
IBM WebSphere CloudBurst Appliance

z/VM hypervisor setup



This presentation covers the steps needed to configure the z/VM hypervisor for use by IBM WebSphere® CloudBurst™.

Agenda

- Overview
- z/VM setup
 - SYSTEM CONFIG updates
 - TCP/IP configuration
 - DIRM configuration
 - Hypervisor user ID setup
 - IBM WebSphere CloudBurst RPM
- Troubleshooting

In addition to the actual IBM WebSphere CloudBurst appliance, some configuration of the z/VM hypervisor is needed. This presentation will go over the required configuration.

Overview

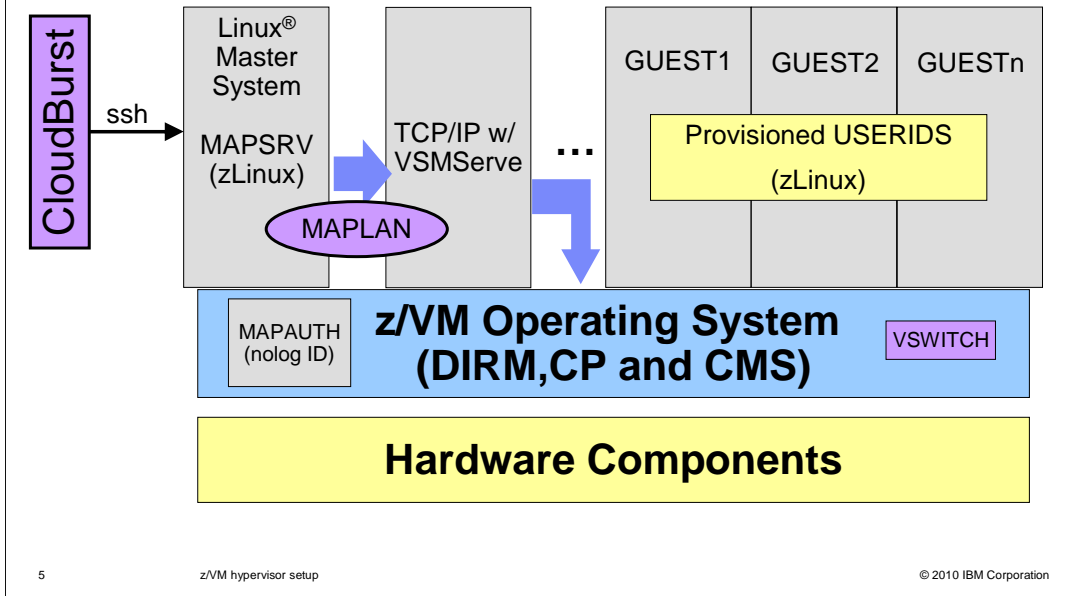
This section will present the IBM WebSphere CloudBurst configuration in the z/VM environment.

Minimum required hardware 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
 - For each provisioned guest / VM
 - 3 3390 Mod 9 DASD (21 GB of mini-disk space)
- VSWITCH configured
- Operating system requirements
 - z/VM 5.3 or higher with DIRM enabled (RACF®, optional)
 - 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. Currently, DIRM is required to be enabled on the machine, with RACF being optional. The MAINT user ID privilege is needed for initial setup of the hypervisor. Once you start provisioning with IBM WebSphere CloudBurst, a Class A user ID that you previously created is used instead.



IBM WebSphere CloudBurst on z/VM setup overview



This slide shows an overview of some of the user IDs used in IBM WebSphere CloudBurst administration and how they interact. The MAPAUTH user ID is the Class A user that issues restricted system commands through the SMAPI (via VSMSERVE) to the z/VM Operating System. This user ID is a nolog ID and is used instead of MAINT.

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. The appliance will then invoke the workflows which logon to VSMSERVE using the MAPAUTH user ID and password. VSMSERVE will take SMAPI commands used by the appliance and convert these to instructions that z/VM will understand (DIRM commands, CP commands and CMS commands). The RPM for IBM WebSphere CloudBurst z/VM is installed on the MAPSRV zLinux guest, and using the MAPAUTH user ID, is allowed to issue SMAPI commands to DirMaint™. IBM WebSphere CloudBurst will create the GUEST zLinux user IDs, as needed, to provision the virtual system you have defined.

