	149,843	
verbosegc_004.log	Text Document	2/24/2011 9:27 AM	1,652,951	92%	140,388	
verbosegc_005.log	Text Document	2/24/2011 9:27 AM	179,646	93%	12,996	
verbosegc_006.log	Text Document	2/24/2011 9:27 AM	1,620,282	91%	139,959	
verbosegc_007.log	Text Document	2/24/2011 9:27 AM	1,615,884	93%	119,039	
verbosegc_008.log	Text Document	2/24/2011 9:27 AM	1,589,292	95%	80,591	
verbosegc_009.log	Text Document	2/24/2011 9:27 AM	239,367	94%	13,664	
verbosegc_010.log	Text Document	2/24/2011 9:27 AM	1,658,144	91%	146,925	
error_2011.02.23_14.54.49.0691.log.gz	WinZip File	2/24/2011 9:27 AM	146,760	4%	140,511	
error_2011.02.23_18.01.31.0801.log.gz	WinZip File	2/24/2011 9:27 AM	146,855	4%	140,645	
error_2011.02.23_20.51.18.0725.log.gz	WinZip File	2/24/2011 9:27 AM	150,123	5%	143,024	
error_2011.02.24_00.04.20.0810.log.gz	WinZip File	2/24/2011 9:27 AM	143,963	4%	137,618	
error_2011.02.24_03.06.44.0617.log.gz	WinZip File	2/24/2011 9:27 AM	151,075	4%	144,437	
trace_2011.02.17_20.59.01.0425.log.gz	WinZip File	2/24/2011 9:27 AM	7,146,683	2%	7,036,...	
trace_2011.02.18_06.19.46.0819.log.gz	WinZip File	2/24/2011 9:27 AM	7,105,440	2%	6,985,...	
trace_2011.02.18_18.50.41.0697.log.gz	WinZip File	2/24/2011 9:27 AM	7,065,108	4%	6,807,...	
trace_2011.02.19_16.55.39.0727.log.gz	WinZip File	2/24/2011 9:27 AM	5,540,470	11%	4,947,...	
trace_2011.02.20_15.02.35.0987.log.gz	WinZip File	2/24/2011 9:27 AM	5,518,920	11%	4,907,...	
trace_2011.02.21_12.15.48.0925.log.gz	WinZip File	2/24/2011 9:27 AM	5,542,440	10%	4,977,...	
trace_2011.02.22_00.43.24.0911.log.gz	WinZip File	2/24/2011 9:27 AM	7,426,170	5%	7,034,...	
trace_2011.02.22_14.00.00.0920.log.gz	WinZip File	2/24/2011 9:27 AM	7,317,965	6%	6,901,...	
trace_2011.02.23_06.25.32.0575.log.gz	WinZip File	2/24/2011 9:27 AM	7,920,277	6%	7,479,...	
trace_2011.02.23_19.55.05.0460.log.gz	WinZip File	2/24/2011 9:27 AM	9,198,388	4%	8,821,...	

Selected 0 files, 0 bytes

Total 27 files, 164,854KB

Again, this is the log where you can look for z/VM-specific information. Note there is another error.log and trace.log here and they also wrap on a regular basis and are also zipped up. The WinZip files are not be included in the ?latest traces.

