IBM Wave

IBM Wave for z/VM: User Guide and Reference

Version 1 Release 2

SC27-6119-14

Note:
Before using this information and the product it supports, read the information in "Notices" on page 253.

This edition applies to Version 1, Release 2 of IBM Wave for z/VM (product number 5648-AE1) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2007, 2018.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

© Copyright IBM Corp. 2007, 2018

Tables

© Copyright IBM Corp. 2007, 2018

Figures

© Copyright IBM Corp. 2007, 2018

About this information

This document supports IBM® Wave for z/VM® (5648-AE1).

IBM Wave for z/VM is a provisioning and productivity management solution for simplifying the control and use of virtual Linux servers and z/VM. IBM Wave significantly reduces the learning curve needed to manage and control z/VM and Linux guests. The information describes how you can maintain and customize IBM Wave for z/VM to meet your requirements.

Intended audience

This information is intended for new and experienced Linux system administrators and z/VM system administrators who are responsible for managing servers with their current skill set. IBM Wave for z/VM significantly reduces the learning curve that is needed to control the sophisticated z/VM environment and manage large Linux installations.

Links to documents and websites

The PDF version of this information contains links to other documents and websites. A link from one PDF file to another PDF file works only when both files are in the same directory or database. Links to websites work when you have internet connectivity. A document link is to a specific edition. If a newer edition of the linked documents is published, ensure that you have the current edition.

© Copyright IBM Corp. 2007, 2018

How to send your comments to IBM

We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or give us any other feedback that you might have.

Use one of the following methods to send us your comments:

1. Send an email to mhvrcfs@us.ibm.com.
2. Mail the comments to the following address:
   IBM Corporation
   Attention: MHVRCFS Reader Comments
   Department H6MA, Building 707
   2455 South Road
   Poughkeepsie, NY 12601-5400
   U.S.A.
3. Fax the comments to us as follows:
   From the United States and Canada: 1+845+432-9405
   From all other countries: Your international access code +1+845+432-9405

Include the following information:

- Your name and address
- Your email address
- Your telephone or fax number
- The publication title and order number:
  IBM Wave for z/VM V1R2 User Guide and Reference
  SC27-6119-14
- The topic name or page number related to your comment
- The text of your comment

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

IBM or any other organizations will use the personal information that you supply only to contact you about the issues that you submit to IBM.

If you have a technical problem

Do not use the feedback methods listed above. Instead, do one of the following:

- Contact your IBM service representative.
- Contact IBM technical support.
- See IBM: z/VM Service Resources (www.ibm.com/vm/service/).
- Go to IBM Support Portal (www.ibm.com/support/entry/portal/Overview/).

If you have a technical problem

Do not use the feedback methods listed in "How to send your comments to IBM." Instead, use one of the following methods:

© Copyright IBM Corp. 2007, 2018

- If you plan to contact IBM Wave Support to open a Service Request or Problem Management Report (PMR), review the "MustGather" instructions for sending data to IBM Wave Support at www.ibm.com/support/docview.wss?uid=isg3T1022347.
- To contact your IBM z/VM Support representative, go to IBM Support Portal (www.ibm.com/support/entry/portal/Overview/).
- See IBM: z/VM Service Resources (www.ibm.com/vm/service/).

Summary of changes for IBM Wave for z/VM: User Guide and Reference

This information includes terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Summary of changes as updated October 2017

This information describes the summary of changes for *IBM Wave for z/VM V1R2: User Guide and Reference* as updated October 2017.

Updated information:

This information includes only terminology, maintenance, and editorial changes.

Summary of changes as updated March 2017

This information describes the summary of changes for *IBM Wave for z/VM V1R2: User Guide and Reference* as updated March 2017.

New information:

- IBM Wave provides a RESTful API service. For information, see Chapter 6, "IBM Wave API," on page 237.
- The BTS Work unit viewer has a new detailed view, which allows filtering the displayed work units. For more information, see "Work Units Viewer" on page 50.

Summary of changes as updated December 2016

This information describes the summary of changes for *IBM Wave for z/VM V1R2: User Guide and Reference* as updated December 2016.

Updated information:

- Updated information is provided about how to update IP addresses when cloning by using a remote IBM Wave. For information, see "Connecting z/VM guests to virtual network segments" on page 136 and "Clone from a Prototype" on page 141.
- After Bare Metal Installtion (BMI), to use Live Guest Relocation (LGR) on Linux guests in an SSI cluster, you must deactivate and then activate the guest. See "Installing Linux with the BMI wizard" on page 88 or "Relocate to..." on page 128.

Deleted information:

- The domain parameters for LDAP and Active Directory are no longer options of the command line interface (CLI).

Summary of changes as updated September 2016

This information describes the summary of changes for *IBM Wave for z/VM V1R2: User Guide and Reference* as updated September 2016.

New information

© Copyright IBM Corp. 2007, 2018

- When you first start IBM Wave, after service pack 5 (SP5), a **What's New** pane displays the new function in IBM Wave. The pane contains a dismiss option. To reopen the pane, select **Help** > **What's New**.
- IBM Wave for z/VM supports LOGONBY access when adding a new z/VM system to IBM Wave for z/VM. For more information, see "Add New System" on page 60.
- You can now use IBM Wave to dynamically increase memory. For more information, see the procedure for "Adding Memory Dynamically" on page 108. Though not explicitly stated in this manual, you can dynamically increase memory when you use the following tasks:
  - "Create New z/VM Guest" on page 72
  - "Clone" on page 81
  - "Display Information" on page 82
  - "Duplicate z/VM User Definition" on page 85
  - "Update z/VM User" on page 101

Summary of changes as updated July 2016

This information describes the summary of changes for *IBM Wave for z/VM V1R2: User Guide and Reference* as updated July 2016.

New information

- IBM Wave for z/VM supports Ubuntu Server 16.04. For more information, see the following topics:
  - "Installing Linux with the BMI wizard" on page 88.
  - "Clone from a Prototype" on page 141.
- You can now use IBM Wave to dynamically increase the number of central processing units (CPUs) that are assigned to an active guest without recycling the guest. For more information, see the procedure to "Add CPUs Dynamically" on page 106 and the following topics:
  - "Create New z/VM Guest" on page 72
  - "Clone" on page 81
  - "Display Information" on page 82
  - "Duplicate z/VM User Definition" on page 85
  - "Update z/VM User" on page 101
  - "Clone from a Prototype" on page 141.
- To ensure accurate values are displayed in the **Storage Viewer**, new messages are issued to notify you about any inconsistencies.
- IBM Wave can issue a warning message if anyone, other than the system level administrator (SLA), tries to activate a guest on a z/VM system that is not the guest's default system. For more information, see the following topics:
  - "Activate" on page 110
  - "Add CPUs Dynamically" on page 106.
- "Toggle Table Viewer" on page 8 is a new topic that explains the available column in the **Table Viewer**, which is available in the **Enterprise Viewer** and **Hardware Viewer**.
- A prerequisite for the command line interface (CLI) is added to "CLI package" on page 231.

- The **Execute Rexx** function no longer times out after 60 seconds. For more information, see "Execute REXX" on page 76.

Changed information:

- Additions to accurately reflect the function in "Detailed performance - Virtual Machines Utilization tab" on page 43.
- "Run Automatic Guest Classification" on page 131 contains updated information, and a link to the complete information in *IBM Wave for z/VM Administration and Customaization*.
- Changes to the Chapter 5, "Command line interface," on page 231 (CLI) information:
  - In "CLI parameters" on page 232, information is added about the CLI parameter order and syntax, and the domain parameter is updated with correct information.
  - The "CLI syntax" on page 232 is updated to reflect the required parameter, optional parameter, and choice items.
  - The "CLI commands" on page 234 contains a new example about parsing out the active guests.

Errata In the "Clone from a Prototype" on page 141 topic, the information erroneously claimed that the **Optional Linux parameters** > **Regenerate SSH Keys** option was not supported. "Clone from a Prototype" on page 141 now contains the correct information.

Summary of changes as updated April 2016

This information describes the summary of changes for *IBM Wave for z/VM V1R2: User Guide and Reference* as updated April 2016.

New information

- Beginning with V1R2 service pack (SP) 3, IBM Wave supports IBM z13s, IBM LinuxONE Rockhopper™, and IBM LinuxONE Emperor™.
- Beginning with V1R2 SP3, IBM Wave supports FTP Secure (FTPS). The FTPS protocol is described in RFC 4217 "Securing FTP with TLS". FTPS is used by default, unless IBM Wave detects that your installation is using standard FTP.
- "Interactive Toolbar" on page 6 is a new topic.
- New information about using authorized credentials for Linux guest actions. See the following topics:
  - "Init Users for IBM Wave use" on page 55
  - "z/VM guest and virtual server functions" on page 78
  - "Refresh Linux Data" on page 128
- "Purge Spool" on page 67 is updated with new options.
- IBM Wave contains a new SSH access menu for Linux users. For information, see "SSH Access" on page 98.
- A new task is added for "Running scripts from guests" on page 229.
- "CLI commands" on page 234 contains new commands for pause, recycle, and resume.

Changed information:

- The following topics are updated
  - "Special considerations for site defined groups" on page 25
  - "Transferring z/VM guests between groups" on page 25

  - "Transferring z/VM guests" on page 70
- You must select the Linux Guests when you use the "Init Users for IBM Wave use" on page 55 process.
- Additional guidance was added for the `fstab` mount point. See "Manage Storage" on page 93 and "Create or extend a DASD partition" on page 94.
- "Activate" on page 110, "Deactivate z/VM Users" on page 118, and "Recycle z/VM users" on page 127 are clarified.
- "Generate Disk Storage Map" on page 123 is updated to contain information that is consistent with the current user interface.
- Clarification about the permissions, and file name rules are added to the Chapter 4, "Script management subsystem," on page 225.
- "CLI syntax" on page 232 contains updated syntax.
- As an intermediate fix, the "CLI parameters" on page 232 clearly identify the required and optional parameters.

Deleted information:

- The topic about "Connect to virtual network" is removed. Instead, see "Connecting z/VM guests to virtual network segments" on page 136.
- The external entities topics, "Create an External Entity" and "Update information", are removed. Instead, see "External entities functions" on page 171 and "Create a new external entity" on page 171.

Chapter 1. IBM Wave User Interface

The IBM Wave user interface contains the IBM Wave Main menu, toolbar, and viewers that are illustrated in Figure 1.

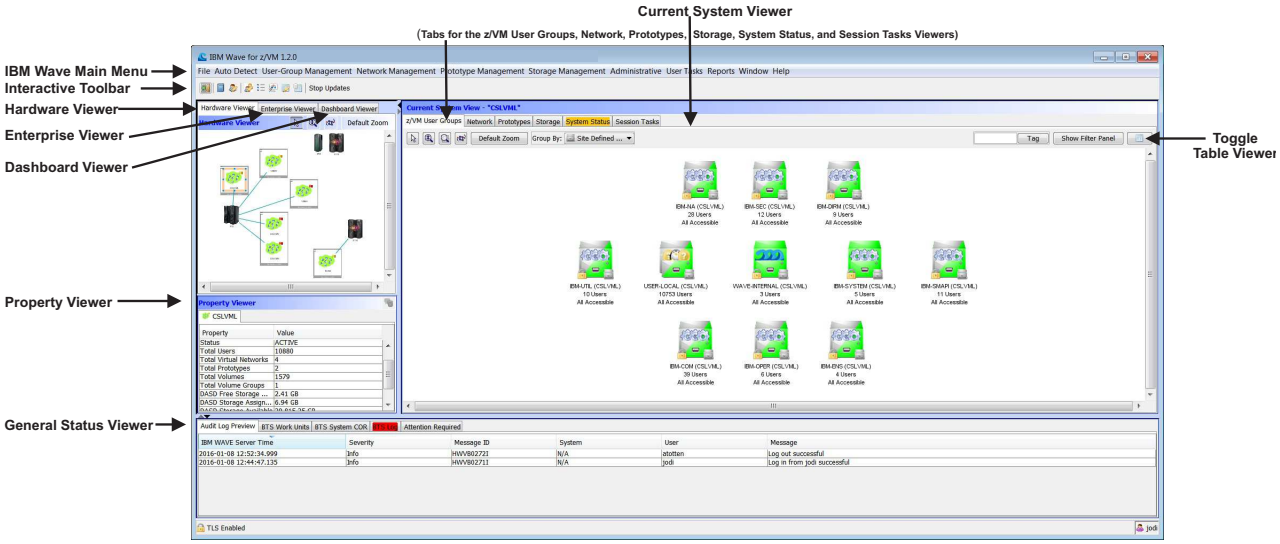

Figure 1. The IBM Wave user interface

For the viewer types, see "IBM Wave viewer types" on page 2.

For the menu items and actions controlled by using the Interactive toolbar, see "Interactive Toolbar" on page 6.

After you select one of the following major viewers, the viewer displays in the **Current System Viewer**:

- **Hardware Viewer**: When you select the **Hardware Viewer**, and an eligible z/VM System is selected, the **Current System Viewer** contains the tabs for z/VM Users Group, Network, Prototypes, Storage, System Status, and the Session Tasks. For more information, see "Hardware Viewer" on page 9. For information about each tab within the **Hardware Viewer**, see the following topics:
  - "z/VM User Groups Viewer" on page 24
  - "Network Viewer" on page 30
  - "Prototype Viewer" on page 35
  - "Storage Viewer" on page 36
  - "Session Tasks Viewer" on page 40
  - "System Status Viewer" on page 41.
- **Enterprise Viewer**: When you select the **Enterprise Viewer**, the **Current System Viewer** contains the **Guests Viewer**. The viewer is populated with guest groups based on the initial filter settings for the **Enterprise Viewer**, but can be adjusted by using the "Group-By" setting. For more information, see "Enterprise Viewer" on page 14.

© Copyright IBM Corp. 2007, 2018

- **Dashboard Viewer**: When you select the **Dashboard Viewer**, the **Current System Viewer** contains the CPC or z/VM System information that matches the "Filter selection" you select in the **Enterprise Dashboard Viewer**. When applicable, the table can contain performance information for the z/VM Systems. For more information, see "Dashboard Viewer" on page 14.

For common tools that you can use in each viewer, including how to **Toggle the Table Viewer**, see "Tools common to Diagram Viewers" on page 3.

Notes:

- The **Current System Viewer** retains its subcomponent view when switching between the **Hardware Viewer**, **Enterprise Viewer**, and **Dashboard Viewer**.