IBM WebSphere CloudBurst – z/VM support



- Linux master system is a z/VM Linux guest:
 - With RPM to enable IBM Websphere CloudBurst to deploy VMs to z/VM system.
- Utilize z/VM System Management API for VM deployment.
 - Use of DIRM, RACF.
- Admin can customize properties of each z/VM system for deployment of VMs by Websphere CloudBurst, such as:
 - What mini-disk pools to be used for VM's dasds; for image disk cache.
 - Which vSwitch to used; networks to which VMs can connect (name, VLAN id).
 - Memory capacity (with over-commit ratio).
- Cache of images on z/VM system: image files, mini-disks (for fast disk cloning/flash-copy).
 - Faster deployment once image is cached.
- Support L2, L3 VLAN (non-)aware vSwitch.
- Minimize usage of mini-disks by sharing code binary disks (R/O).
- Extend/Capture images for customization
 - Captured image can be re-deploy to different z/VM systems
 - Export/import of images between WebSphere CloudBurst Appliances

Shown here is a summary of features available for z/VM support for the IBM WebSphere Cloudburst appliance.

***z/VM setup
SYSTEM CONFIG updates***

This section will present SYSTEM CONFIG updates that can be required on your system.

SYSTEM CONFIG updates

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

```

CP_Owned Slot 1 530RES
CP_Owned Slot 2 530SPL
CP_Owned Slot 3 530SP1
CP_Owned Slot 4 530PG1
CP_Owned Slot 5 530PG2
CP_Owned Slot 6 530PG3
CP_Owned Slot 7 530W01
CP_Owned Slot 8 530W02
CP_Owned Slot 10 530PG4
CP_Owned Slot 11 530PG5
CP_Owned Slot 12 RESERVED
CP_Owned Slot 13 RESERVED
CP_Owned Slot 14 RESERVED

```

- USER Volumes (3390 mod 9) – DASDPOOL

```

/*****
/*                               User_Volume_List                               */
/*****
USER_VOLUME_LIST CB*                /* All CloudBurst Minidisks          */

```

In order to improve system performance for z/VM, you should define five page datasets defined 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. DASDPOOL is a group of minidisks managed by DIRM and the z/VM system. For the initial release, you will use MOD9s as the DASD model type for the DASDPOOL. Use the CPFMTXA command to format the minidisks for the DASDPOOL. As a shortcut, you can format all the minidisks with the same first two characters (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 (1 of 3)

- Networking definitions

- MAPLAN

- Allows communication between the MAPSRV guest and the VSMERVE server in the TCP/IP stack

```
DEFINE LAN MAPLAN OWNERID SYSTEM MAXCONN 2 RESTRICTED TYPE QDIO IP
MODIFY LAN MAPLAN OWNERID SYSTEM GRANT TCPIP
MODIFY LAN MAPLAN OWNERID SYSTEM GRANT MAPSRV
```

- VSWITCH (Layer 2 or Layer 3; VLAN 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 NS27VSW1 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>
```

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 VSMERVE server in the TCP/IP stack. The MAXCONN statement limits this to two connections.

The VSWITCH 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, defaults owners DTCVSW1 and DTCVSW2 are used. These guests should be added to the AUTOLOG PROFILE EXEC.

SYSTEM CONFIG updates (2 of 3)

- FEATURES statement
 - Passwords_on_Cmds
 - Tells CP which commands allow passwords

```
Passwords_on_Cmds ,
Autolog yes ,
Link yes ,
Logon yes ,
```
 - Disconnect_Timeout
 - Controls whether a virtual machine is logged off after it has been forced to disconnect

```
Disconnect_Timeout off,
```
- SET statement
 - ShutdownTime/Signal ShutdownTime
 - Allows Linux to shut down cleanly before z/VM shuts down

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

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, XAUTOLOG, 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.

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 (3 of 3)

- 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      SUBC *         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 MAPSRV Linux system 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

z/VM setup TCP/IP configuration

This section will present the TCP/IP configuration changes that can be required on your system.

.

TCP/IP configuration (1 of 3)

- Define the NIC for MAPLAN to the TCP/IP stack
 - Add to 'SYSTEM DTCPARMS' file:

```
: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.

TCP/IP configuration (2 of 3)

- Add MAPLAN NIC device to the TCP/IP profile

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

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; portname NICC300 is used in the example above.

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.