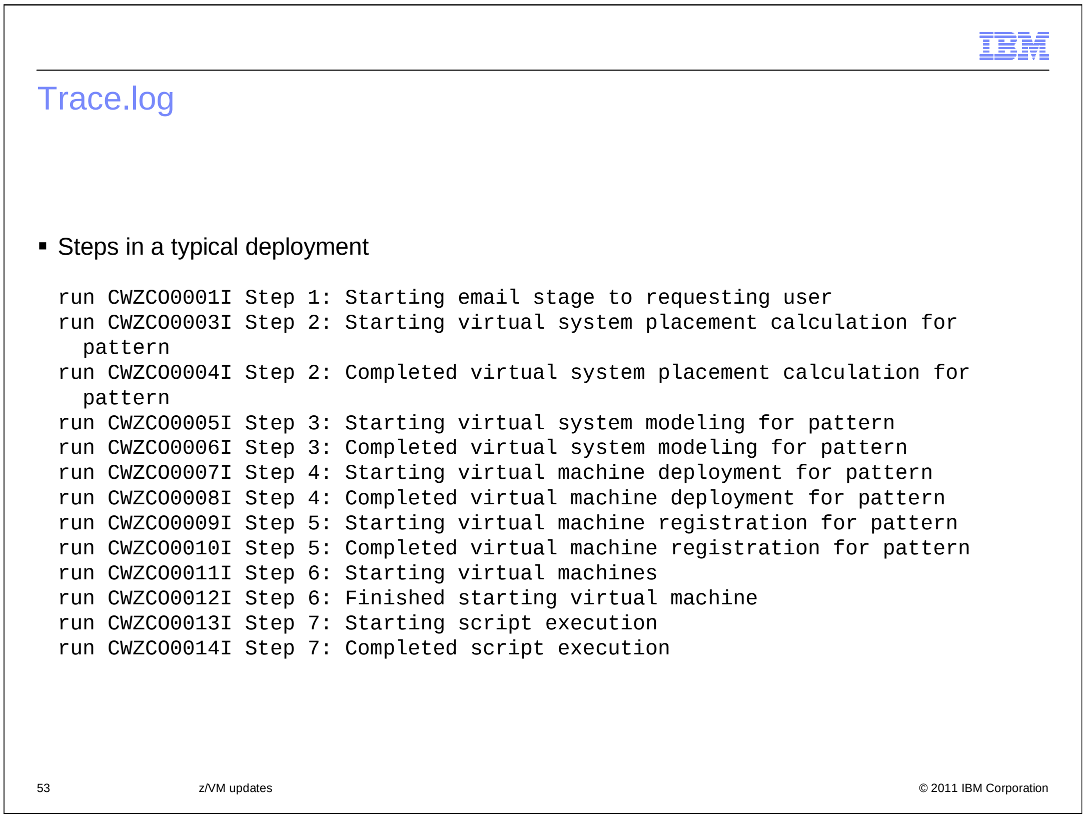

Trace.log

- Steps in a typical deployment

```
run CWZCO0001I Step 1: Starting email stage to requesting user
run CWZCO0003I Step 2: Starting virtual system placement calculation for
  pattern
run CWZCO0004I Step 2: Completed virtual system placement calculation for
  pattern
run CWZCO0005I Step 3: Starting virtual system modeling for pattern
run CWZCO0006I Step 3: Completed virtual system modeling for pattern
run CWZCO0007I Step 4: Starting virtual machine deployment for pattern
run CWZCO0008I Step 4: Completed virtual machine deployment for pattern
run CWZCO0009I Step 5: Starting virtual machine registration for pattern
run CWZCO0010I Step 5: Completed virtual machine registration for pattern
run CWZCO0011I Step 6: Starting virtual machines
run CWZCO0012I Step 6: Finished starting virtual machine
run CWZCO0013I Step 7: Starting script execution
run CWZCO0014I Step 7: Completed script execution
```

This shows some records from the non-ensemble trace.log. It gives you an idea of the steps that you should expect in a typical deployment. You can find these records by searching on ‘Step’ in the log to see where the deployment might have failed.