IBM Wave viewer types

The IBM Wave Viewers comprise the main window of the IBM Wave client. Each viewer shows a different aspect of IBM Wave management to help you manage the virtualized environment in an intuitive and efficient manner.

IBM Wave includes three types of viewers in the user interface:

Diagram Viewers
: The **Diagram Viewer** displays a diagram of objects. The **Diagram Viewer** can also have tagging and filtering capabilities, and some viewers include minor viewers. There are two major **Diagram Viewers** in the IBM Wave interface:

  Hardware Viewer
  : Provides you with a method to interact with a managed CPC and z/VM Systems.

  Current System Viewer
  : Provides you with a method to interact with the selected z/VM System.

Each **Diagram Viewer**, except for the **Network Viewer**, has an alternative **Table Viewer**. The **Table Viewer** retains the tagging or filtering done on the diagram viewer. Conversely, the diagram viewer retains tagging or filtering done on the table viewer.

Table Viewers
: Each **Table Viewer** displays a table that contains objects with attributes. You can rearrange and sort columns in many table viewers. You can select which columns are visible or hidden, except for the first two columns that indicate the objects ID and z/VM System. Many of the data tables can also be exported to a comma-separated value .csv file by using the **Export to CSV** function.

Regular Viewers
: The Regular Viewers display informational data. There are two major **Regular Viewers** in IBM Wave:

  General Status Viewer
  : Displays the viewers for the **Audit Log Preview**, **BTS Work Units**, **BTS System COR**, **BTS Log**, and **Attention Required**.

  Property Viewer
  : Provides you with a method to see the key properties of any selected object.

Each major or minor **Diagram Viewer** displays a diagram of objects. Each of these objects represents a z/VM object or an IBM Wave object. Each object is represented by an icon that provides information about the object. Because each **Diagram Viewer** displays different types of z/VM objects, the types of objects for each viewer are documented. Each diagram viewer has a toolbar with tools for manipulating the view and interacting with the objects in the viewer. Some tools are common to all viewers and others are viewer-specific.

In each **Diagram Viewer**, except for the **Hardware Viewer**, you can select multiple objects, and then act on all of the selected objects. To select multiple objects, either:

- Click an object while also holding the SHIFT key.
- Drag a marquee to select all objects inside the marquee square.

Tools common to Diagram Viewers

The following tools are common to all **Diagram Viewers**:

Figure 2. Diagram Viewers: common tools

The Select Tool
The **Select Tool** allows you to select an object, and then right-click to display its menu.

The Zoom Tool
The **Zoom Tool** allows you to zoom in and out of a diagram. To use the zoom tool, click and drag. To zoom in, drag the mouse down. To zoom out, drag the mouse up. The center point of the zoom is the part if the screen that you first click.

Note: It is also possible to zoom by using the mouse wheel (if applicable). To zoom by using the mouse wheel, rotate the mouse wheel while you hold the SHIFT key.

The Marquee Zoom Tool
The **Marquee Zoom** tool allows the IBM Wave User to zoom in on a specific rectangle.

The Rotate Tool
The **Rotate Tool** allows the IBM Wave User to rotate the diagram. The rotation is done by click and drag. After dragging starts, the diagram rotates in the direction the mouse is moved. For the best rotation results, rotate the mouse in a circular manner.

The Default Zoom Button
The **Default Zoom** button resets the zoom level of the diagram to fit all elements on the screen.

The Toggle Table Viewer Button
The **Toggle Table Viewer** button toggles between the Diagram Viewer and the Table Viewer for the current Viewer.

Common options for diagram viewers

All diagram viewers provide an option to right-click within the viewer and select the following options:

Arrange Items
The **Arrange Items** option arranges the objects on the screen.

Select All
The **Select All** option is viewer-dependent and provides a method for a user to select all objects of a certain type. For example, in the Network Viewer, you can select all z/VM Guests, all Virtual Networks, all Virtual Network Segments, and more.

Sort The **Sort** option is viewer-dependent and provides a method for a user to sort all objects of a certain type in the **Hardware Viewer** and **Enterprise Viewer**. For example, in the z/VM User Group, you can select all z/VM Guests, and then click to sort them by name in ascending or descending order.

Icons common to all diagram viewers

Generally, icons that are in viewers and minor viewers are built by using a basic icon which represent the type of the object such as z/VM Guest, Virtual Network, z/VM Prototype, and others. Badges are added to different corners of an icon to represent status such as type or warning condition.

The following badges are common to all objects in the **Current System Viewer**:

IAN badge
The note badge, in Figure 3, indicates that there is an intelligent active note (IAN) is attached to the object.

Figure 3. IAN badge

Lock badge
The lock badge, in Figure 4, indicates that the object is locked. If the object is locked because the assigned group or project is locked, another indication appears inside the lock.

Figure 4. Lock badge

Sorting capability for viewers

The optional sort feature can help you organize objects within views.

To provide greater control and better organization of the IBM Wave user interface, you can sort the order in which objects appear in the Hardware, Enterprise, and Network viewers. The sort order is based on the objects' attributes. For example, to make guests easier to locate, you can sort them in alphanumeric order based on the name. To make groups easier to identify, you can sort them by name or by the number of users in the group. The sort order for objects is persistent for each IBM Wave user.

The sort order is left-to-right and top-to-bottom for all views except for the Network viewer.

In the Network viewer, the sort order for the top level set of objects is left-to-right. Subsequent sets of child objects are sorted as distinct groups. For example, a group of guests that are connected to *VNSA* are sorted separately from a group of guests that are connected to *VNSB*. In the Network viewer, the sort capability is available only in hierarchical layouts.

Table 1 lists the sortable objects that are available in the Hardware and Enterprise viewers. An object's attributes can be sorted in ascending or descending order, and all objects can be reset.

To sort objects, right-click the white space in the view, and then select the sort type.

Table 1. Sort capability for the Hardware and Enterprise Viewers

Viewer name and sortable object	Sort capability by attribute
Hardware viewer > z/VM User Groups	Name Group size Reset
Hardware and Enterprise viewers > Guests	Name Number of Users Reset
Hardware viewer > Network	Name Reset
Hardware viewer > Disk Storage map	Name Reset
Hardware viewer > Prototype	Name Reset
Hardware viewer > Storage view: Group tab	Free space Name Total size Reset
Hardware viewer > Storage view: Volumes tab	Address Free space Name Total size Reset

Tagging capability

IBM Wave is designed to manage even the largest z/VM environments. Large environments can contain thousands of virtual servers and storage volumes, dozens of Virtual Networks, and more. Tagging is an intuitive way to locate objects, so some of the diagram viewers have tagging capabilities to help an IBM Wave user quickly locate an object by name.

To tag an object, input a search string into the tag box, and press "Tag". The wildcard "*" is automatically appended to the beginning and end of the tag string. All objects in the current viewer are highlighted and selected. All of the object properties appear in the Property Viewer, and you have the ability to perform

multiple actions on them (where applicable). If the objects are inside groups, expand the group to see the tagged objects.

Filtering capability

All of the diagram viewers have filtering capabilities. Filtering capability allows you to display only the objects of interest and hide objects that are of no interest. For example, you can set a filter to view all of the z/VM Guests that are connected to a specific Virtual Network or to view only the storage volumes that have a minimum of 2 GB of free space.

Filtering is done with filter strings. All filter strings can contain the "*" wildcard to indicate "match 0 or all characters".

Because each of the viewers display a different set of objects, filtering is viewer specific. Filters are saved across sessions. When you exit IBM Wave, the filters are saved and are reapplied the next you login to IBM Wave.

Filters are saved for each of the minor viewers per z/VM system and the Enterprise Viewer, Dashboard Viewer, and General Status Viewer. For example, you select z/VM System A using the Hardware Viewer and set a filter for the Guests and Groups Viewer. The filter is saved and applied every subsequent time you select z/VM System A. If you select z/VM System B no filter is applied because you did not set a filter for System B.

Filtering is used in the Enterprise Viewer and the Dashboard Viewer to determine the initial selection of the z/VM Guests or z/VM Systems to populate in the Main Viewer.

As shown in "Filtering capability," you can chain filters together. IBM Wave uses the logical "AND" operation for chaining.

For example, you select the "z/VM Guest Name" filter, and then enter the value "D*". Next, you select the "z/VM Guest Status" filter, and then enter the value "Active". IBM Wave displays only the "active" z/VM Guests whose name starts with "D".

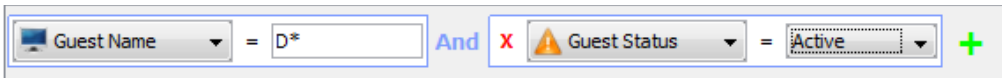

Figure 5. Filtering capability

Each filter setting is based on the filtering criteria from the first menu, and an assigned value from the second menu or input field. Some of the filters have predefined values that appear in the menu, and some are free-text that can be entered in the input field. Pressing the green "+" opens an additional filter box. Pressing the red "X" removes the filter box from the filter settings.

To manipulate the filters for the current Viewer, display the filter panel by clicking "Show Filter Panel", which is on the top right of the viewer.

Interactive Toolbar

The **Interactive Toolbar** contains shortcuts to frequently used menu items.

Use the **Interactive Toolbar** to quickly open items that are common to the **IBM Wave Main Menu**. For example, clicking the blue book icon opens the **IBM Wave Audit Log Display**.

Figure 6. Interactive Toolbar

Each icon illustrated in Table 2 is followed with a description of the action or menu item:

Table 2. Interactive Toolbar icons and actions

Icon	Action
	Exit IBM Wave
	Opens **IBM Wave Audit Log Display** For more information, see "View Audit Log" on page 50.
	Opens **IBM Wave Workunit Viewer** For more information, see "Work Units Viewer" on page 50.
	Opens **Change User Preferences** menu For more information, see "Change IBM Wave User Preferences" on page 51.
	Opens **IBM Wave User Manager** For more information, see the topic about "Creating and updating IBM Wave users" in .
	Opens **IBM Wave User Profile Manager** For more information, see the topic about "Creating and updating IBM Wave User Profiles" in .
	Opens **Project Manager** For more information, see "Display projects" on page 52.
	Opens **IBM Wave Parameters** For more information, see the topic about "System customization" in .
	Opens **IBM Wave Resource Serialization Elements** For more information, see the topic about "Wave resource serialization (WRS)" in .
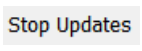	Stops processing of IBM Wave requests **Note:** When you click **Stop Updates**, all IBM Wave logging actions are disabled.
	Resumes processing of IBM Wave requests (*number of requests in the queue*) **Note:** You must click **Process Updates** to resume normal processing. When you resume the session, the changes are processed in FIFO order. For more information, see the topic about "Stop Updates" in .

Toggle Table Viewer

When you toggle the **Table Viewer**, in the **Enterprise Viewer** or the **Hardware Viewer**, you can customize, view, and export system data in table format.

When you toggle the **Table Viewer** in the **Enterprise Viewer**, all of the systems are presented in table format. When you toggle the **Table Viewer** the **Hardware Viewer**, only the currently selected system is presented in table format.

To toggle the **Table Viewer**, select the table icon in the upper right of the viewer as highlighted in Figure 7. To toggle back to the **Diagram Viewer**, select the diagram icon.

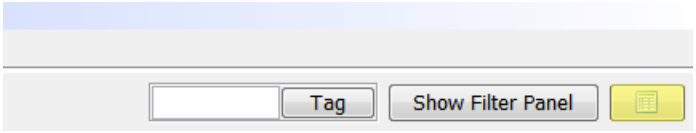

Figure 7. Toggle the Table Viewer

To choose which columns are visible or hidden, in the **Table Viewer**, click **Select Fields...**. The **Table Column Selector for Guests** opens.

- Use the left and right arrows to add or remove one or multiple fields from the table.
- Use the up and down arrows to change how the columns are displayed in the table.

The first two columns, "z/VM Guest Name" and "z/VM System Name" cannot be removed, but you can sort the "z/VM Guest Name" in ascending or descending order. When you use the **Enterprise Viewer**, you can also sort the "z/VM System Name" in ascending or descending order.

To export the table, click **Export to CSV** and save the file to your local workstation.

The table contains the following fields:

Default z/VM System
The default z/VM System to which the guest belongs.

Site Defined Group
The Site Defined Group (SDG) to which the guest belongs.

Hostname
The host name for the Linux guest.

Locker
The name of the user who is using the guest (the guest is locked).

Distribution
The name of the Linux distribution, if applicable.

Status Indicates the guest status.

Project
The name of the "Project," if applicable (if you are using projects).

Functionality
The name of the functionality level that is defined, if applicable.

Connectable
Indicates whether the z/VM Guest has IP connectivity to the Wave server (WAVESRV).

Description
The description of the guest, if applicable.

Prototype
The name of the clone prototype, if applicable.

Original User ID
The name of original guest that is used for the clone operation, if applicable.

Memory Min(MB)
The minimum amount of memory that is allocated to the guest.

Memory Max(MB)
The maximum amount of memory that is allocated to the guest.

Number of CPUs
The number of CPUs that are assigned to the guest.

Directory maximum CPUs
The maximum number of CPUs that can be assigned to a guest.

Disk Space(GB)
The total disk space that is allocated for the guest.

Account
The name of the account, if applicable.

AGC Status
Indicates whether Automatic Guest Classification (AGC) is set. When set, the table displays the AGC status.

Created By
The name of the user or task, such as auto-detect, who created the guest or started the task.

Modified By
The name of the user or task, such as the periodic guest update, who last modified the guest or ran the task.

IP 1, 2, or 3
One of three possible IP addresses for the guest.

VNS Name 1, 2, or 3
One of three possible VNS names for the guest.

Virtual Network 1, 2, or 3
One of three possible virtual networks for the guest.

Custom Attributes
The value of the user-defined custom attribute, if applicable.

Hardware Viewer

The Hardware Viewer displays the z Systems® (CPC) and z/VM Systems and SSI clusters that are managed by IBM Wave (that are within your scopes and permissions).

Hardware Viewer

The **Hardware Viewer**, in the upper-left main IBM Wave window, provides information about the z/VM Systems, z/VM LPARs, and z Systems (CPC) that are managed by IBM Wave.