TCP/IP configuration (3 of 3)

- Ensure that *VM SMAPI SERVER* and *PORTMAP* servers are started automatically

```
AUTOLOG
        PORTMAP      0          ; PORTMAP SERVER
        VSMSEERVE    0          ; VM SMAPI SERVER
        . . .
ENDAUTOLOG
```

- Add PORT definitions for PORTMAP and VSMSEERVE

```
111 TCP PORTMAP          ; Portmap Server
111 UDP PORTMAP          ; Portmap Server
172.16.0.1 845 TCP VSMSEERVE ; VM SMAPI SERVER
```

Two servers need to be added to the AUTOLOG section of the TCP/IP profile. These servers are PORTMAP and VSMSEERVE. This will ensure the applications are started when the TCP stack is started. These servers are in addition to the existing servers defined in the AUTOLOG section.

In addition to adding the servers to the TCP/IP profile, you need to add the PORT definitions. PORTMAP uses port 111 and listens on both TCP and UDP sockets.

VSMSEERVE uses port 845 and listens on a TCP socket. Also note, in the example shown above, VSMSEERVE is bound to the IP address of the MAPLAN NIC only. This prevents VSMSEERVE from accepting commands from all other interfaces on the TCP stack.

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

z/VM setup DIRM configuration

This section will present the configuration changes needed to configure DIRM for use with IBM WebSphere CloudBurst.

DIRM configuration (1 of 6)

- DIRM must be enabled!
- Update the CONFIGxx DATADVH member
 - Issue these commands:

```
ALLOW_ASUSER_NOPASS_FROM=VSMSSERVE *
ASYNCHRONOUS_UPDATE_NOTIFICATION_EXIT.xxx=DVHXNE EXEC
      (where xxx=UDP or TCP)
```
 - Add or verify these entries:

```
RUNMODE=OPERATIONAL
ONLINE=IMMED
DASD_ALLOCATE=EXACT_FF
DATAMOVE_MACHINE=DATAMOVE * *
DVHDXD_FLASHCOPY_BEHAVIOR=2
DVHDXD_FLASHCOPY_COMPLETION_WAIT=0 0
MAXIMUM_UNASSIGNED_WORKUNITS=100
```
 - Issue these to update DIRM:

```
dirm file configxx datadvh
dirm rldcode
dirm rlddata
dirm rldextn
```

DIRM needs to be enabled as part of the z/VM setup. In addition to the base configuration of DIRM, the CONFIGxx DATADVH file must be customized to allow VSMSSERVE to talk to DIRMAINT. The z/VM system can contain multiple CONFIGxx DATADVH files. The files are read by the system in reverse alphabetical order. You should name the file CONFIGZZ DATADVH to ensure IBM WebSphere CloudBurst required changes are picked up first. All DIRM commands issued must complete with RC=0. Follow the steps shown to create a customized CONFIGxx DATADVH file.

DIRM configuration (2 of 6)

- Update the EXTENT CONTROL file with IBM WebSphere CloudBurst minidisks

```

:Regions.:
  RegionId      VolSer      RegStart      RegEnd      Dev-Type      Comments
  000001        CB2001      00001         10016       3390-09
  000002        CB2002      00001         10016       3390-09
  000003        CB2003      00001         10016       3390-09
  000004        CB2004      00001         10016       3390-09
  000005        CB2005      00001         10016       3390-09

:Groups.:
*GroupName RegionList
POOL0 (ALLOCATE ROTATING)
POOL0 000001 000002 000003 000004 000005

```

Allocation groups are used to define DASDPOOLS to the z/VM system. The DASDPOOL is defined in two sections – regions and groups.

Regions consist of the pre-defined minidisks to be used by the various DASDPOOLS, or Groups. The regions are defined using a region ID value. All DASD must be formatted as a minidisk and defined in the SYSTEM CONFIG as User_Volumes.

Groups are used to define the DASDPOOLS. Each Group will consist of several region IDs.

In the SYSTEM CONFIG example shown, all of the DASD for the allocation group have a volume label of CBxxxx. In this example, CB2001-CB2005 is used as the minidisks.