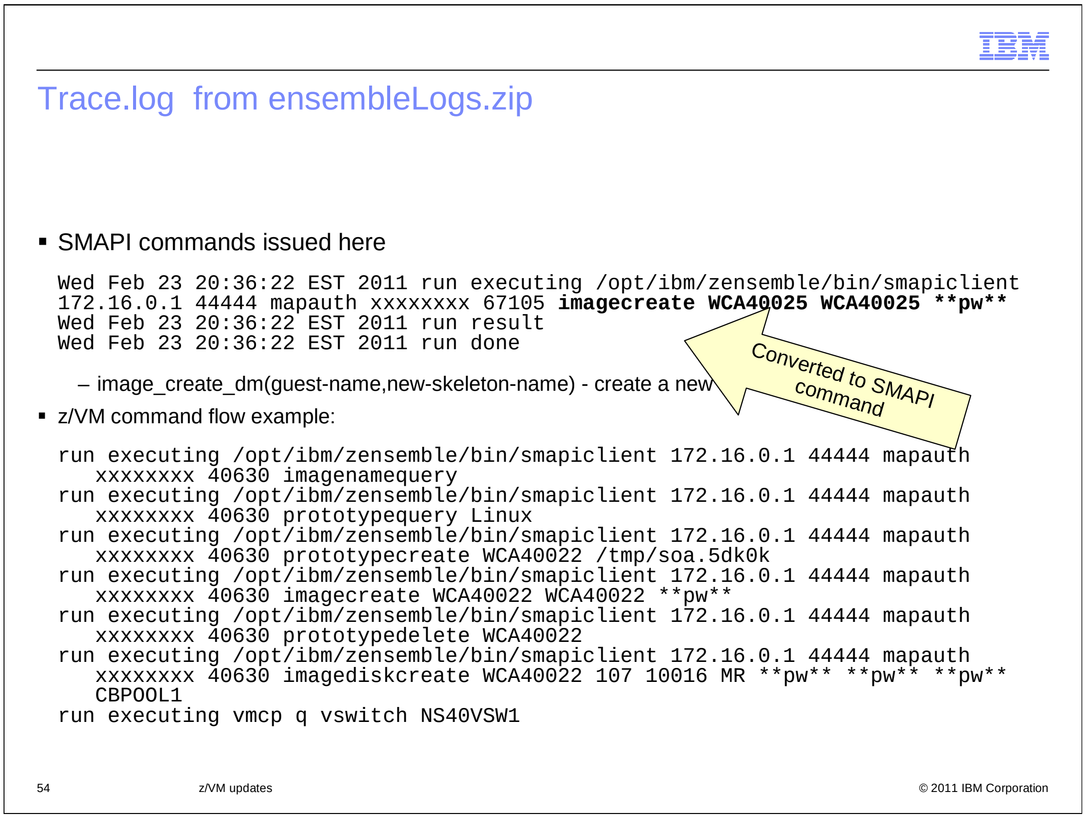

Trace.log from ensembleLogs.zip

- SMAPI commands issued here

```
Wed Feb 23 20:36:22 EST 2011 run executing /opt/ibm/zensemble/bin/smapiclient
172.16.0.1 44444 mapauth xxxxxxxx 67105 imagecreate WCA40025 WCA40025 **pw**
Wed Feb 23 20:36:22 EST 2011 run result
Wed Feb 23 20:36:22 EST 2011 run done
```

  - image_create_dm(guest-name,new-skeleton-name) - create a new

Converted to SMAPI command

- z/VM command flow example:

```
run executing /opt/ibm/zensemble/bin/smapiclient 172.16.0.1 44444 mapauth
   xxxxxxxx 40630 imagenamequery
run executing /opt/ibm/zensemble/bin/smapiclient 172.16.0.1 44444 mapauth
   xxxxxxxx 40630 prototypequery Linux
run executing /opt/ibm/zensemble/bin/smapiclient 172.16.0.1 44444 mapauth
   xxxxxxxx 40630 prototypecreate WCA40022 /tmp/soa.5dk0k
run executing /opt/ibm/zensemble/bin/smapiclient 172.16.0.1 44444 mapauth
   xxxxxxxx 40630 imagecreate WCA40022 WCA40022 **pw**
run executing /opt/ibm/zensemble/bin/smapiclient 172.16.0.1 44444 mapauth
   xxxxxxxx 40630 prototypedelete WCA40022
run executing /opt/ibm/zensemble/bin/smapiclient 172.16.0.1 44444 mapauth
   xxxxxxxx 40630 imagediskcreate WCA40022 107 10016 MR **pw** **pw** **pw**
   CBPOOL1
run executing vmcp q vswitch NS40VSW1
```

54 z/VM updates © 2011 IBM Corporation

This shows some records from the ensemble trace.log. This slide shows an imagecreate execution with no error. It translates to the image_create_dm SMAPI command. To give you an idea of the z/VM flow, you can do a search on the phrase ‘run executing’..