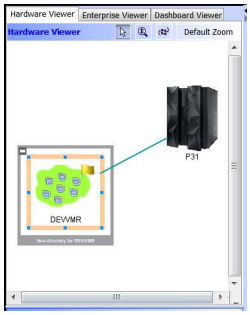

Figure 8. Hardware Viewer: z/VM system

All CPCs that are managed by IBM Wave appear in the **Hardware Viewer**. Managed z/VM Systems appear when the users scope is defined to include objects from the z/VM System.

For example, in Figure 8, the image shows that the current IBM Wave user has at least one scope entry defined that contains objects from the "**DEVVMR**" z/VM System. It is also possible, that there are other z/VM Systems managed by IBM Wave, but the user does not have scope entries defined for them.

z/VM Systems in the **Hardware Viewer** are grouped based on the z/VM Guest directory that they are using. For example, if SYSA and SYSB are in a shared directory complex, they are grouped together.

z/VM Systems that are members of an SSI cluster are grouped in a similar manner (because they are fundamentally sharing a directory). However, the color of the group (light blue) indicates that they are members of an SSI cluster. For example, in Figure 9 on page 11, the name of the group is "**SSITEST**". Instead of the name of the directory, the name of the SSI cluster is used.

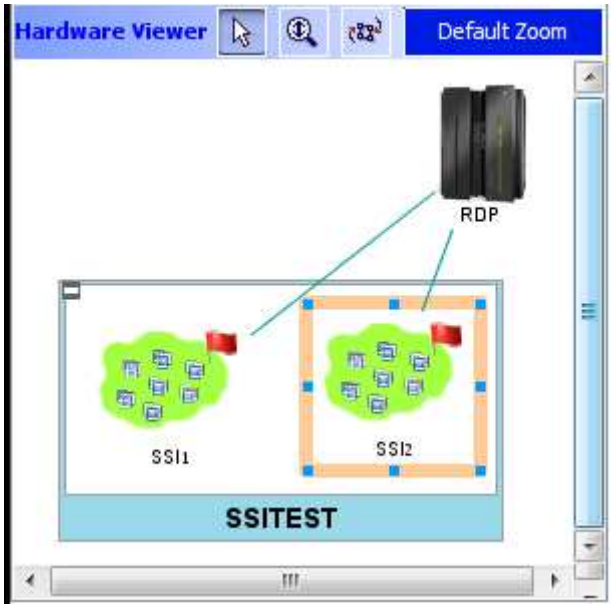

Figure 9. Hardware Viewer: SSI cluster

You can use the output of the **CP Q SSI** command to find the name of an SSI cluster.

Using the **Hardware Viewer**, what you can manage depends on your scopes and permissions.

- IBM Wave users can interact with the managed z/VM Systems and CPCs.
- IBM Wave Site Level Administrators (SLA) can add and remove CPCs, and add and remove z/VM Systems from CPCs. (Only IBM Wave users in the SLA role can add, remove, and update CPCs and z/VM Systems.)

Special tools: Hardware Viewer

There are no special tools available in this viewer.

Related information:

"Hardware functions" on page 58

Hardware Viewer icons

The Hardware Viewer can display CPC and z/VM System (z/VM LPAR) objects:

- "CPC icons"
- "z/VM System icons" on page 12.

CPC icons

Each processor (CPC) is represented by an icon that matches the machine type. As shown in Figure 10 on page 12, the IBM z Systems® mainframe, the 2964 series CPC, is represented by a z13® icon.

Figure 10. Example of an IBM z Systems icon

z/VM System icons

Every z/VM System in IBM wave has an icon that indicates the status of the system. The system status can be active or inactive. Extra icons indicate system conditions such as suspended, paused, suspended, and others as shown in Table 3.

Table 3. z/VM System icons

z/VM System status	z/VM System Icon
An active z/VM System that is accessible and functioning properly is represented by a green cloud.	 *Figure 11. Green cloud icon*
An inactive z/VM System that is not accessible is represented by a gray cloud.	 *Figure 12. Gray cloud icon*
When a z/VM System is suspended, a pause icon is visible.	 *Figure 13. Pause icon*
When an item is locked, a padlock icon displays in the lower-left corner of the z/VM System icon.	 *Figure 14. Padlock icon*

Table 3. z/VM System icons (continued)

z/VM System status	z/VM System Icon
When there is a warning condition for a z/VM System, a yellow flag is displayed in the upper-right corner of the window. For example: • The page/spool usage exceeds the warning threshold that is defined in the IBM Wave Parameters. • There is an attention required object with a severity equal-to or greater-than 50.	 *Figure 15. Yellow flag icon*
When there is an error condition for a z/VM System, a red flag icon appears in the upper-right corner of the window. For example: • The page/spool usage exceeds the error threshold that you defined in the IBM Wave Parameters • There is an attention required object with a severity equal-to or greater-than 70.	 *Figure 16. Red flag icon*
When a z/VM System is being shut down, a down arrow appears in the upper-left corner of the window.	 *Figure 17. Yellow down arrow icon*
When the z/VM System's API server is unreachable or TCP/IP connectivity to the z/VM LPAR is down, a no-entry sign appears in the lower-right corner of the window.	 *Figure 18. No entry icon*

Enterprise Viewer

Use the **Enterprise Viewer** to interact with z/VM Guests across the enterprise.

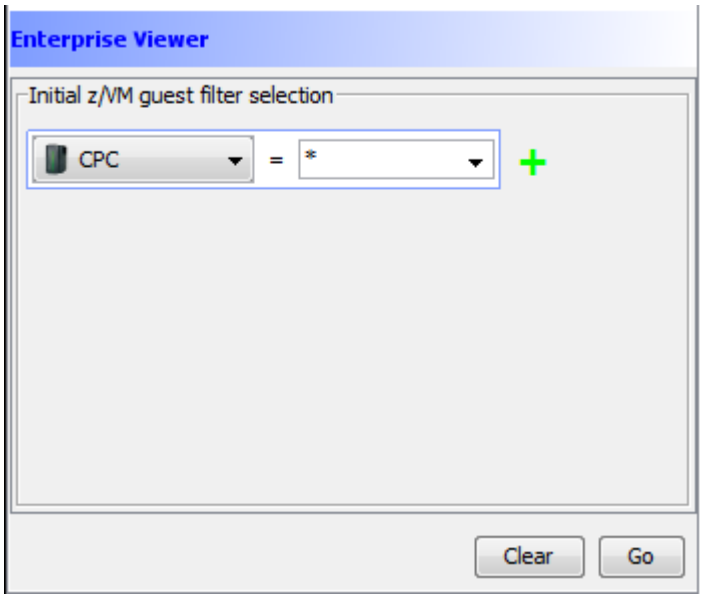

Figure 19. Enterprise viewer

Set an initial filter, and then click **Go** to populate the **Main Viewer** with the criteria that matches the filter. The guests are grouped by using one of the available "Group-By" options.

When you log out of IBM Wave, your filter settings are saved. The next time you log in, IBM Wave loads the last set of filters you used.

Dashboard Viewer

Use the **Dashboard Viewer** to see the **Enterprise Status Viewer**.

You can use the **Dashboard Viewer** to view performance-related information across the enterprise in a unified manner. After you set the **Filter selection**, the **Main Viewer** is populated with all of the managed z/VM Systems or CPCs that match the filter.

Your filter settings are persistent. The next time you log in, IBM Wave loads the last set of filters you used.

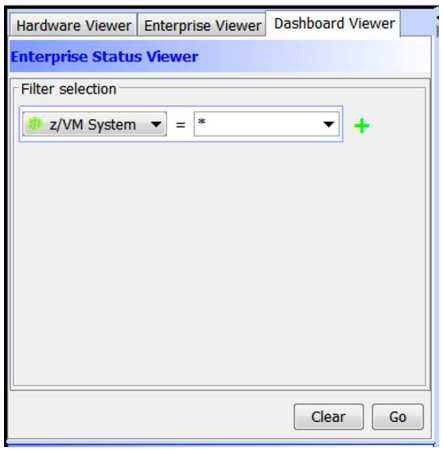

Figure 20. Dashboard Viewer: Filter selection

You can use the **Dashboard Viewer** to monitor all of the managed z/VM Systems across your enterprise. The z/VM Systems that are displayed in the **Dashboard Viewer** match the initial filter set in the **Enterprise Dashboard Viewer**.

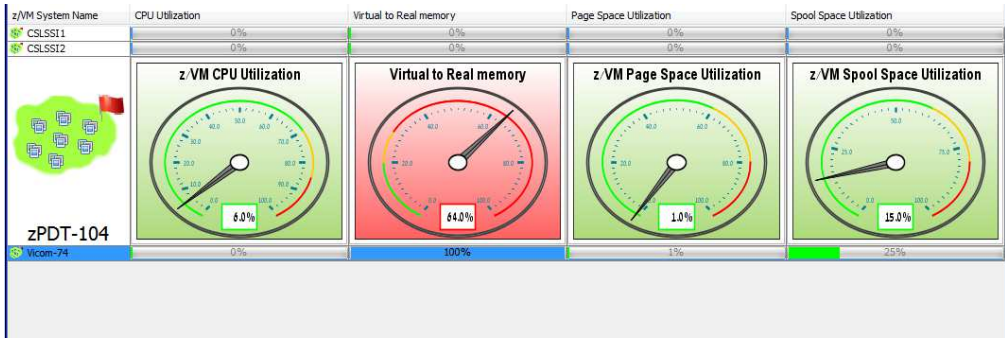

Figure 21. Main Viewer: Dashboard

The table contains following columns:

- **z/VM System Name** - The name of the z/VM System, along with an icon that depicts its status.
- **CPU Utilization** - The CPU utilization of the z/VM System.
- **Virtual to Real Memory** - The ratio between virtual memory that is used in the z/VM System to the actual physical memory allocated to the LPAR.
- **Page Space Utilization** - The utilization of the Page DASD devices that are defined to the z/VM System.
- **Spool Space Utilization** - The utilization of the Spool DASD devices that are defined to the z/VM System.

The table has two display modes:

- **Collapsed** - In the collapsed mode, all data is displayed as a bar graph. The color of the bar indicates the threshold of the specific element. You can define the thresholds in the **Administrative** > **Manage Parameters** > **Threshold tab**.
- **Expanded** - In the expanded mode, all data is displayed as Dial Plots that illustrate the utilization percent rate.

To switch from collapsed to expanded mode, double-click the corresponding line.

To increase or decrease the size, press the SHIFT key and scroll the mouse wheel.

- Increasing the size of the rows in the table makes the size of the Dial Charts larger. As the row size increases, fewer z/VM Systems are visible and scrolling might be necessary.
- Decreasing the size of the rows in the table allows more z/VM Systems to be visible without the need to scroll.

Note: The size of the rows cannot be decreased past the row size in the collapsed mode.

Property Viewer

The **Property Viewer** is a Regular Viewer that displays key properties of a selected object. When you select an object, a tab is opened in the **Property Viewer** table that shows you the key properties and the values. When no object is selected, the **Property Viewer** is empty.

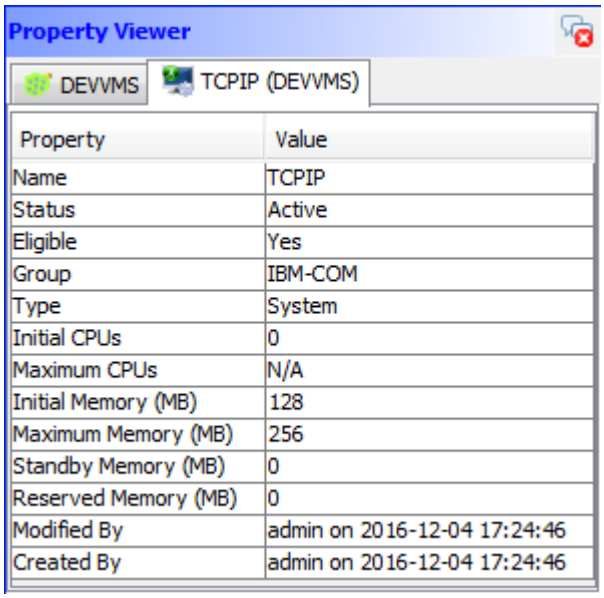

Figure 22. Property Viewer

Each property tab has a miniature icon that indicates the type of the object it describes. The icon and the properties are automatically updated when a property changes in the object.

Even when you switch viewers, properties are kept across viewers. For example, if a z/VM Virtual Server is selected from the z/VM Guest and Groups Viewer, and then a Virtual Network is selected from the Network Viewer, both objects' properties are displayed in the **Property Viewer**.

To make comparing objects in different z/VM Systems easier, properties are also kept across selected hardware elements.

Special tools: Property Viewer

Clear All

Allows you to clear all of the properties from the Property Viewer, except

for the tab with the current selected system or CPC. Using Clear All to remove the property tabs does not deselect an object.

Icons in this viewer

There are no icons in this viewer, except for the icons appearing in each property tab, which represent the object selected.

General Status Viewer

The **General Status Viewer** displays information about events, actions, and objects in the z/VM and IBM Wave environment. Many of the object and events apply to all systems. When an object or event depends on a specific z/VM System (other than the one you select), IBM Wave notifies you to switch to the correct system.

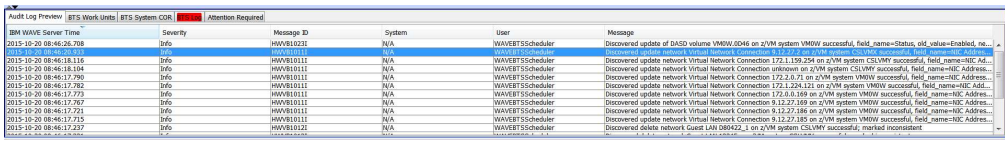

Figure 23. General Status Viewer

The **General Status Viewer** contains the following five minor viewers.

Audit Log Preview
The **Audit Log Preview** displays up to 100 of the most current audit log entries by time stamp. To sort the entries, double-click the **Audit Log Preview** tab to load the **Audit Log Display**. For more information, see "Audit Log Preview" on page 19.

For complete details about using the **Audit Log Display**, see the "Audit Log Feature" in *IBM Wave for z/VM: Administration and Customization*.

BTS Work Units Viewer
The **BTS Work Units** viewer displays the background task scheduler (BTS) Work Units that are submitted by you or other users when the work unit objects are within your scope. The table in the **BTS Work Units** is sortable. Click a table row to see the work unit details. For more information, see "BTS Work Units Viewer" on page 19.