DIRM configuration (3 of 6)

- Create a default directory entry that is used to define common definitions for all provisioned guests (LINDFLT DIRECT)

```
PROFILE LINDFLT
CLASS G
STORAGE 512M
MAXSTORAGE 2047M
IPL 201
IUCV ALLOW
MACHINE ESA
OPTION QUICKDSP
CONSOLE 0009 3215 T
SPOOL 000C 2540 READER *
SPOOL 000D 2540 PUNCH A
SPOOL 000E 1403 A
LINK MAINT 0190 0190 RR
LINK MAINT 019D 019D RR
LINK MAINT 019E 019E RR
LINK TCPMAINT 0592 0592 RR
```

DIRM ADD LINDFLT

- Add the new profile using the DIRM ADD command

```
dirm add lindflt
```

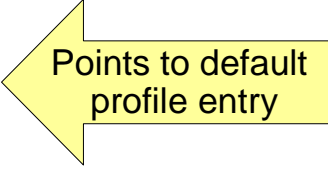
Default directory entries are used to define common definitions for all provisioned guests. This entry is referenced by the Linux prototype directory entry.

Create a file called LINDFLT DIRECT using Xedit. Shown here is an example of a default directory entry for Linux. Once created, add the new profile using the DIRM ADD command.

DIRM configuration (4 of 6)

- Create a directory prototype that will allow provisioned guests to all share common statements (LINUX PROTODIR)

```
USER LINUX NOLOG 512M 2G G  
INCLUDE LINDFLT
```



Points to default
profile entry



DIRM FILE LINDFLT PROTODIR

A directory prototype allows provisioned guests to all share common statements. This slide shows the creation of a directory prototype using the LINDFLT directory entry you saw on the previous slide. Note that multiple directory prototypes can be created to support as many configurations as are needed.

DIRM configuration (5 of 6)

- Define a Class A user that will issue restricted system commands through the SMAPI interface (MAPAUTH)

- Create MAPAUTH DIRECT:

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

- Add the guest to the z/VM directory

```
dirm add mapauth
```

- Authorize user ID to issue SMAPI commands on behalf of IBM WebSphere CloudBurst:

```
DIRM FOR ALL AUTHFOR MAPAUTH CMDLEVEL 140a CMDSET adghmops
DIRM FOR ALL AUTHFOR MAPAUTH CMDLEVEL 150a CMDSET adghmops
```

You need to define a Class A user that is allowed to issue restricted system commands through the SMAPI. This slide shows the MAPAUTH user ID being created for this purpose. The DIRM commands shown on the bottom of the slide are issued with the AUTHFOR operand to authorize one user to act on behalf of another user. Recall that the IBM WebSphere CloudBurst appliance interface to z/VM is through the MAPSRV guest.

DIRM configuration (6 of 6)

- Update the VSMERVE 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
VSMERVE	ALL	ALL
MAPAUTH	ALL	ALL

Recycle VSMERVE

Finally, you must update the VSMERVE AUTHLIST file to give the new MAPAUTH user ID permission to issue Systems Management APIs through VSMERVE. You must re-cycle VSMERVE for this change to take effect.

***z/VM setup
Hypervisor user ID setup***

This section will discuss the steps needed to create a user ID that will serve as the entry point into z/VM from the IBM WebSphere CloudBurst appliance.

MAPSRV user ID (1 of 4)

- SLES10 Linux 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 DirMaint

The MAPSRV ID is a SLES10 Linux guest which is used as the entry point into the z/VM LPAR. 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 DirMaint.

MAPSRV user ID (2 of 4)

- MAPSRV DIRECT file

```

USER MAPSRV PASSWORD 512M 1G GT
INCLUDE LINDFLT
IPL 201
MACHINE ESA
OPTION LNKNOPAS LANG AMENG
*
* DEDICATE 0201 <Addr_of Dedicated DASD for Linux System>
* MDISK 0201 3390 1 10016 CBxxxx MR PASSWORD PASSWORD PASSWORD
*
NICDEF Vswitch_OSA_Addr TYPE QDIO LAN SYSTEM NS27VSW1
NICDEF MAPLAN_OSA_Addr TYPE QDIO LAN SYSTEM MAPLAN

```



- Save the file and add the guest to the z/VM user directory

```
dirm add mapsrv
```

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

This slide shows an external of a MAPSRV user ID definition. Note that the default profile, LINDFLT that you saw on an earlier slide is included in this definition. Also shown on the slide, is a dedicated volume being used for the Linux operating system. This can also be a mini-disk, shown in the example in the commented out statement. In either case, the volume is mounted as virtual DASD address 201 and is IPL'd at start. There are two NICDEF statements needed; the first one is for the external network and the second one is for MAPLAN to talk to VSMSSERVE. Finally, save the file and add the guest to the z/VM user directory using DIRM.