SMAPI call documentation

- Calls are documented here:
  - http://publib.boulder.ibm.com/infocenter/zvm/v5r4/topic/com.ibm.zvm.v54.dmse6/sok.htm#sok

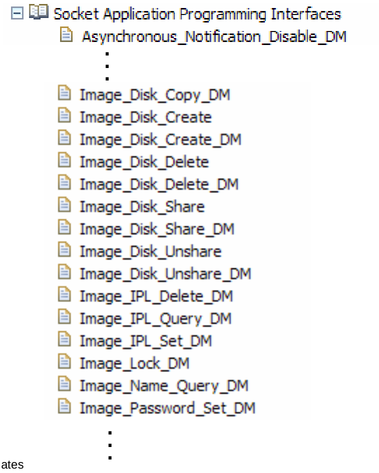

- Socket Application Programming Interfaces
  - Asynchronous_Notification_Disable_DM
  - ⋮
  - Image_Disk_Copy_DM
  - Image_Disk_Create
  - Image_Disk_Create_DM
  - Image_Disk_Delete
  - Image_Disk_Delete_DM
  - Image_Disk_Share
  - Image_Disk_Share_DM
  - Image_Disk_Unshare
  - Image_Disk_Unshare_DM
  - Image_IPL_Delete_DM
  - Image_IPL_Query_DM
  - Image_IPL_Set_DM
  - Image_Lock_DM
  - Image_Name_Query_DM
  - Image_Password_Set_DM
  - ⋮

The SMAPI calls are documented at the URL shown on the slide. For most errors, you can use the URL shown on the next slide.

Return codes documented

- All return codes for SMAPI calls
  - – http://publib.boulder.ibm.com/infocenter/zvm/v6r1/topic/com.ibm.zvm.v610.dmse6/hcsl8c0153.htm#wq1271

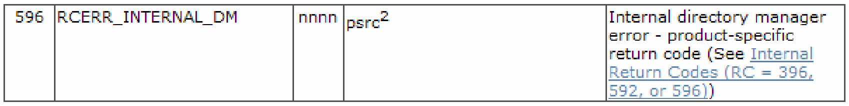

596	RCERR_INTERNAL_DM	nnnn	psrc[2]	Internal directory manager error - product-specific return code (See Internal Return Codes (RC = 396, 592, or 596))

  - – With VM:Secure, the 596 return code says that VM:Secure returned an error
    - • Will see a return code in the 5XXXXX - 9XXXXX range; ERROR maps to a VM:Secure function
      - – For example:

```
Failed. rc = 596, rs = 1399899
*ERROR*:1399899

rc=64512
```

The URL shown here is a collection of all return codes that are returned from the SMAPI calls. The slide is showing the 596 error code. The error returned maps to a product-specifc VM:Secure error.