BTS System COR Viewer
The **BTS System COR Viewer** displays the common output repository (COR) for the current system. Double-click the **BTS System COR Viewer** tab to load the **System COR Viewer**, and then view an entry by selecting the row. You can sort the entries by time stamp in both the "Available System COR Entries" table and the "System COR Details" table. The "System COR Details" table can also sort the Data header. For more information, see "BTS System COR Viewer" on page 21.

BTS Log Viewer
The **BTS Log** Viewer displays a subset of current tasks from the complete BTS Log. Double-click the **BTS Log** tab to load the **BTS Manager**, which contains a complete view of scheduled tasks and status. For more information, see "BTS Log Viewer" on page 22. For more information about the **BTS Manager**, see *IBM Wave for z/VM: Administration and Customization*.

Attention Required
The **Attention Required Viewer** displays all objects in all systems that are

defined to IBM Wave that require attention. When you double-click an entry in the **Attention Required Viewer** table, you are switched immediately to the object that requires your attention. If the attention is for a system other than the one you are viewing, you are notified to change to the system that requires attention. For more information, see "Attention Required Viewer" on page 23.

Audit Log Preview

The Audit Log Preview displays an overview of what you can find in the complete Audit Log Display.

The **Audit Log Preview** is a minor viewer. Click any row in the table to see the full message description.

By default, the **Audit Log Preview** previews up to 100 of the most recent messages that are issued to the **Audit Log Display**. To change the amount of preview messages, go to **Administrative** > **Manage Parameters** > **GUI**, and enter a number between 10 and 100 in the Audit Log Preview option.

The **Audit Log Preview** table contains the following information:

IBM WAVE Server Time
The time stamp for the message.

Severity
The message severity can be informational or error. The background color for the error message is red.

Message ID
The complete message identifier. Click the row to see the complete message.

System
The system to which the message applies.

User The user ID or the task, such as the BTS Scheduler, from which the message originates.

Double-click the **Audit Log Preview** tab to open the **Audit Log Display** window.

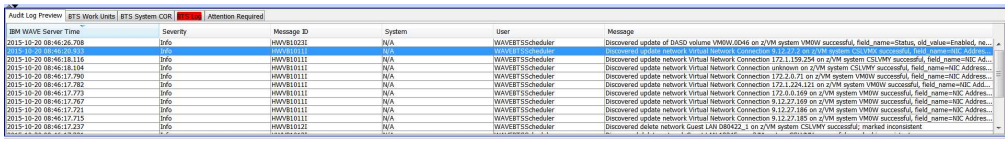

Figure 24. Audit Log Preview

The **Audit Log Display** tracks all of the auditable events in IBM Wave.

To open the complete **Audit Log Display**, you can either:

- Double-click the **Audit Log Preview** tab.
- Click **User Tasks > View Audit Log** from the **IBM Wave Main Menu**.

The **Audit Log Display** viewer can display all of the audit log entries from all systems. For more information, see the “Audit Log Feature” in .

BTS Work Units Viewer

The **BTS Work Units** viewer is a minor viewer that shows BTS work units that are submitted by IBM Wave users when the referenced objects are in your scope.

Work units that are submitted by IBM Wave API users and the background task scheduler do not appear in the **BTS Work Units** viewer.

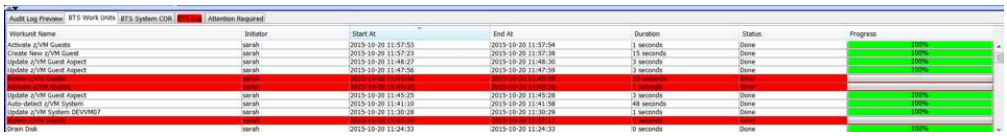

Figure 25. BTS Work Units viewer

From the **BTS Work Units** viewer, you can go to specific BTS work units to view the associated BTS requests and common output repository (COR) entries.

All columns in the **BTS Work Units** table are sortable. Click a column header to sort. To see the work unit details, click a specific row as shown in Figure 26. To delete a BTS work unit from the IBM Wave database, right-click, and select **Delete**.

Double-click the tab heading of the BTS Work Units to view and filter work units. This View opens a new window and is available from the **User Tasks** menu and the Interactive toolbar. The new window also includes work units that are originated by API, GUI, and CLI users for objects that are in your scope.

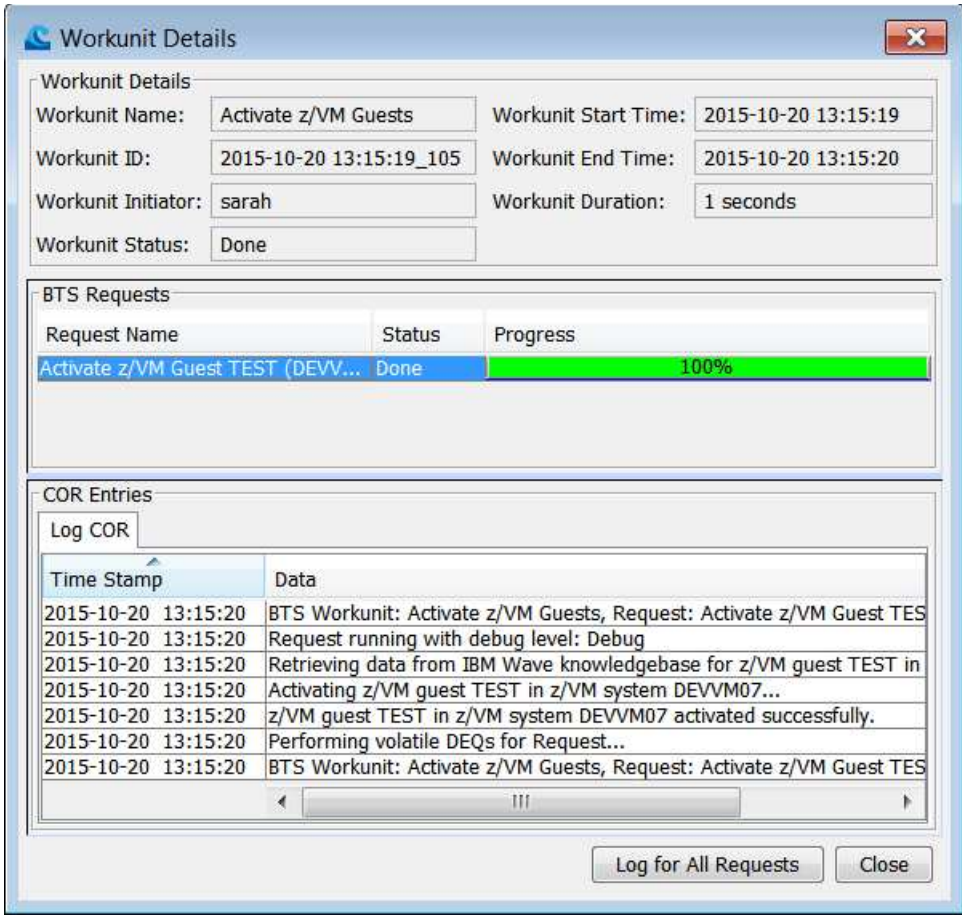

Figure 26. Work unit details

The BTS **Work Unit Details** viewer provides information about the currently selected BTS work unit. The viewer displays key properties of the BTS work unit and includes the name, start time, end time, ID, initiator, duration, status, and progress indication.

To populate the **Log COR** entry, select a **BTS Request**. To see the complete **Log COR** for the work unit, click **Log for All Requests**.

BTS System COR Viewer

Use the **BTS System COR** Viewer to view the common output repository (COR) entries.

The **BTS System COR** minor viewer that shows you the current system COR entry. The **BTS System COR** is a special type of COR entry that contains information about events that occurred within the IBM Wave System.

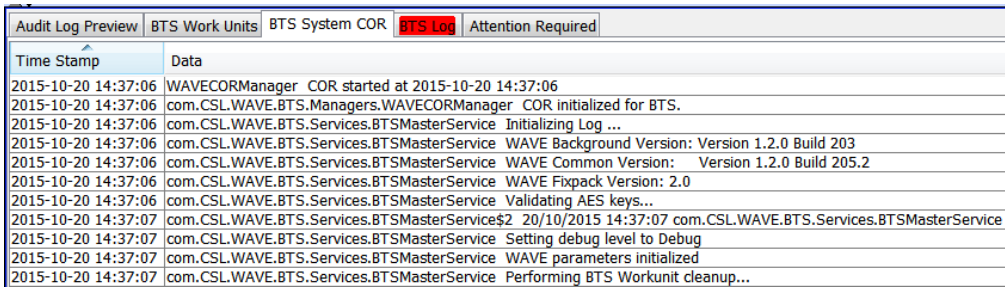

Figure 27. BTS System COR

To open the **System COR Viewer**, double-click the **BTS System COR** Viewer tab. You can use the **System COR Viewer** for the following tasks:

- View the active system COR entries which is highlighted green in Figure 28 on page 22.
- View the inactive system COR entries. You can preview the entry as shown in Figure 28 on page 22.
- Rotate the current system COR.
- Delete **BTS System COR** entries that are saved in the IBM Wave database.

You can sort each of the column headers in the **System COR Viewer**. In the "Available System COR Entries" table, you can sort by time stamp. In the "System COR Details" table, you can sort the "Time Stamp" or "Data" table header.

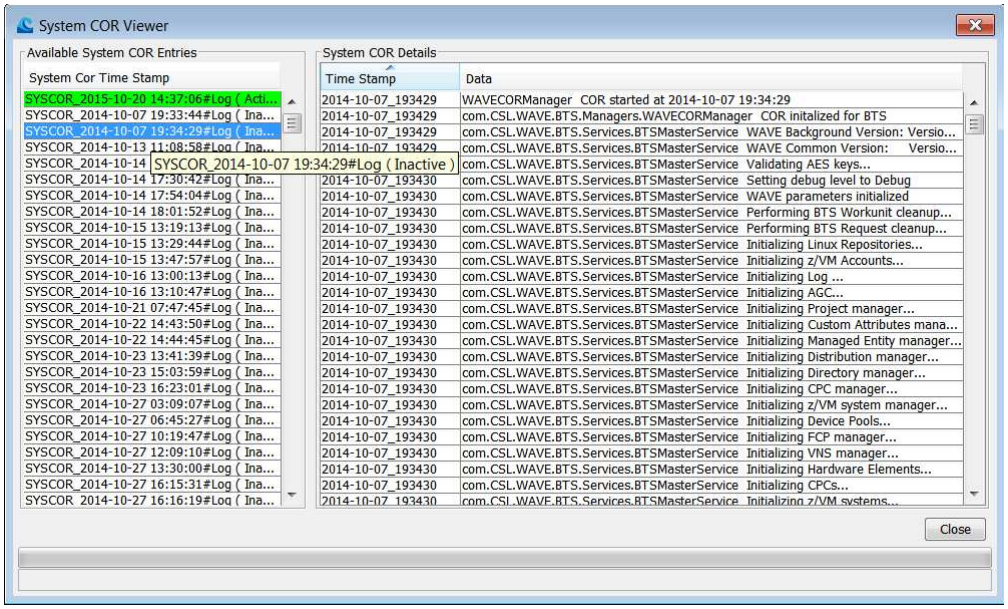

Figure 28. System COR Viewer

BTS Log Viewer

The **BTS Log** Viewer shows you a subset of scheduled tasks that are currently running in IBM Wave.

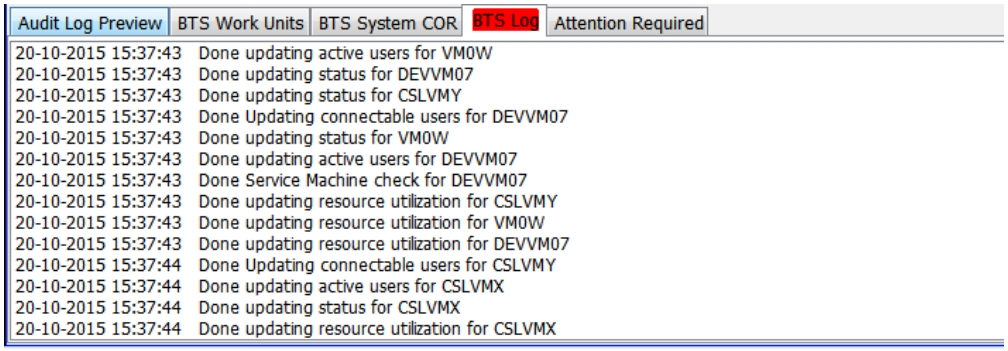

Figure 29. BTS Log Viewer

The complete background task scheduler (BTS) log is in the **BTS Manager** along with other tools to manage scheduled tasks in IBM Wave. To access the **BTS Manager**:

- From the **General Status Viewer**, double-click the **BTS Log** tab.
- From the **IBM Wave main menu**, select **Administrative** > **BTS Manager**.

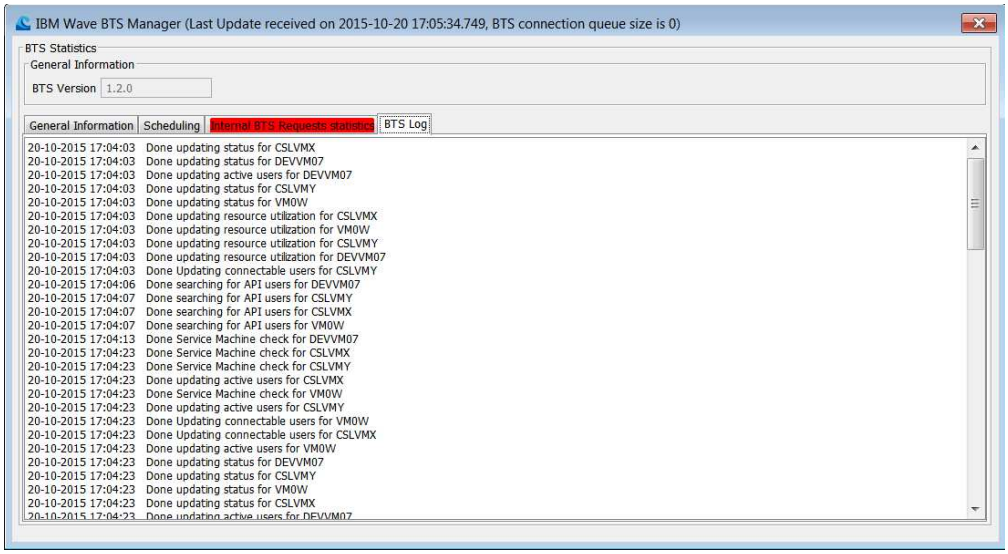

Figure 30. BTS Log in the BTS Manager

For complete information about the **BTS Manager**, see .