MAPSRV user ID (3 of 4)

- Install SLES10 on the MAPSRV ID (see Chapter 7 of *Getting Started with Linux on System z® (v5r3)* - <http://publibz.boulder.ibm.com/epubs/pdf/hcsx0b20.pdf>)
- When installing the Linux operating system, make sure to install these pre-requisites

```
perl-libwww-perl
perl-HTML-Parser
perl-HTML-Tagset
perl-XML-DOM
perl-XML-Generator
perl-XML-RegExp
perl-XML-Writer
```
- Add these to the `/etc/sysconfig/kernel` file after install complete

```
MODULES_LOADED_ON_BOOT="vmcp"
```

SLES10 is necessary for the MAPSRV guest. You should reference the “Getting Started with Linux on System z” 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 (4 of 4)

- Use `yast` to configure MAPSRV ID to communicate with the VSMERVE application 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






Finally, you need to configure your MAPSRV guest to communicate with the VSMERVE application in 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 VSMERVE as shown on the slide. The "**z/VM System Setup**" document has more detailed instructions on how to configure this if needed.

***z/VM setup
IBM Websphere CloudBurst RPM***

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 2)

- 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 a clickable link titled **Download agent**
 - Save and upload the file to your MAPSRV id

RALNS21 Cloud 	
Description:	Ralns21 Cloud for zVM
Hypervisor type:	zVM
Type:	 Custom cloud group
Required z/VM agent:	 Download agent
Use shared minidisks:	Enable 
Status:	 All hypervisors available
Hypervisors:	RALNS21 - MAPSRV21 <input type="text" value="Add more..."/>
Access granted to:	Administrator [owner] <input type="text" value="Add more..."/>

29

z/VM hypervisor setup

© 2010 IBM Corporation

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 a clickable link titled Download agent. Save the file to your workstation, upload the file to your MAPSRV ID and install it.

IBM WebSphere CloudBurst RPM (2 of 2)

- Once installed, need to run **setProperties** to gather system data/information 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 which needs to be run in order to gather system/data information 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 2)

```

mapsrv21:/opt/ibm/zensemble # ./setProperties
Verifying vmcp...Success
Enter the IP Address of the SMAPI management server [172.16.0.1]:
Verifying ip address...Success
Enter the TCP port of the SMAPI management server [845]:
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

```

From port definition for VSMSEVERE

Class A user that is able to issue DIRM commands on behalf of IBM WebSphere CloudBurst through VSMSEVERE

```

Enter the storage over-commit factor (1 : no over commit -- max 10) [10]:
Use vswitch NS21VSW1 (y/n) [y]:
The vswitch is vlan unaware. No further configuration needed for vswitch. Hit
Enter to continue
Enter the default prototype to build guests on [Linux]:
Verifying proto...Success

```

VSWITCH defined and associated with the MAPSRV id

Name of the directory prototype to use for the VM Guests that are created (LINUX PROTODIR)

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 will go to 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. This comes from the port definition you defined earlier for VSMSEVERE. 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 DIRM command through the SMAPI. You will also need to know its 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 (for example WS pages) and estimates usage of memory for placement. Finally, if you defined a VSWITCH for use by the MAPSRV user ID, the VSWITCH ID can be discovered. You created a default prototype file to use to build the Linux guests that get provisioned to run the WebSphere Application Server. Input that here.

IBM WebSphere CloudBurst RPM...setProperty (2 of 2)

```
Enter the naming convention for new guests. Trailing Xs will be replaced with numbers [WCA21XXX]:
```

```
Enter the poolname to use [POOL0]:
```

```
Verifying pool...Success
```

```
Should mini disks be used to cache images (y/n) [y]:
```

```
Enter the poolname to use for cache [POOL1]:
```

```
Verifying pool...Success
```

```
Configuration Finished Hit Enter to continue
```

} DASDPOOL
set up for
Linux Guests

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

The first prompt shown on this slide asks you to enter the naming convention for your new guests. The trailing XXX's is replaced with a numeric value. This will result in the ID that your virtual system is registered as. Earlier you added a list of pre-defined minidisks to a group. This is used as your DASDPOOL to create the Linux guests needed for provisioning. On the slide, it is the second prompt, POOL0. The third prompt is asking you if you want to use minidisks to cache the master OVM images on the z/VM LPAR after an initial DD copy. This allows subsequent copies to use FLASHPOINT which is much quicker. A different POOL can be used for this, as shown on the slide.