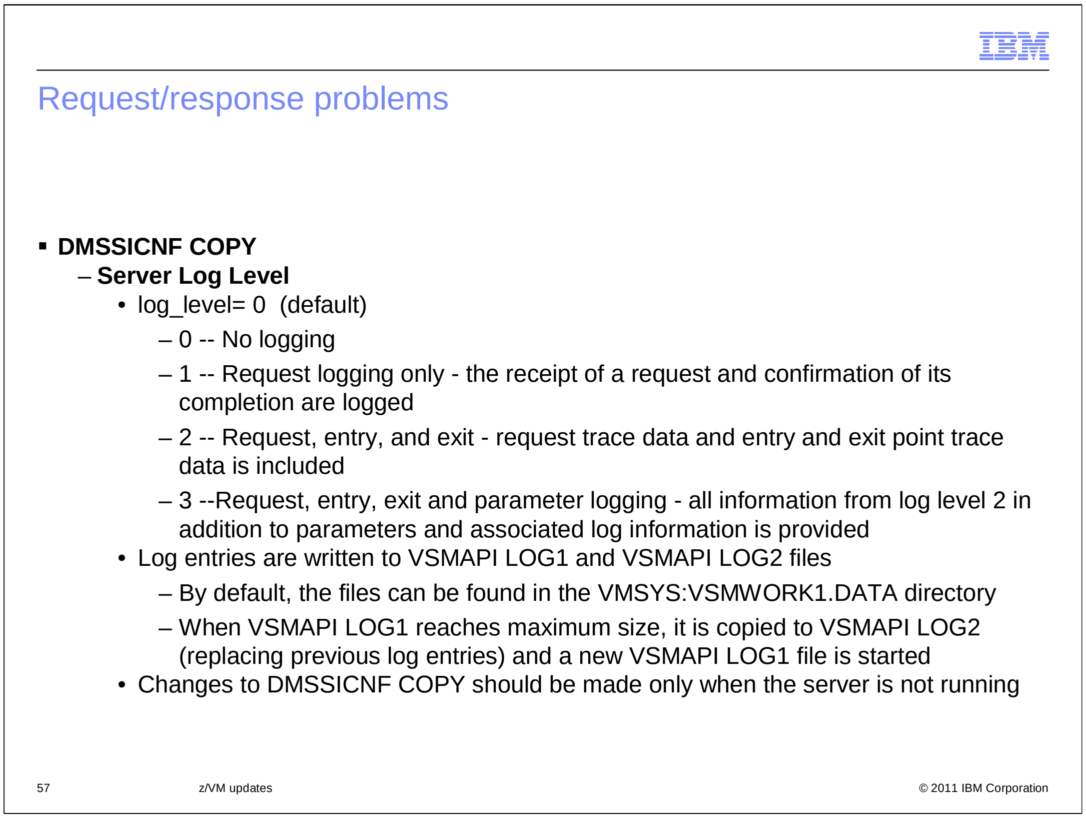

Request/response problems

- **DMSSICNF COPY**
  - **Server Log Level**
    - log_level= 0 (default)
      - 0 -- No logging
      - 1 -- Request logging only - the receipt of a request and confirmation of its completion are logged
      - 2 -- Request, entry, and exit - request trace data and entry and exit point trace data is included
      - 3 --Request, entry, exit and parameter logging - all information from log level 2 in addition to parameters and associated log information is provided
    - Log entries are written to VSMAPI LOG1 and VSMAPI LOG2 files
      - By default, the files can be found in the VMSYS:VSMWORK1.DATA directory
      - When VSMAPI LOG1 reaches maximum size, it is copied to VSMAPI LOG2 (replacing previous log entries) and a new VSMAPI LOG1 file is started
    - Changes to DMSSICNF COPY should be made only when the server is not running

If you are experiencing problems with the server environment, you can turn on logging in the DMSSICNF COPY file. The log levels are listed on the slide. The log entries are written to the VSMAPI LOG1 and the VSMAPI LOG2 files. When LOG1 reaches its maximum size, it is copied to LOG2 and a new LOG1 file is started. Be default, these files are found in the VMSYS:VSMWORK1.DATA directory.