Attention Required Viewer

The **Attention Required Viewer** is a minor viewer that displays all the objects in IBM Wave that are marked "Attention Required". The **Attention Required Viewer** describes in plain text the reason why the object requires attention and the severity of the error.

The background color of the table entry reflects the severity of the entry:

- A yellow background indicates a low severity (less than 50).
- An orange background reflects a mild severity (greater than 50, but less than 80).
- A red background reflects a high severity (greater than 80).

The color of the attention entry is based on the user severity value assigned.

Figure 31. Attention Required Viewer

The original user severity is calculated internally by IBM Wave. The severity depends on the type of object and the type of error. By default, IBM Wave sets the severity to the original severity, and it cannot be modified by an IBM Wave user.

When you double-click an entry in the **Attention Required Viewer** table:

- When the attention required object is in the current viewer, you are switched to the object that requires your attention.
- When the attention required object is in a different viewer, you are asked to select which viewer to open. After you make your selection, you are switched to the object that requires your attention.

If the attention is for a system other than the one you are viewing, you receive a message to change to the system that requires attention.

Attention Required Viewer

You can filter the Attention Required table by using the “Filters” below the table. You can filter the type of objects that are displayed, and the z/VM System to which the objects belong. The filter is dynamic. It changes the table immediately. To disable all of the filters and return the **Attention Required Viewer** to its original value, press “Reset”.

Right-click to select one or more Attention Required table entries to change the severity, reset the severity, mark entries as “Ignored”, or reset “Ignored” entries to the original state. The ignored entries reappear in the table only when the “Ignore” filter is checked.

The title of the table, “The following objects...” displays the number of objects that match the selected filter versus the total number of objects that display when the filter is cleared. For example, in Figure 31 on page 23 (195/496 match filter and current z/VM system selected).

z/VM User Groups Viewer

The **z/VM User Groups Viewer** is the first minor diagram viewer in the **Current System Viewer**.

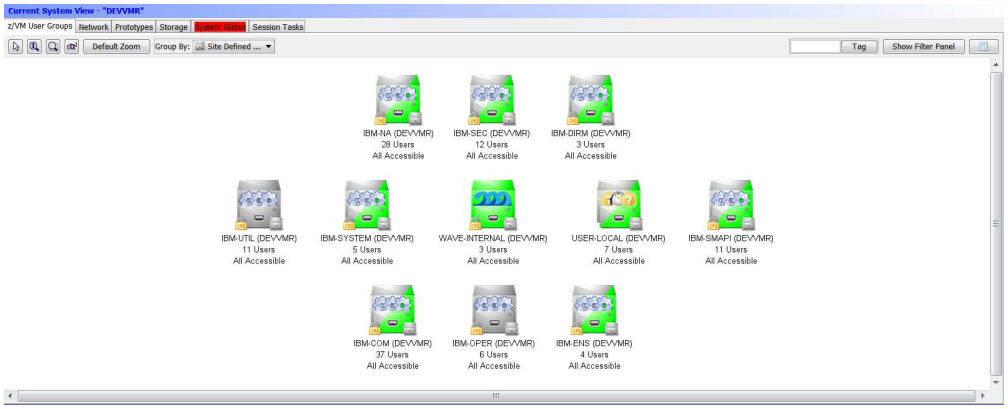

Figure 32. z/VM User Groups Viewer

The **z/VM User Groups Viewer** allows you to interact with z/VM Guests (guests that are in the IBM Wave User's Scope).

The guests are grouped together based on the value specified in the “Group-By” drop down list.

The z/VM Guests that display in the **z/VM User Groups Viewer** depend on whether the **Hardware Viewer** or **Enterprise Viewer** are used. When you use the **Hardware Viewer**, the viewer displays only the z/VM Guests belonging to the selected z/VM System. When you use the **Enterprise Viewer**, the viewer displays all of the z/VM Guests in all managed z/VM Systems that match the initial filter set in the **Enterprise Viewer**.

The amount of function available is based on the viewer you use. When you use the **Enterprise Viewer**, limited z/VM guest function is provided .

Working with z/VM User Groups

Working with z/VM User Groups explains the available options that are available to group guests in IBM Wave.

The **z/VM User Groups** Viewer displays z/VM Guests that are grouped based on the value that is selected in the**Group-By** menu. z/VM Guests can be grouped in the following categories:

- **CPC** Only available when you are using the Enterprise Viewer - z/VM Guests are grouped based on the CPC to which they belong.
- **z/VM System** Only available when you are using the Enterprise Viewer - z/VM Guests are grouped based on the z/VM System to which they belong.
- **SDG** - z/VM Guests are grouped based on the Site Defined Groups (SDG) to which they belong.
- **Project** - z/VM Guests are grouped based on the Project to which they belong.
- **OS Distribution** - z/VM Guests are grouped based on the operating system that is hosted on the z/VM Guest (Linux distributions, CMS, z/VM, or z/OS®).
- **z/VM Account** - z/VM Guests are grouped according to the account to which they are assigned.
- **Custom Attributes** - z/VM Guests are grouped based on possible values and default values of user-defined Custom Attributes.

  Note: Type of grouping is only available when **Custom Attributes** are defined.

You can expand or collapse a group by right-clicking on it and selecting the appropriate option from the menu. You can also double-click a group to expand it, and double-clicking the "-" sign in the upper left corner of an expanded group collapses it.

Special considerations for site defined groups

A user with sufficient privileges can add, update, and delete a Site Defined Group (SDG).

Certain SDG is locked. z/VM guests cannot be transferred in or out of the following groups:

- Default IBM Wave SDG - Default IBM Wave SDG names contain the "IBM-" prefix.
- IBM Wave internal SDG - IBM Wave internal SDG names are labeled "WAVE-INTERNAL".

Transferring z/VM guests between groups

With IBM Wave, you can transfer z/VM Guests from one group to another. The outcome of the transfer depends on which "Group By" value you select in the **Current System Viewer**.

To transfer one z/VM guests or a group of guests, select the guests to transfer, and press CTRL while dragging them to the new group. To signify a valid transfer, the group you are moving the guest or guests to is highlighted with a green border.

The following list of "Group By" values define the outcome you get when you transfer the guest or guests.

- **z/VM System** - When transferred, IBM Wave opens the Live Guest Relocation (LGR) window.

Note: The LGR window only opens if both the source and target z/VM Systems are running z/VM 6.2 and above and are members of the same SSI cluster. Otherwise, you get an error that explains the failure.

- **Site Defined Groups** - By default, you cannot transfer default IBM or Wave Internal SDG. Default SDG names are labeled "IBM-" and "WAVE-INTERNAL".
- **OS Distribution** - When transferred, the OS Distribution value of the z/VM Users are changed.
- **Project** - When transferred, the Project value of the z/VM guests are changed
- **z/VM Account** - When transferred, the primary account number for the z/VM guests are changed.
- **Custom Attributes** - When transferred, all z/VM guests are assigned the new value for the selected custom attribute.

Group By tool

Use the "Group By" tool to view the z/VM Users by site defined, operating system, project and custom attributes groups.

Use the "Group By" tool to view the z/VM Users in a selected z/VM System according to the following cross-sections:

- Site Defined Group (SDG)
- OS Distribution
- Project
- Custom Attribute

After the grouping is selected, the z/VM Users in the selected z/VM System are grouped by the selection criteria. When you use the Custom Attribute to "Group By", the groups that display reflect the possible values and default value for the selected custom attribute.

Note: The transfer action (CTRL + Drag) changes according to the "Group By" selection. For example, when the "Group By" selection is by Site Defined Group, by using the transfer (CTRL+ Drag) the guests are transferred from one group to another. However, when the "Group By" selection is a specific Custom Attribute, the transfer (CTRL+ Drag) assigns a new value for the custom attribute of the transferred guests.

z/VM User Group Icons

Group icons provide information about the type and status of z/VM User inside them. Each group is displayed as a file cabinet drawer. When the drawer is closed, the group is empty. When the drawer is open, the group has z/VM Guests inside it.

The background color of the group icon indicates the status of the z/VM User:

- A solid green background indicates that all the z/VM Guests inside the group are active.
- A solid gray background indicates that all the z/VM Guests inside that group are inactive.
- A half green and half gray background indicates a mix of active and inactive z/VM Guests inside the group.

Table 4. z/VM Group Icons

Group Icon	Description
	Empty group icon
	Active z/VM Guests icon.
	Inactive z/VM Guests icon.
	Mix of active and inactive z/VM Guests icon.

The content of the drawer indicates the types of z/VM Guests inside that group. A group can have a combination of the following icons:

Table 5. Icons within a drawer

Icon	Description
	IBM Wave Internal virtual server
	System Service Machine
	Linux Virtual Server
	CMS z/VM Guest
	OS running on this z/VM Guest is unknown

Table 5. Icons within a drawer (continued)

Icon	Description
	z/OS Virtual Server
	z/TPF Virtual Server
	z/VM Virtual Server
	z/VSE Virtual Server

For Example, in Figure 33, there is a combination of CMS, Linux, and system service machines with some active and inactive z/VM Guests inside the group.

Figure 33. Typical Group with icons

z/VM guest icons

The z/VM guest icons indicate the type and status of the z/VM Guest.

- A status indicator appears in the top-left corner of the icon (except for inconsistency icons).
- An inconsistency icon appears in the middle of the icon.
- A type indicator is located on the bottom-right of the icon.
- A warning or error indicator is located on the bottom-left of the icon.

Table 6. Status icons

Icon	Status	Description
	Active icon	The z/VM Guest is logged on in the z/VM System.
	Inactive status	The Linux inactive icon is gray.
	Starting Up icon	The Activate action has finished successfully.

Table 6. Status icons (continued)

Icon	Status	Description
	Shutting Down icon	The Deactivate action has finished successfully.
	Recycling icon	The Recycle action has finished successfully.
	Paused icon	The Pause action has finished successfully. The guest is in a CP READ state in which it does not consume any CPU and memory can be swapped out.
	Dormant Icon	When A Linux Guest has low CPU usage and the **Disable connectivity check for dormant guests** in the IBM Wave Parameters is checked, a guest will be marked as dormant.
	Clone Pending	The z/VM Guest is currently in the cloning process, but MINIDISK cloning has not yet started.
	Clone Source	The z/VM Guest is currently used as a source for cloning.
	Cloning indicator	This is indicated with the same icon as clone-pending, but with an added progress bar at the bottom of the icon to indicate the progress of the clone process. This indicates that the cloning process has started MINIDISK copy.
	Relocation Source indicator	This indicates that the z/VM Guest is currently relocating to a target z/VM System in the SSI cluster (outbound relocation).
	Relocation Target indicator	This indicates that the z/VM Guest is currently relocating from a source z/VM System in the SSI cluster (inbound relocation).
	Linux Installation	This indicates that an Interactive Linux Installation is in progress for this z/VM Guest.
	Inconsistent icon	The z/VM Guest was marked as inconsistent by the BTS.
	Unknown warning icon	IBM Wave is unaware of the status of the z/VM Guest.

Type icons

For z/VM Guest type icons, see the table on the previous page regarding the contents of the z/VM Guest Groups.

Warning and error icons

The following icons represent warnings and errors.

Table 7. Warning and error icons

Icon	Meaning	Description
	Clone error	A clone error is indicated with the same icon as clone-pending, but with an added red flag.
	Not connectable	A not connectable icon indicates the WAVESRV server does not have IP connectivity to the z/VM Guest.
	Not initialized for IBM Wave	The not initialized for IBM Wave icon indicates one of the following: • The specific user has not been inited for IBM Wave use yet. • The initialization for IBM Wave action was only partially completed.

Network Viewer

The **Network Viewer**, as shown in Figure 35 on page 32, is the second minor Diagram viewer in the **Current System Viewer**.

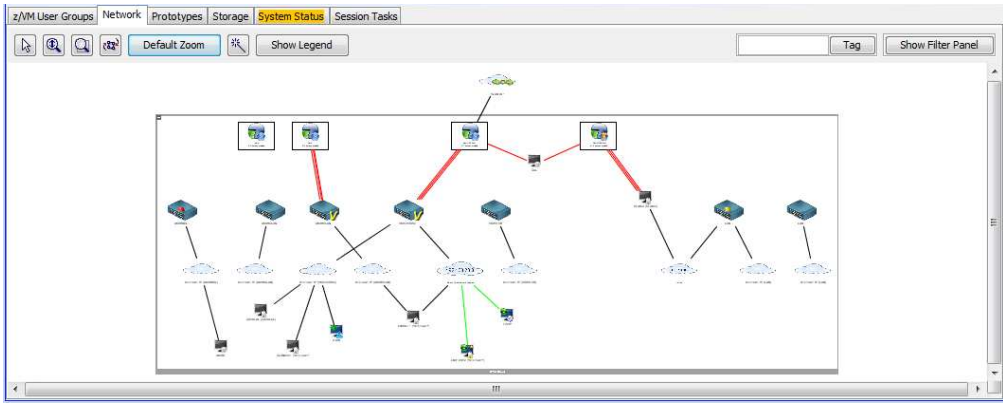

Figure 34. Network Viewer

Use the **Network Viewer** to interact with the network of the currently select z/VM System. From this view, you can create, update, and delete Virtual Networks (Guest LANs and VSwitches), and connect and disconnect z/VM Guests to Virtual Network Segments. Only Virtual Networks that are your scope are visible in this viewer. Also, if there are z/VM Guests that are connected to more than one Virtual Network, and one of the Virtual Networks is not in your scope, it is not shown in the **Network Viewer**.

The **Network Viewer** can also show connections outside the currently selected z/VM System (provided that IBM Wave has enough metadata information to present it). For example, Managed Entities that are associated with Device Pools and VNSs are present in the **Network Viewer**.

The **Network Viewer** is comprised of five layers, which each contain a type of IBM Wave managed object. The layers are as follows:

1. **The zVM Guest Layer** - This layer contains z/VM Guests that are connected To Virtual network Segments, or directly to Device Pools.
2. **The VNS Layer** - This layer contain Virtual network Segments that are connected to Virtual Networks and/or IBM Wave Manage Entities.
3. **The Virtual Network Layer** - This layer contain Virtual Networks (z/VM Guest LANs and VSwitches) that are connected to VNSs and may be connected to Device Pools (A VSwitch that is directly connected to an OSA real device is connected to a Device Pool that owns the real device).
4. **The Physical Layer** - This layer contains Device Pools visible to the IBM Wave User.
5. **The VLAN Layer** - This layer displays a representation of VLANs defined to IBM Wave. VLANs are defined to IBM Wave by virtue of the VNS to which they belong.
6. **The External Layer** - This layer contains IBM Wave Managed Entities which that connected to Device Pools, or VNSs.