setProperties (1 of 2)

```
Networks currently defined
-----

Networks discovered from zVM
-----

***1***
Name:VLAN0001
Vlan:0001
***2***
Name:VLAN0209
IP:9.42.89.97
Vlan:0209
***3***
Name:VLAN0213
IP:9.42.89.220
Vlan:0213

(A)dd a discovered network
(D)efine a new network
(R)emove a network definition
(F)inish
Choice [F]: a
```

Select 'add' to
add the
hypervisor's
VLAN

The setProperties command will try to discover VLANS that have been defined. You then have the option of adding a discovered network or defining a new one.

setProperties (2 of 2)

```

Networks discovered from zVM
-----
***1***
Name:VLAN0001
Vlan:0001
***2***
Name:VLAN0209
IP:9.42.89.97
Vlan:0209
***3***
Name:VLAN0213
IP:9.42.89.220
Vlan:0213

Select 1..3 from discovered networks. anything else to go back: 2
Vlan:0209

. . .

(A)dd a discovered network
(D)efine a new network
(R)emove a network definition
(F)inish
Choice [F]: f
  
```

(F)inish and choices are written to `/opt/ibm/zensemble/properties.xml`
 Ready for provisioning!!


Since saying you wanted to add a discovered network, you are given the option of adding it by number. Once you are satisfied, you specify you are finished and your choices are written to the `/opt/ibm/zensemble/properties.xml` file. IBM WebSphere CloudBurst is ready for provisioning now.

Troubleshooting

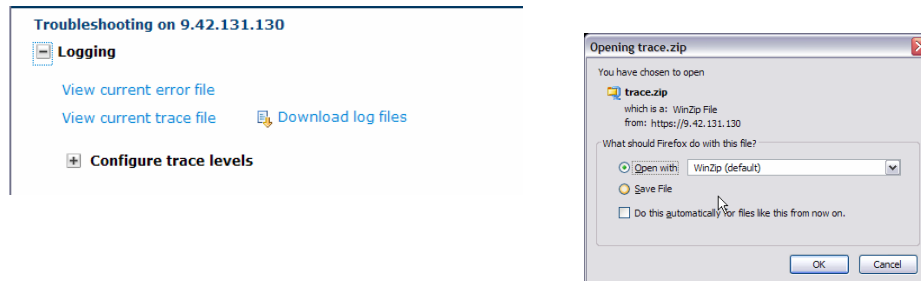
Finally, you will briefly look at where to find information if things go wrong.

Log files (1 of 6)

- Logs found on the IBM WebSphere CloudBurst appliance under 'Troubleshooting'



Expand logging; select Download log files



36

z/VM hypervisor setup

© 2010 IBM Corporation

There are logs available for you to look on from the IBM WebSphere CloudBurst appliance under the Appliance tab as shown on the slide. This logs the various CP/DIRM commands that are issued while provisioning your WebSphere Application Server environments. The "latest" logs can be downloaded by using this command, `https://<HOST_NAME>/resources/trace.zip?latest`.



Log files (2 of 6)