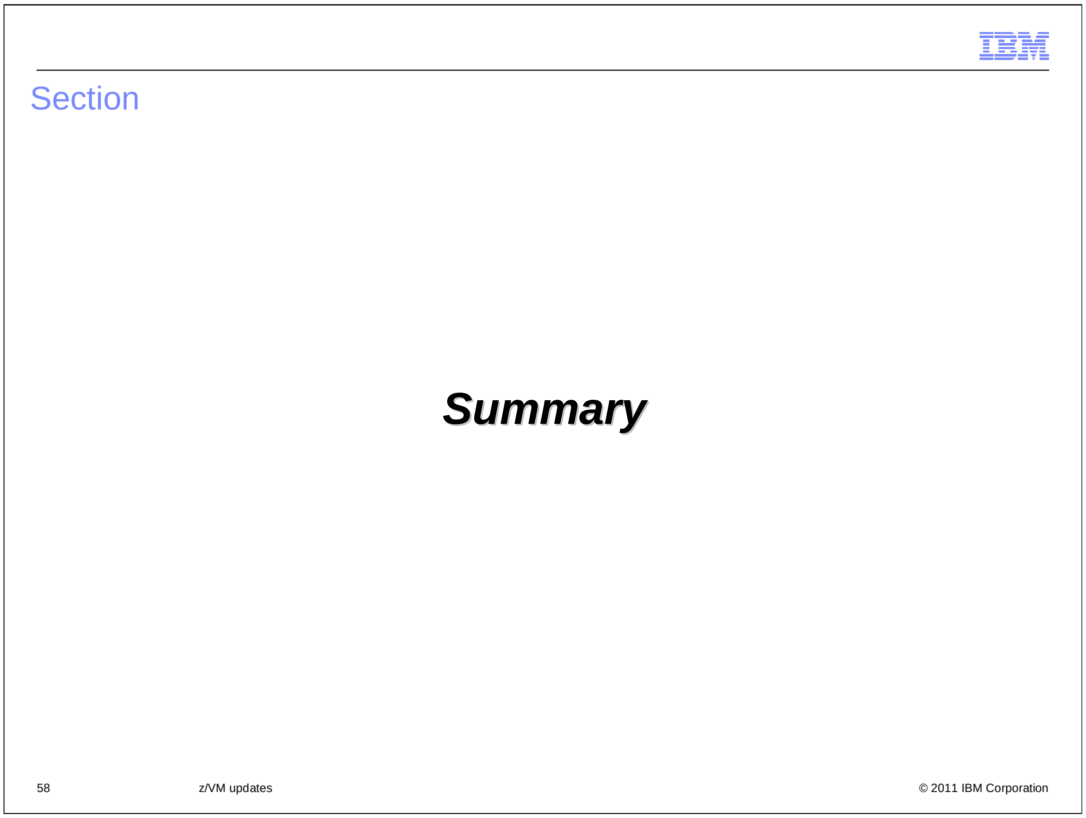

Section

Summary

58 z/VM updates © 2011 IBM Corporation

This section summarizes the IBM WebSphere CloudBurst V2.0.0.3 z/VM updates.