Depending on the visible layers, IBM Wave draws connections between the layers, based on internal IBM Wave computations. For example, if a z/VM Guest A is connected to VNS 1, which is associated with a VSwitch V, which uses an OSA device owned by Device Pool D, and all layers are visible, IBM Wave will draw connections from the z/VM Guest to the VNS, from the VNS to the VSwitch and from the VSwitch to the Device Pool. However, if only the z/VM Guest Layer and the Physical Layer are visible, IBM Wave will draw a connection directly from the z/VM Guest to the Device Pool. The visibility of the layers is controlled by the IBM Wave User from the popup menu accessible by right clicking an empty space in the Network Viewer. IBM Wave provides three predefined layer combinations that can be manually turned on or off.

Network connections from z/VM Guests to other objects such as VNSs, Virtual Networks, Device Pools and others can be shown (depending on the NIC address configured for the connection). Each Virtual NIC defined to the z/VM Guest has its own “connector”.

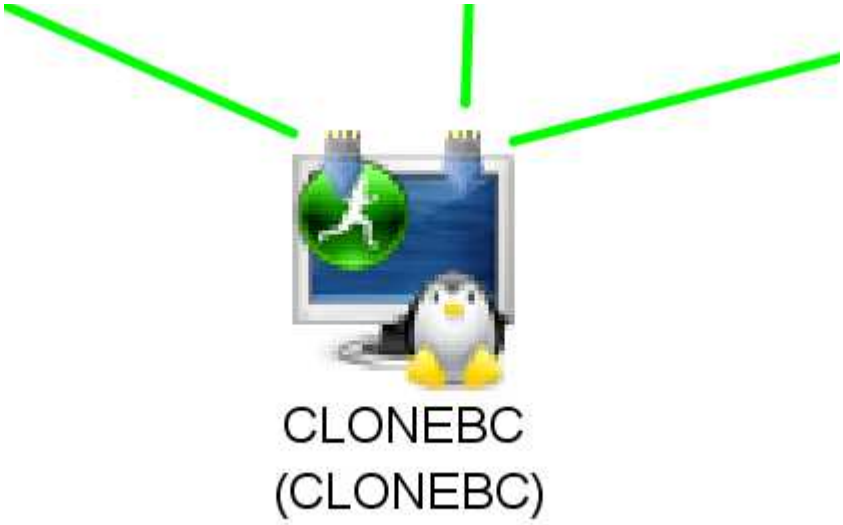

Figure 35. Network Connection example

In Figure 35, the z/VM Guest has two Virtual NICs defined. One virtual NIC connects to one VNS, and the other connects two different VNSs.

Note: IBM Wave automatically detects if multiple IP Addresses are defined on a NICDEF connecting the z/VM Guest to a Virtual Network. IBM Wave will display multiple connections between the z/VM Guest and the Virtual Network (or VNS, depending on the visible layers) in this case.

There are three predefined layer selections:

- **Logical View** - The logical view is comprised of the z/VM Guest Layer, VNS Layer, VLAN Layer, and External Layer.
- **Physical View** - The physical view is comprised of the z/VM Guest Layer, Virtual Network Layer, Physical Layer, and External Layer.
- **VLAN View** - This VLAN view is comprised of the z/VM Guest Layer and the VLAN Layer.

Note: The term Virtual Network is used to reference a LAN defined in the z/VM System (either a Guest LAN or a VSWITCH). It does not reference a VLAN (for example, a group of ports in a router or routers acting as one LAN) in the traditional sense.

Special tools available in this viewer

Magic Wand - The Magic Wand tool allows you to connect the following objects within the Network Viewer:

- A z/VM Guest to a Virtual Network Segment
- A Virtual Network to Segment to Virtual Network
- A Device Pool to an IBM Wave Managed Entity

When selecting the Magic Wand tool, the cursor changes according to the eligibility of the object it hovers over. Only objects that are eligible for connections turn the cursor to the Magic Wand. After clicking and dragging the Magic Wand, objects that are eligible for a connection as targets are highlighted in green when hovering over them (while dragging) with the cursor.

The object eligibility for the Magic Wand use is as follows:

- **z/VM Guests** - The z/VM Guest must be active. It must not in an inconsistent status, and not locked. The z/VM guest must be defined as a CMS or Linux Guest; it cannot be an internal IBM Wave z/VM Guest or an IBM Service Machine.
- **Virtual Network Segments** - The VNS cannot be an "Unknown IP" Segment.
- **Virtual Network** - The Virtual Network must not be in an inconsistent state and it must not be locked.
- **Device Pool** - The Device Pool must not be locked.
- **Managed Entity** - The Managed Entity type must be appropriate for the Device Pool type from which the Magic Wand originated. For example, it is impossible to connect a HIPERSOCKET Device Pool to a Router type Managed Entity.

To cancel a connection while it is being drawn, point to a white space in the diagram and release the cursor.

The **Show Legend** widget displays a window with a legend of the colors and shapes of connections in the diagram. For more information, see "Connection type indications" on page 35.

Icons in this viewer

This Viewer contains 3 types of objects:

- Virtual Networks
- z/VM Guests
- Connections

Each icon or connection provides information about the type and the status of the object it represents.

Virtual network icons

The following virtual network icons appear in IBM Wave for z/VM:

Figure 36. Guest LAN functioning properly

Figure 37. A VSwitch operating properly

The following warning and error indicators can also appear on the Virtual Network icon:

Table 8. Warning and error indicators on the Virtual Network icon

Icon	Meaning
	The Virtual Network is inconsistent, but there are still NICDEFs defined to it.

Table 8. Warning and error indicators on the Virtual Network icon (continued)

Icon	Meaning
	The Virtual Network is inconsistent.

Virtual network segment icons

If the VNS is not the "Unknown IP" VNS, the network defined to the VNS is displayed in the middle of the icon.

Figure 38. Virtual Network Segment icon

Device pool icons

Table 9. Device pool icons

Icon	Description
	An OSA Device Pool – This is a Device Pool which owns OSA devices.
	A HIPERSOCKET Device Pool – This is a Device Pool which owns HIPERSOCKET devices.

Managed entity icons

The icons in Table 10 represent entities that are managed by IBM Wave.

Table 10. Managed entity icons

Icon	Meaning
	Router Entity - This entity represents a Router.
	Storage Controller - This Entity represents a Storage Controller.
	z/OS System - This Entity represents z/OS System.
	z/VM System - This Entity represents a z/VM System.

VLAN icons

The virtual local area network (VLAN) ID is presented in the name of the VLAN.

Figure 39. Virtual local area network icon

z/VM guest icons

See "z/VM guest icons" on page 28.

Connection type indications

z/VM Guest to virtual network segment (VNS) connections are indicated by the color and solidity of the connector line:

Table 11. Device pool connector lines and descriptions

Color and solidity	Description
Figure 40. Solid black line	A permanent, inactive connection (The z/VM Guest connected by this connection is inactive).
Figure 41. Solid green line	A permanent, active connection (The z/VM Guest connected by this connection is active).
Figure 42. Dotted black line	A temporary connection, inactive.
Figure 43. Dotted green line	A temporary connection, active.
Figure 44. Solid gray line	A permanent connection, but IP address not detected.
Figure 45. Dotted black line	A connection that is being created (not finalized).
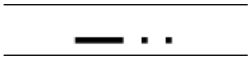 *Figure 46. Partial black line*	A Locked Connection.

Prototype Viewer

The term *Prototype* is used throughout IBM Wave for z/VM to signify either a DIRMAINT Prototype, or a z/VM: Secure Skeleton.

The **Prototype Viewer** is the third minor viewer in the **Current System Viewer**. Use the **Prototype Viewer** to interact with Directory Manager Prototypes. You can use the **Prototype Viewer** to clone from a selected prototype, edit the prototype's details, delete a prototype, and assign or remove z/VM Guests to and from a prototype.

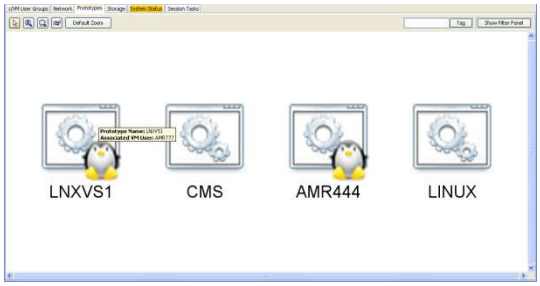

Figure 47. Prototype viewer

Note: The prototypes displayed for a specific IBM Wave User are prototypes whose project assignment or assigned z/VM Guest's project falls within the IBM Wave User's scope.

Special tools available in this viewer

There are no special tools available in this viewer.

Icons in this viewer

An Unassigned Prototype.

An Assigned Prototype with a CMS z/VM Guest assigned to it.

An Assigned Prototype with a Linux virtual server assigned to it.

If the prototype is inconsistent, a small red flag

will appear in the bottom left corner of the icon.

Storage Viewer

The **Storage Viewer** contains multiple viewers to help control storage allocation and distribution.

The **Storage Viewer** is the fourth minor viewer in the **Current System Viewer**. Use the **Storage Viewer** to interact with the storage that defined in the currently selected z/VM System.

The **Storage Viewer** contains three views:

- **Storage Distribution View** - Shows you an overview, in an interactive pie chart format, of the storage status for the z/VM System. For more information, see "Storage Distribution View."
- **Groups View** - Provides a way for you to interact with all of the storage groups that are defined for the z/VM System. For more information, see "Storage Group View" on page 38.
- **Volumes View** - Provides a way for you to interact with the storage volumes defined for the z/VM System. For more information, see "Volumes tab" on page 39.

Storage Distribution View

To view the z/VM **Storage Distribution** pane, select the z/VM system from the **Hardware Viewer**. From the **Current System Viewer**, click **Storage Distribution**.

The z/VM **Storage Distribution view** displays a pie chart that describes the storage utilization across the z/VM System. The graph displays the following types of storage usage:

- Enabled storage
- Used and Free Defined storage
- Used and Free Assigned storage

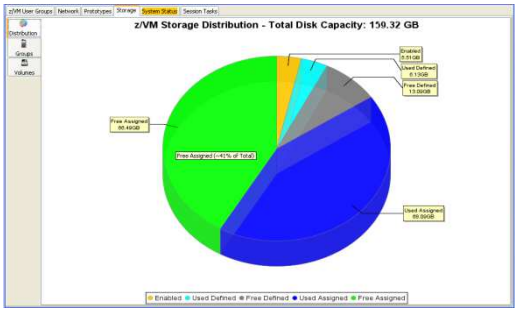

The title of the pie chart contains the total amount of storage that is available in the currently selected z/VM System.

Double-click on any section of the pie chart to display the **Volumes View** with a filter set to display only the selected volumes. For example, if you double-click the Green section (Free Assigned) the **Volumes View** displays only the volumes that are free and assigned.

To print the pie chart, right-click on it and select "Print...".

Special tools available in this viewer

There are no special tools available in this viewer.

Icons in this viewer

There are no special icons in this viewer.

Storage Group View

To view the storage **Group** tab, select the z/VM system in the **Hardware Viewer**. From the **Current System Viewer**, click **Storage** > **Group** tab.

Use the **Group View** to interact with the Directory Manager Storage Groups that are defined in the z/VM System. The **Group View** gives you a method to view the available storage groups, see the total capacity, and the ratio of free to used space.

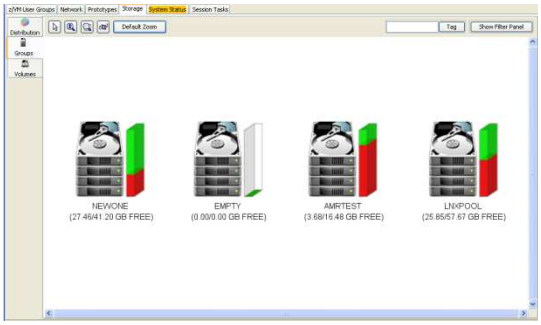

Figure 48. Storage viewer: Group View

In the **Group View**, you can create, update, and delete Directory Manager Storage Groups.

Double-click a group to display the DASD Volume View with an initial filter set to display only the volumes that are assigned to the group.

Note: The term "Directory Manager Storage Group" can refer to DASD Groups in DIRMAINT and subpools in z/VM: Secure.

Special tools available in this viewer

There are no special tools available in this viewer.

Icons in this viewer

The following icon represents a Storage Group.

Figure 49. Storage group icon

The bar to the right of the storage group icon represents the ratio of free to used space. Used space is marked in red. Free space is marked in green.

When a DASD Storage group is marked as inconsistent, a small red flag displays, as shown in Figure 50 on page 39.

Figure 50. Inconsistent flag icon

Volumes tab

To view the storage **Volumes** tab, select the z/VM system in the **Hardware Viewer**. From the **Current System Viewer**, click **Storage** > **Volumes** tab. To help you view more effectively, you can select the table view from the **Volumes** tab.

The **Volumes** tab in the **Storage Viewer** provides a way for you to interact with volumes in the currently selected z/VM system. Use the **Volumes** tab to assign volumes to storage groups, define and undefine regions on volumes, and more.