trace.log	Text Document	3/11/2010 2:12 PM	68,973,...	95%	3,287,...
trace_2010.03.05_18_41.42.0872.log.gz	WinZip File	3/11/2010 2:12 PM	4,150,241	26%	3,070,...
trace_2010.03.08_03.43.05.0795.log.gz	WinZip File	3/11/2010 2:12 PM	4,144,465	27%	3,037,...
trace_2010.03.09_08.13.51.0305.log.gz	WinZip File	3/11/2010 2:12 PM	4,154,337	26%	3,068,...
trace_2010.03.09_22.56.57.0870.log.gz	WinZip File	3/11/2010 2:12 PM	5,104,376	9%	4,632,...
trace_2010.03.10_13.43.01.0029.log.gz	WinZip File	3/11/2010 2:12 PM	5,013,181	9%	4,539,...
trace_2010.03.11_04.29.07.0085.log.gz	WinZip File	3/11/2010 2:12 PM	5,022,605	9%	4,555,...
trace_2010.03.03_06.51.37.0738.log.gz	WinZip File	3/11/2010 2:13 PM	4,423,735	22%	3,459,...
trace_2010.03.04_14.11.19.0768.log.gz	WinZip File	3/11/2010 2:13 PM	4,224,747	23%	3,322,...
trace_2010.03.06_23.12.30.0242.log.gz	WinZip File	3/11/2010 2:13 PM	4,142,690	27%	3,041,...
trace_2010.03.01_23.59.07.0375.log.gz	WinZip File	3/11/2010 2:13 PM	4,349,420	22%	3,385,...
error.log	Text Document	3/11/2010 2:13 PM	1,946,128	91%	168,037
appliance.log	Text Document	3/11/2010 2:13 PM	348	68%	111
audit-log.1	File	3/11/2010 2:13 PM	256,132	95%	12,671
audit-log	File	3/11/2010 2:13 PM	158,747	95%	8,068
derby.log	Text Document	3/11/2010 2:13 PM	1,734	76%	411
access.log	Text Document	3/11/2010 2:13 PM	7,250,956	92%	595,980
global-context-snapshot.log	Text Document	3/11/2010 2:13 PM	1,463,997	81%	280,392
ensembleLogs.zip	WinZip File	3/11/2010 2:13 PM	44,559,...	0%	44,35,...
default-log	File	3/11/2010 2:14 PM	53,468	95%	2,702
cli-log	File	3/11/2010 2:14 PM	109,268	92%	8,936
xact-log	File	3/11/2010 2:14 PM	37,125	96%	1,500
usage-log	File	3/11/2010 2:14 PM	0	0%	2
compile-log	File	3/11/2010 2:14 PM	0	0%	2
diag-log	File	3/11/2010 2:14 PM	100,266	95%	5,055
dbstart.log	Text Document	3/11/2010 2:14 PM	11,662	88%	1,373
appstart_rainmaker_app.log	Text Document	3/11/2010 2:14 PM	6,475	87%	873
appstart_rainmaker_pd.log	Text Document	3/11/2010 2:14 PM	6,364	87%	828
appstart_rainmaker_serialconsole.log	Text Document	3/11/2010 2:14 PM	18,321	93%	1,281
appstart_rainmaker_serialconsole_setup.log	Text Document	3/11/2010 2:14 PM	116	6%	109
appstart_rainmaker_pd_setup.log	Text Document	3/11/2010 2:14 PM	699	71%	206
appstart_rainmaker_app_setup.log	Text Document	3/11/2010 2:14 PM	1,288	68%	418
appstart_rainmaker_serialconsole_start.log	Text Document	3/11/2010 2:14 PM	255	32%	174
appstart_rainmaker_app_start.log	Text Document	3/11/2010 2:14 PM	128	6%	120
appstart_rainmaker_app_stop.log	Text Document	3/11/2010 2:14 PM	127	6%	120
appstart_rainmaker_serialconsole_stop.log	Text Document	3/11/2010 2:14 PM	127	6%	119
appstart_rainmaker_pd_start.log	Text Document	3/11/2010 2:14 PM	22,362,...	99%	177,877
appstart_rainmaker_ensemble.log	Text Document	3/11/2010 2:14 PM	6,393	87%	834
appstart_rainmaker_pd_stop.log	Text Document	3/11/2010 2:14 PM	127	6%	120
supervisor.log	Text Document	3/11/2010 2:14 PM	7,221	87%	911
cbHealth.log	Text Document	3/11/2010 2:14 PM	16,287	90%	1,651
setupfs.log	Text Document	3/11/2010 2:14 PM	60,221	91%	5,209
cbserial.log	Text Document	3/11/2010 2:14 PM	1,906	77%	435
appliance-get-status-snapshot.log	Text Document	3/11/2010 2:14 PM	86,414	92%	6,900
error-report.txt	Text Document	3/11/2010 2:14 PM	616,172	79%	128,521

37

z/VM hypervisor setup

© 2010 IBM Corporation

Shown here is an example of the extracted files that are part of the trace.zip file that you just downloaded.