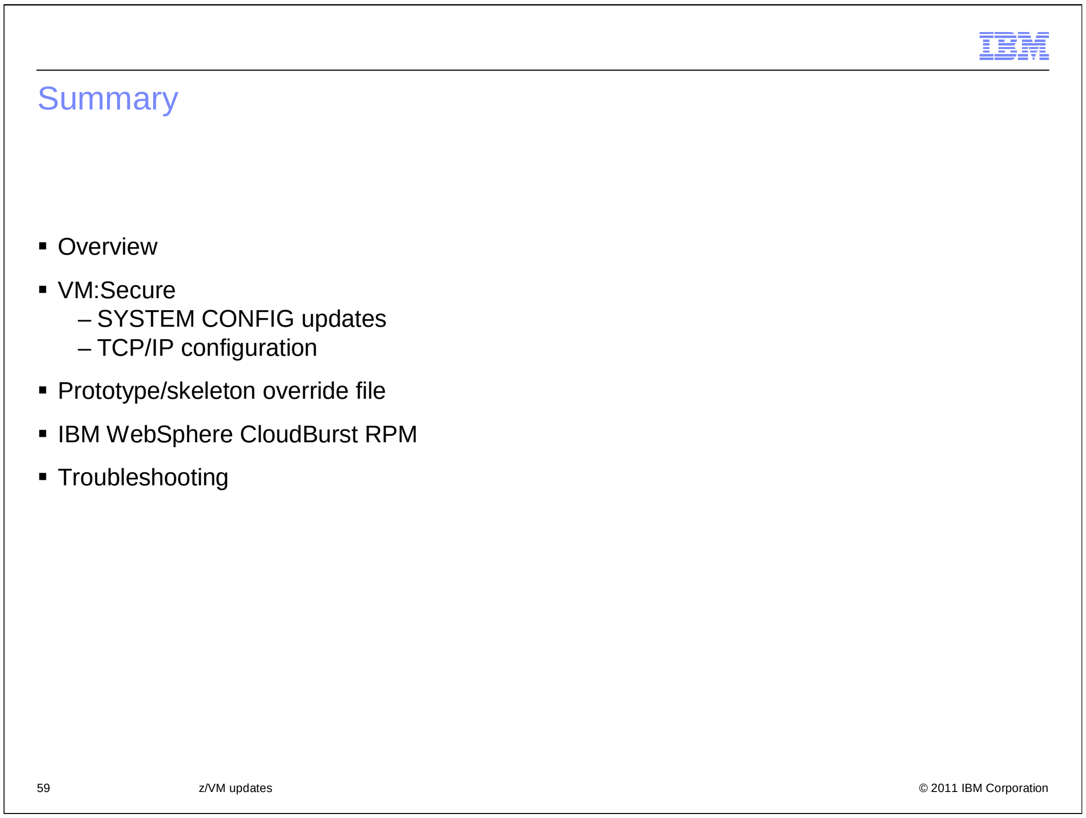

Summary

- Overview
- VM:Secure
  - SYSTEM CONFIG updates
  - TCP/IP configuration
- Prototype/skeleton override file
- IBM WebSphere CloudBurst RPM
- Troubleshooting

59 z/VM updates © 2011 IBM Corporation

In summary, this presentation looked at the configuration needed on the z/VM system in order for IBM WebSphere CloudBurst to successfully provision guests. New with V2.0.0.3, you can use VM:Secure as your directory manager. This presentation focused on the configuration needed to use VM:Secure in the IBM WebSphere CloudBurst environment. You also saw the new prototype override file that was introduced in V2.0.0.3. Finally, you ended up looking at some basic troubleshooting information.