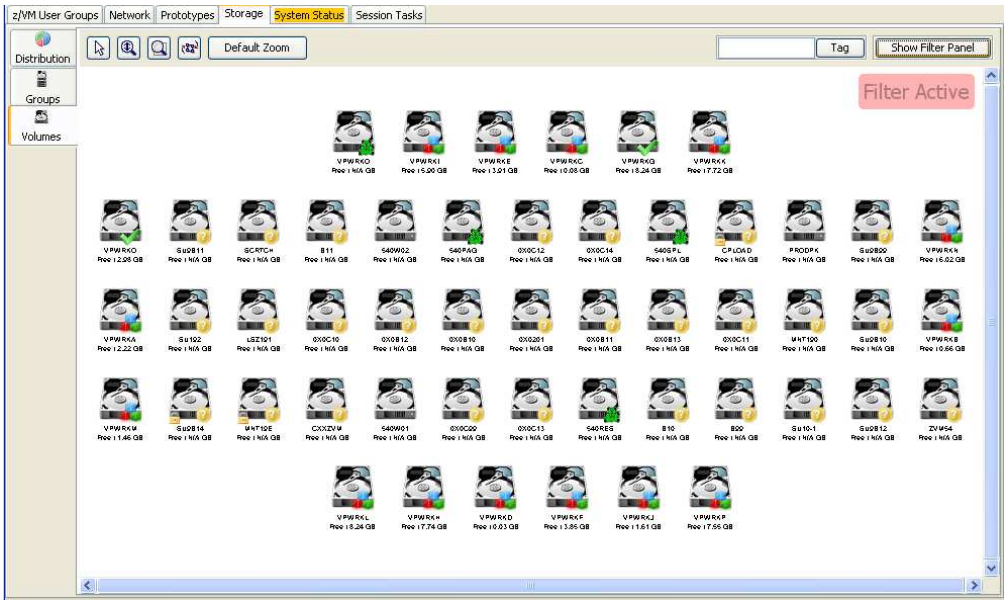

Figure 51. Storage Viewer: Volumes View tab

Compared to the other views in IBM Wave, the first time you use the **Volumes** tab is unique. You must first set a filter to view the DASD volumes. In other views, the objects that are relevant to the view and that are in your scope are presented. Because even a medium-sized installation can have 3,000 - 9,000 model 3 DASD volumes, displaying all of the volumes is inefficient.

When you first open the **Volumes** tab, a message displays requesting that you select an initial filter. To populate the view with the DASD volumes that match the filter:

1. Select an initial filter.
2. Click **Go** on the filter pane to display the view

To view all of the DASD volumes in the installation, use the general filter setting. For example, select "*" in the filter field "DASD Volume Name".

Remember: If your installation has more than 100 DASD volumes, the general filter setting is not advised.

Special tools: Volumes

There are no special tools available in the **Volumes View**.

Volume icons and indicators

Each icon in the **Volume** tab has a base icon with indicator status.

Table 12 displays the status indicators for icons in the **Storage Viewer > Volumes** tab.

Table 12. Volume status indicators

Volume Icons	Meaning
	Online volume indicator
	Offline volume indicator
	Defined volume indicator
	Assign volume indicator
	Spool volume indicator
	Page volume indicator
	Spool and page volume indicator
	CP-owned volume indicator
	Free volume indicator
	User-attached volume indicator
	Inconsistent volume indicator

Session Tasks Viewer

The **Session Tasks** viewer contain a table that lists all of the currently running, pending, or finished multiple task actions.

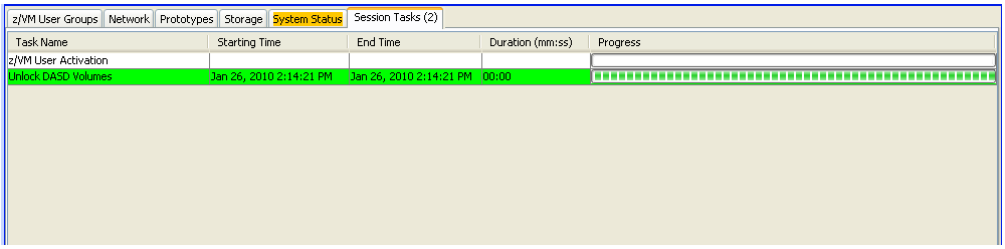

Figure 52. Session tasks viewer

Each entry in the table corresponds to hidden or visible multiple task action. The task name along with its start time, end time, duration, and progress indicator is displayed. Each row is color-coded according to the status of the task:

- **White** - The task did not start.
- **Orange** - The task is running.
- **Green** - The task is complete with no errors.
- **Red** - The task is complete with errors.

To expand a window that contains a multiple task action, double-click the corresponding line, or right-click on the line and select “Expand”.

Multiple task actions that are complete can be removed. To remove a multiple task action from the viewer, right-click and select the “Remove Process” option.

To remove a task that is in progress or did not start, expand the window and press “Cancel”. The cancel action is followed by several warning messages.

System Status Viewer

The System Status view is divided into two sections:

1. The z/VM System statistics, which includes the following options:
   - “Detailed performance panel” on page 42
   - “Detailed performance - Virtual Machines Utilization tab” on page 43
   - “Detailed virtual to real ratio pane” on page 44
   - “Detailed Page Utilization Statistics” on page 46
2. “Attention Required: Single System Viewer” on page 47

The tab header of the **System Status Viewer** reflects the color of the highest error or warning level in the viewer. For example, if the page utilization is in the error range, spool is in the warning range, and there are no “Attention Required” objects, the tab is red.

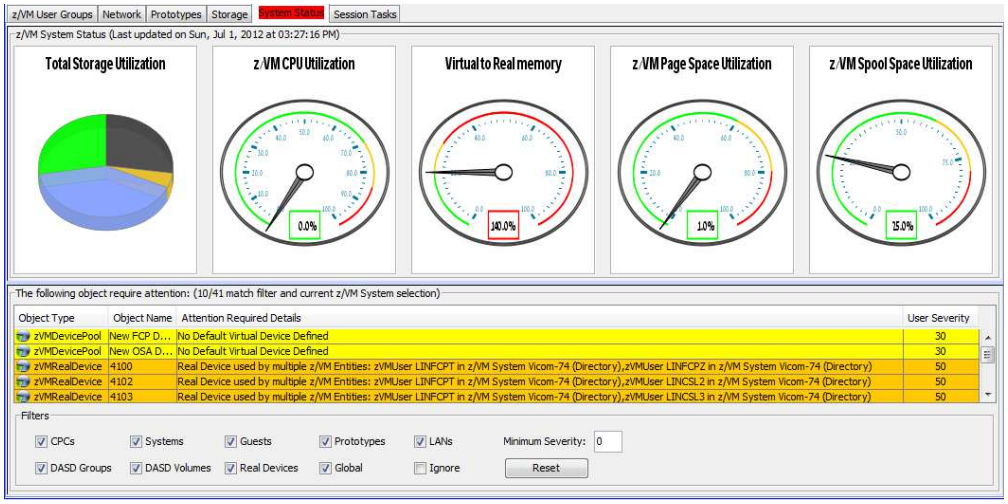

Figure 53. System Status Viewer

z/VM statistics

IBM Wave monitors the CPU utilization, the Virtual to Real Memory ratio, and the Page and Spool file usage in the z/VM LPAR. The monitoring is done by the Background Task Scheduler (BTS). The values are stored in the IBM Wave Database.

- The **Total Storage Utilization** pie chart contains the same information that is in the Storage Viewer.
- The CPU, Virtual to Real Memory and Page and Spool utilization dial charts show the current utilization of the CPU and Page and Spool files in the z/VM LPAR. The dial charts also indicate the warning and error thresholds defined in the IBM Wave Parameters.
  - The range between the warning threshold and error threshold is colored orange.
  - The range between the error threshold and 100% is colored red.

  The dial reflects the current utilization of the Page/Spool file and the Average CPU Utilization. The actual percentage is displayed in a field just below the dial itself. The color of the border of that field also reflects the threshold.

Double-click the dial plots to drill down into the detail panes.

Detailed performance panel

The detailed performance panel provides information about the current CPU and Paging utilization in the z/VM System. In a multi-CPU environment it is possible to view the utilization of each processor in a separate bar graph. The “Update Every” combo box allows to start a trend graph which will sample the system every 1,5, or 10 seconds and generate a scrolling graph.

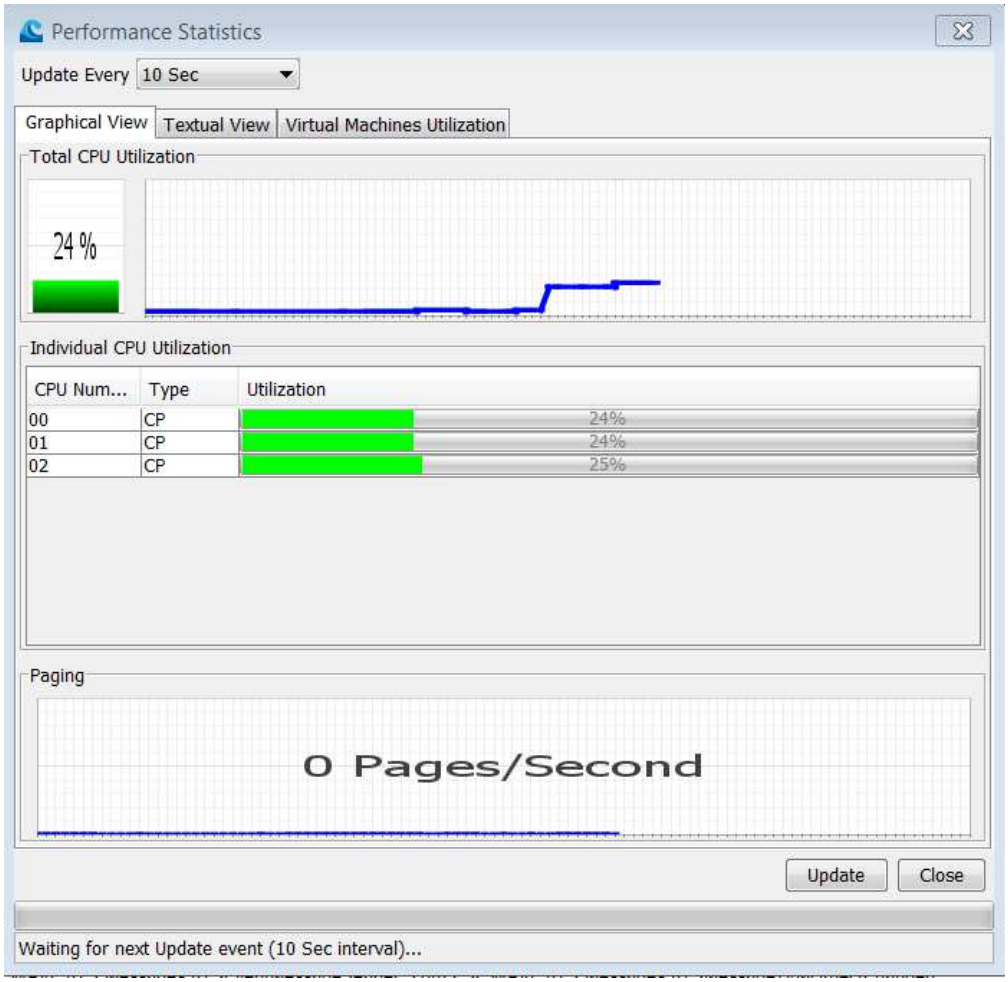

Detailed performance - Virtual Machines Utilization tab

The “Virtual Machines Utilization” tab shows detailed information about all z/VM Guests in the z/VM System. You can select a z/VM Guest and view data specific to its performance in the “User Detailed Information” pane, which is at the bottom of the window.

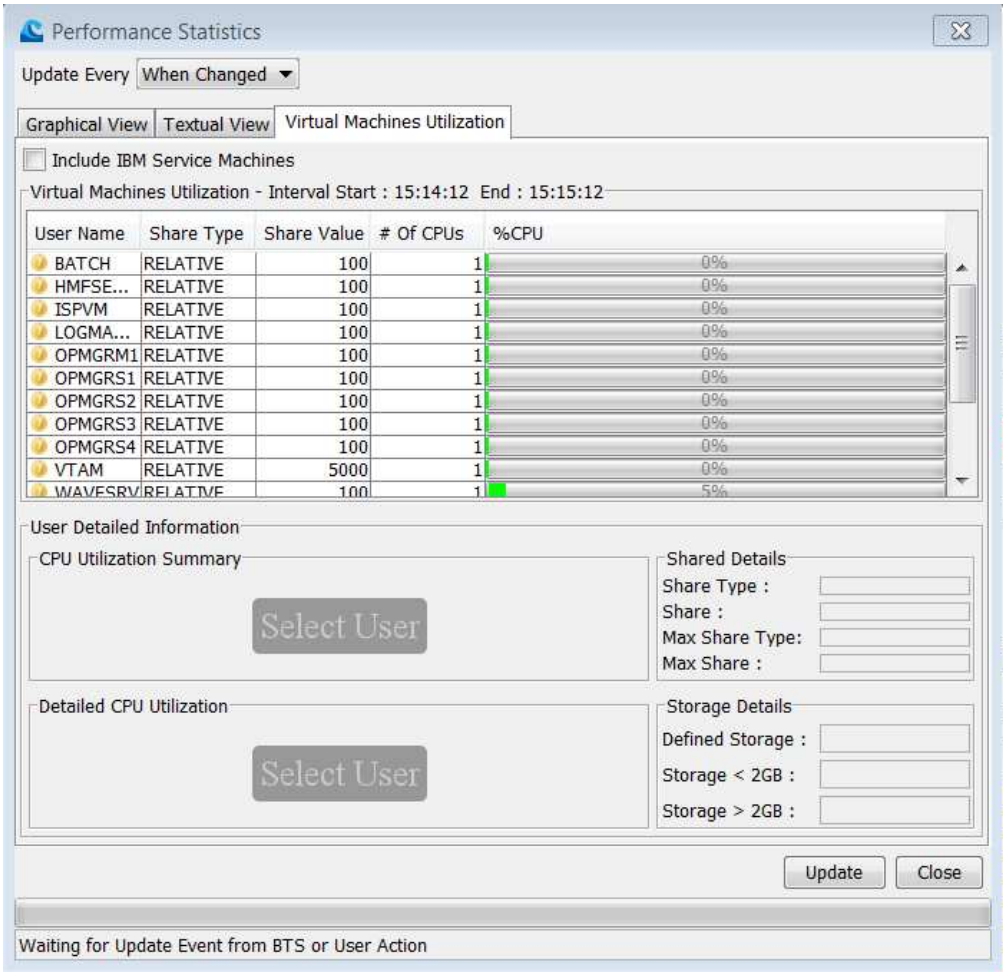

Figure 54. Virtual Machines Utilization tab

Notes:

1. The Virtual Machines Utilization tab is available only when the z/VM Performance Toolkit is used by IBM Wave to retrieve performance data for the z/VM System.
2. Detailed performance is shown only for z/VM Guests who are under the IBM Wave User's scope. The aggregated performance data for z/VM Guests that are not under the IBM Wave User's scope are displayed in the row marked “NOT IN SCOPE USERS”.