Log files (3 of 6)

- Example of trace log

```
1 ***** Start Display Current Environment *****
2 WebSphere CloudBurst version number is 1.1.1.0-19309
3 Java Version = J2RE 1.6.0 IBM J9 2.4 Linux x86-32 jvmpi3260-20081105_25433 (JIT enabled, AOT enabled)
4 J9VM - 20081105_025433_1HdSMr
5 JIT - r9_20081031_1330
6 GC - 20081027_AB,Java Compiler = j9jit24,Java VM name = IBM J9 VM
7 Application home = /opt/ibm/rainmaker/rainmaker.app
8 Disk root = /drouter/ramdisk2/mnt/raid-volume/raid0
9 Current trace specification: {zero.core.security.TokenService=OFF, com.ibm.vespa.util=ALL, com.ibm.ws.vm.
```

On this slide is an example of an entry from the trace.log file.

Log files (4 of 6)

- Sample Steps from the trace.log file

```
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
```

39

z/VM hypervisor setup

© 2010 IBM Corporation

Here is an example of the various steps that are run during the IBM WebSphere CloudBurst initialization process. In a full log there are many messages between each step. You will also see failure messages, if there are any, during the step processing.

Log files (5 of 6)

- Example of trace log from zensemble.zip

```
1 ***** Start Display Current Environment *****
2 WebSphere CloudBurst version number is 1.1.1.0-19309
3 Java Version = J2RE 1.6.0 IBM J9 2.4 Linux x86-32 jvmpi3260-20081105_25433 (JIT enabled, AOT enabled)
4 J9VM - 20081105_025433_1HdSMr
5 JIT - r9_20081031_1330
6 GC - 20081027_AB,Java Compiler = j9jit24,Java VM name = IBM J9 VM
7 Application home = /opt/ibm/rainmaker/rainmaker.ensemble
8 Disk root = /drouter/ramdisk2/mnt/raid-volume/raid0/rainmaker.ensemble
9 Current trace specification: null
10 ***** End Display Current Environment *****
```

- Example of error.log

```
W app.scripts.groovy.rainmaker.cloud.PhysicalMachines.groovy getPhysicalMachineNetworks CWZCL3024W: Get failed. Valid physical machine id not specified
I app.scripts.groovy.rainmaker.instances.stages.Stage1Email.groovy run CWZC00001I Step 1: Starting email stage to requesting user
I app.scripts.groovy.rainmaker.instances.stages.Stage2Placement.groovy run CWZC00003I Step 2: Starting virtual system placement calculation for pattern
E com.ibm.websphere.rainmaker.patterns.PatternModelPatternUtils updatePartForDeploymentInternal CWZCL2622E: Unable to execute requested update - attribute
```

This slide shows some additional logs that can be examined in the event of a failure. The zensemble.zip is a platform specific file that contains logs used for discovery and deployment information.

Log files (6 of 6)

```
[]: ===== Tue Jul 21 18:29:27 EDT 2009 =====  
[]: /opt/ibm/zensemble/workflow/zlinuxInstall.sh  
[]: =====  
[]: parse_config_file input is /tmp/zemapi6317378220663381521.xml  
[]: stack push xmlTags config0  
.  
.  
[96347]: run executing /opt/ibm/zensemble/bin/uhubrpcclient 172.16.0.1 845 mapauth xxxxxxxx  
96347 dirm mapauth ns21cb01 ipl 0107  
[96347] Common parms: VSMSERVE 172.16.0.1, Port 845, ID mapauth, PW *****, MsgID 96347  
[96347] dirm command  
[96347] dirm parameters entered: targetID=ns21cb01, dirmCmd=ipl 0107  
[96347] Connecting...  
[96347] Connected.  
[96347] Loggin In...  
[96347] Logged In.  
[96347] Now, invoking dirm command...  
[96347] dirmCommand returned, rc = 0  
[96347] Disconnecting...
```

You will see here that a DIRM command completed successfully. If there was a problem, you will see that as well and can hopefully fix whatever might be preventing the command from completing successfully.

Summary

This section summarizes the appliance initialization presentation.

Summary

- Overview
- z/VM setup
- Troubleshooting

In summary, this presentation looked at the configuration needed on the z/VM system (the hypervisor) in order for IBM WebSphere CloudBurst to successfully provision WebSphere Application Server environments. You started by looking at an overall interaction of the various pieces and then drilled down into how to configure each of the pieces. You ended up with 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 e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_CB111_zVMHypervisorConfiguration.ppt

This module is also available in PDF format at: [../CB111_zVMHypervisorConfiguration.pdf](http://CB111_zVMHypervisorConfiguration.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, DirMaint, RACF, System z, and WebSphere 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](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

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.

© Copyright International Business Machines Corporation 2010. All rights reserved.