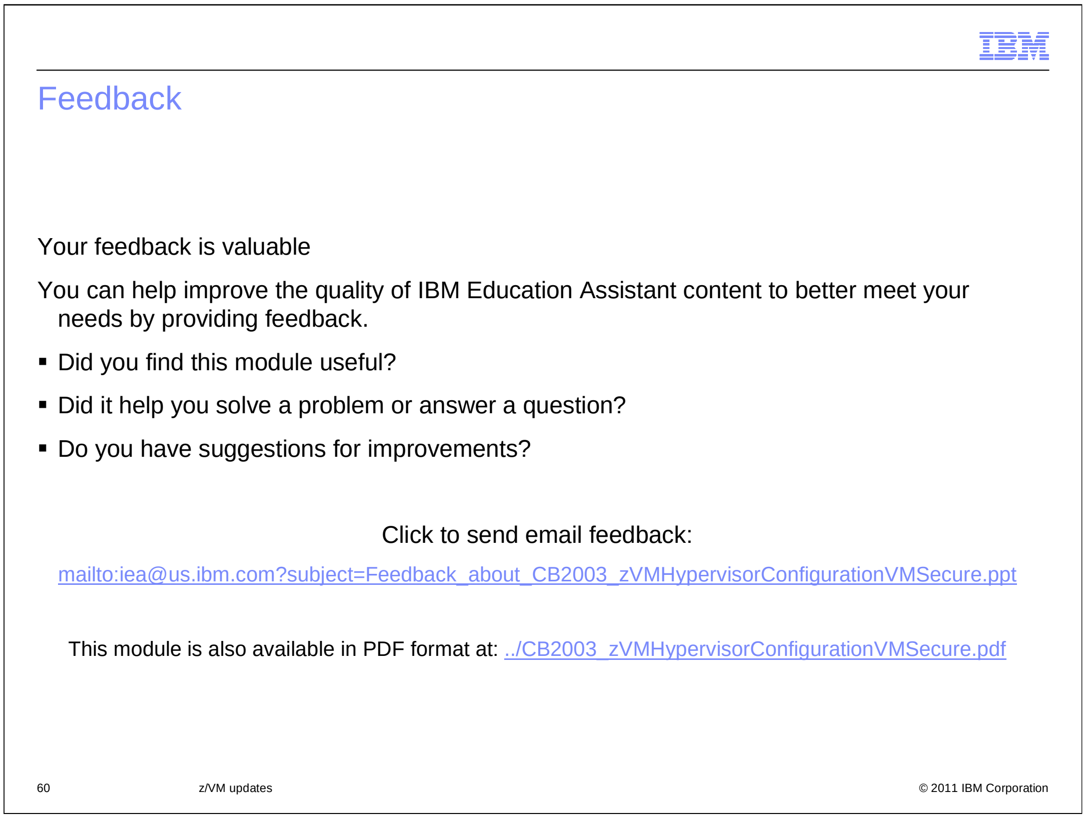

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_CB2003_zVMHypervisorConfigurationVMSecure.ppt

This module is also available in PDF format at: ../CB2003_zVMHypervisorConfigurationVMSecure.pdf

You can help improve the quality of IBM Education Assistant content by providing feedback.

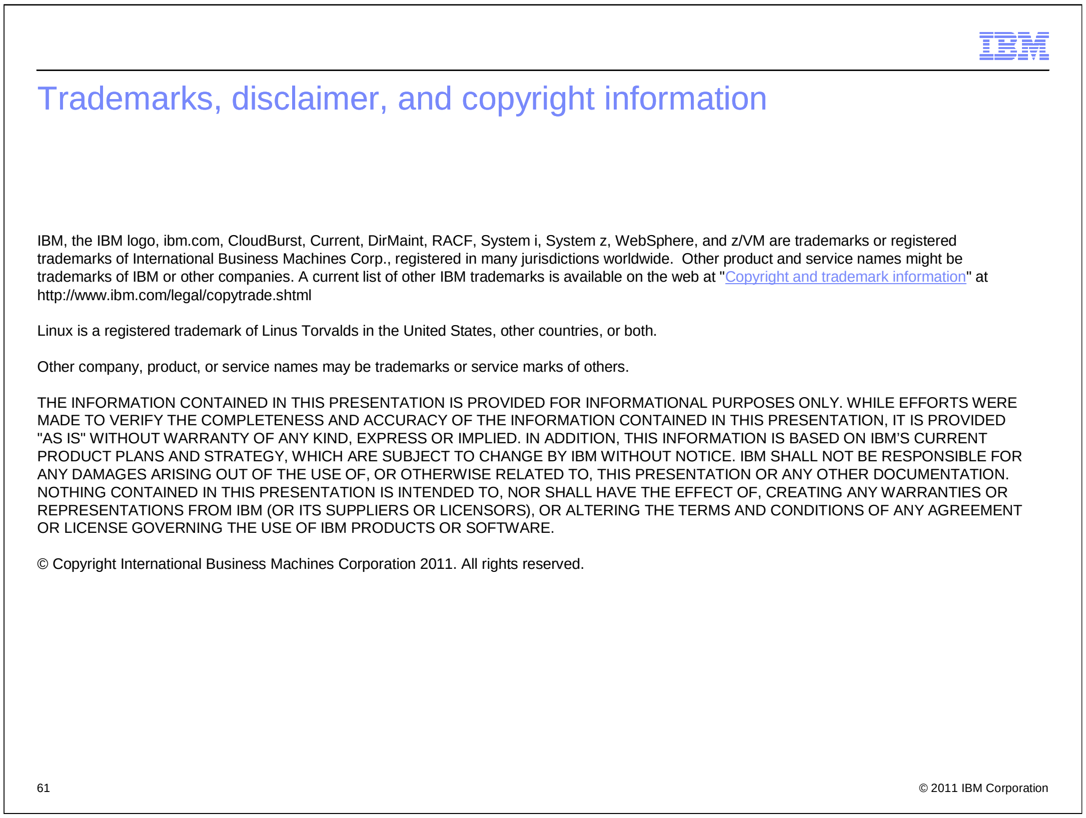

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, CloudBurst, Current, DirMaint, RACF, System i, System z, WebSphere, and z/VM are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.