Detailed virtual to real ratio pane

The detailed **Virtual to Real Memory Utilization Statistics** window can be used to view detailed information about the virtual memory usage in the z/VM System.

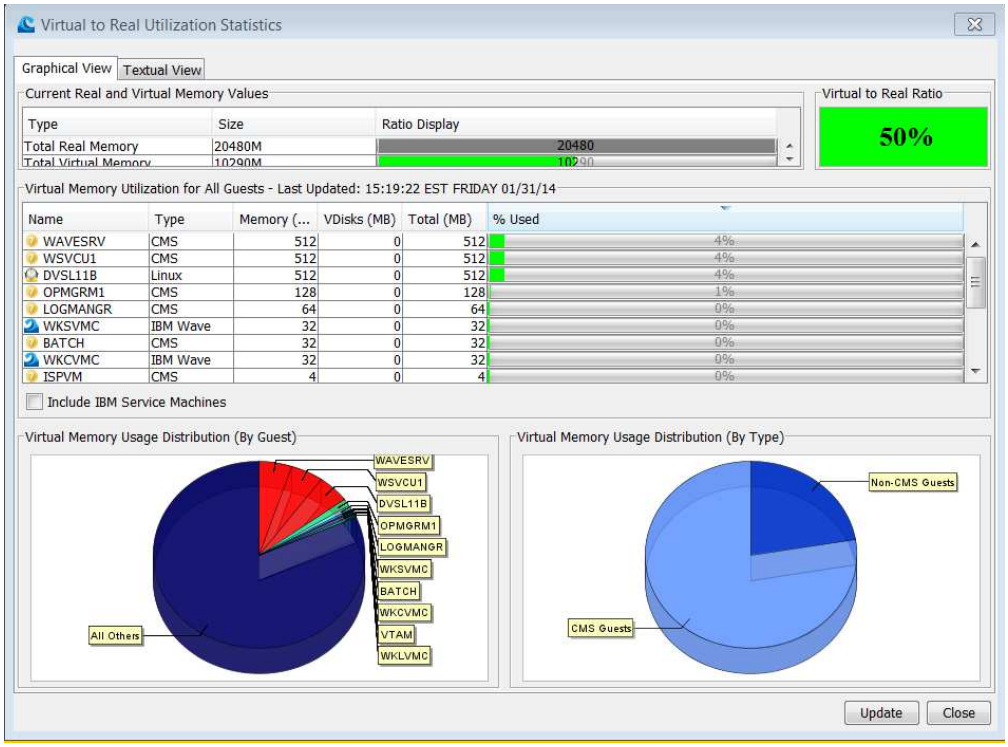

The “Current Real and Virtual Memory Values” pane displays the status of real and virtual memory and the ratio. The virtual memory bar and the ratio bar are color-coded based on the thresholds that are defined in the IBM Wave Parameters.

The “Virtual Memory Utilization for All Guests” pane contains a table with all logged-on guests and detailed information about the virtual storage for each one. The table contains six columns:

- **Name** - The name of the z/VM Guest.
- **Type** - Type of guest such as Linux, Service Machine, CMS, and others.
- **Memory** - The potential amount of virtual memory this guest can be assigned (value of the STOR= parameter in the output of the INDICATE command for the guest).
- **VDisks** - The amount of memory in megabytes which is allocated to VDisks for this z/VM Guest.
- **Total** - The total amount of potential virtual memory that can be allocated to the guest where this value is the sum of the memory field and the VDisks fields.
- **% Used** - The ratio between the total amount of potential virtual memory for this guest as part of the total potential virtual memory allocation in the system.

The “Virtual Memory Usage Distribution” pane located at the bottom contains two pie charts. The “By Guest ” pie chart located to the left shows the top 10 z/VM Guests who have the highest potential for an impact on virtual memory allocation in the system. This pie chart also contains an extra section for all other virtual guests.

The “By Type” pie chart located on the right shows the distribution of virtual memory by target usage and includes CMS guests, Non-CMS guests, VDisk memory.

Detailed Page Utilization Statistics

The **Detailed Page Utilization Statistics** window provides information about the utilization and occupancy of the DASD volumes that make up the Page Space.

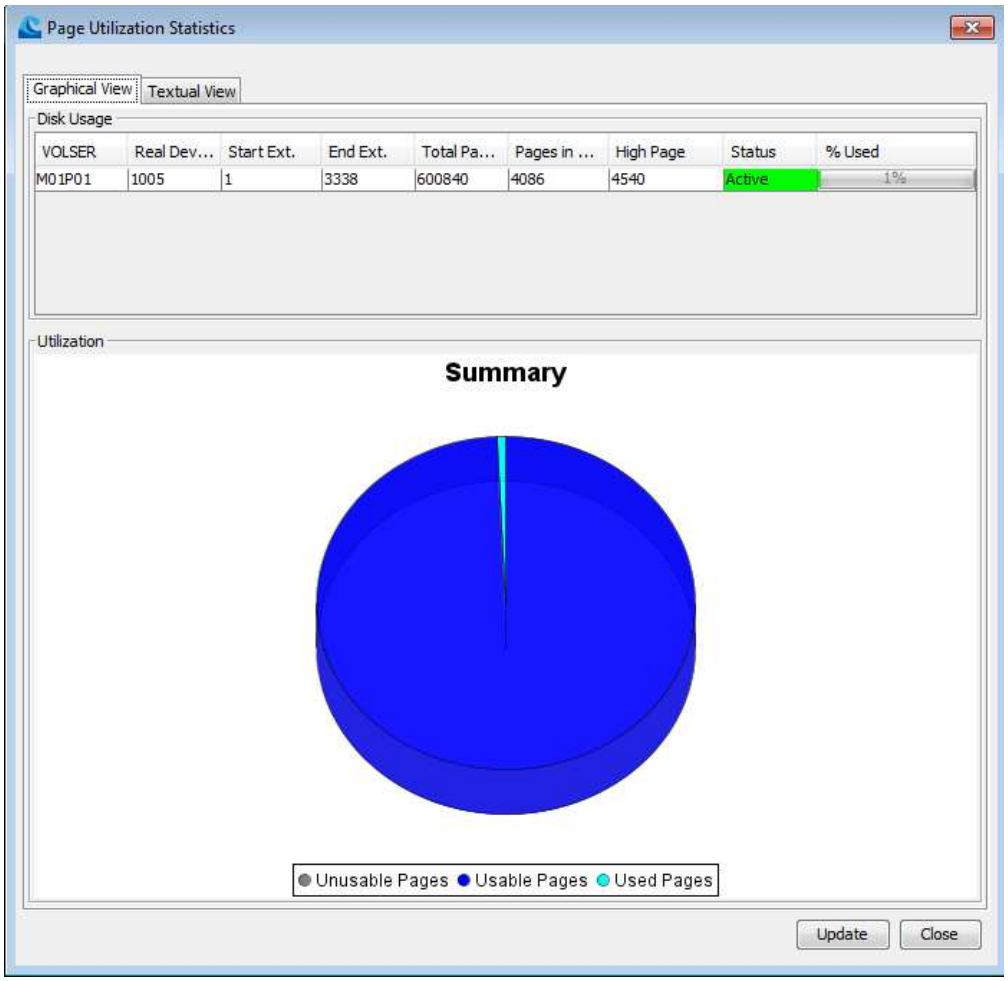

Figure 55. Page Utilization Statistics

Right-click on the table row provides the following actions

- **Drain** action drains a selected volume
- **Start** action starts a selected drained volume
- **Add Page** action adds a page volume to the z/VM system. For more information, see "Add Page and Add Spool" on page 65

A general pie chart is displayed in the "Utilization" pane to indicate the usable, unusable, and used page space. The "Textual View" pane shows the output of the query for page usage that is received from z/VM.

Note: If only one device is left, the option to drain is blocked since it can cause a z/VM outage.

Detailed Spool Utilization Statistics

The **Detailed Spool Utilization Statistics** window provides information about the utilization and occupancy of the DASD volumes that make up the Spool Space.

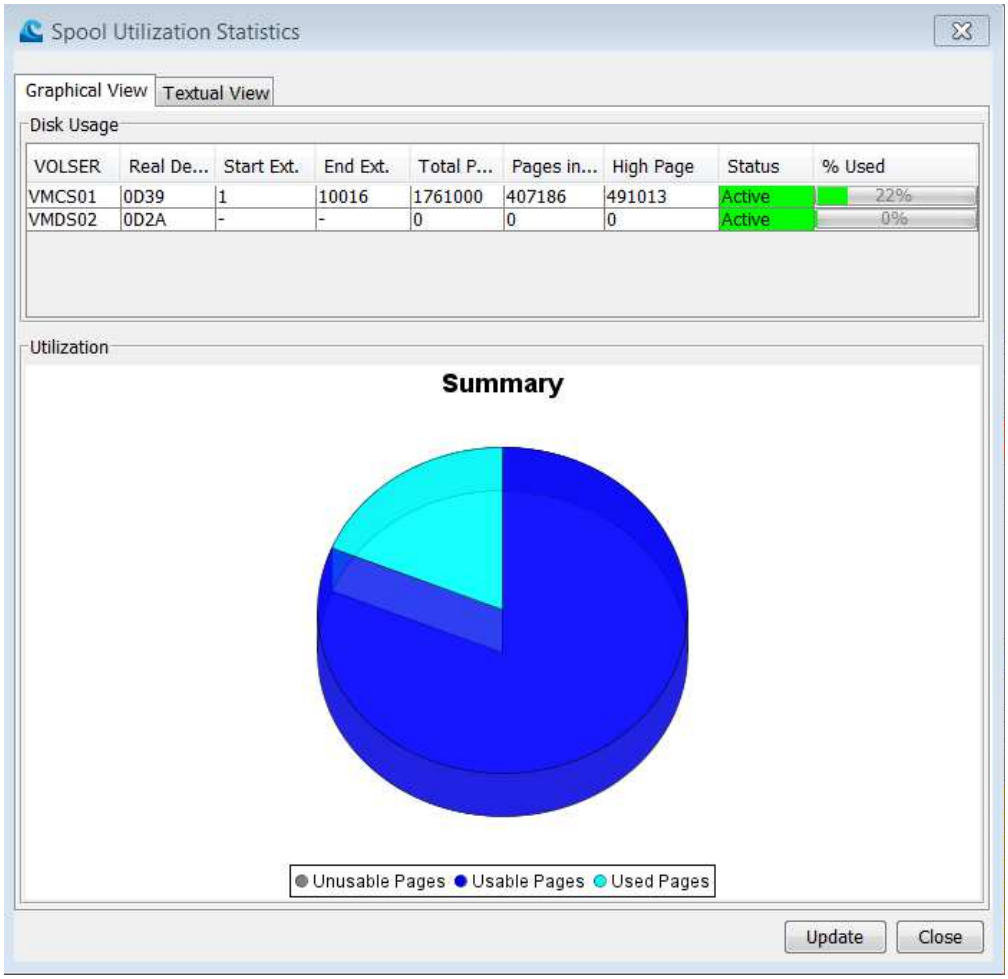

Figure 56. Spool Utilization Statistics

Right-click on the table row provides the following actions

- **Drain** action drains a selected volume
- **Start** action starts a selected drained volume
- **Add Spool** action adds a spool volume to the z/VM system. For more information, see "Add Page and Add Spool" on page 65

A general pie chart is displayed in the "Utilization" pane to indicate the usable, unusable, and used spool space. The "Textual View" pane shows the output of the query for spool usage that is received from z/VM.

Note: If only one device is left, the option to drain is blocked since it can cause a z/VM outage.

Attention Required: Single System Viewer

The second viewer in the **System Status Viewer** displays a system level version of the **Attention Required Viewer**.

The second section of the **System Status Viewer** is a single system version of the **Attention Required Viewer**. The **Attention Required** "Single System Viewer" displays only the objects that belong to the currently selected z/VM System.

Attention Required: Single System

Use the **Attention Required Viewer** to monitor objects that have an error or warning status such as inconsistent objects, guests that are not connectable, prototypes that are disassociated, and others.

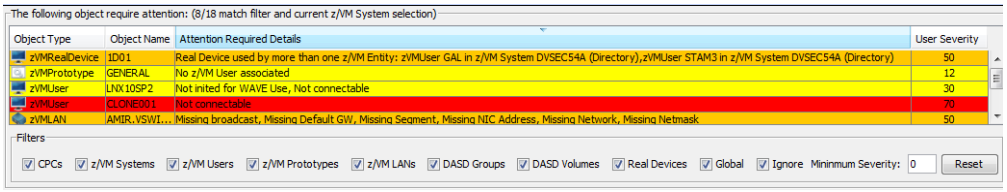

Figure 57. Attention Required: Single system

The table contains the following columns:

- **Object Type** - The type of object.
- **Object Name** - The name of the object.
- **Attention Required Details** - The text that defines the error or warning condition for which this object is flagged. If there are several error or warning conditions, they are listed on the same line separated by commas.
- **User severity** - The field indicates the user-defined severity of the error or warning. If there are several errors or warnings against the object, the severity column displays the highest severity.

Double-click an entry in the "Attention Required Details" table to switch views to the appropriate diagram view, which highlights the object with the error or warning. Right-click the object to take actions to resolve the error or warning condition. The system viewer is a subset of the **General Status Viewer** > **Attention Required Viewer**, all the same actions are available in this viewer.

For more information about the **General Status Viewer** > **Attention Required Viewer** actions, see "Attention Required Viewer" on page 23.

Chapter 2. User tasks and actions

The information in this section explains the user actions available in IBM Wave. It describes tasks and actions for objects that you can manage. Where possible, actions are sorted in alphabetical order. In some instances, such as the **User Tasks** menu, actions are sorted in the order in which they appear in the IBM Wave interface.

Note: All actions in IBM Wave depend on the scopes in permissions your administrator defines.

The following topics are covered:

- "User and Administrative Tasks" - General actions and user tasks.
- "Hardware functions" on page 58 - Actions and tasks for an IBM z Systems CPC and z/VM systems management.
- "z/VM User Groups functions" on page 68 - Actions and tasks for Site Defined Groups (SDG) and general grouping details.
- "z/VM guest and virtual server functions" on page 78 - Actions and tasks for z/VM and Linux user management.
- "Virtual networks functions" on page 132 - Actions and tasks for Virtual Network management, connections, and more.
- "Prototype functions" on page 139 - Actions and tasks for prototype management.
- "Storage functions" on page 150 - Actions and tasks for storage management.
- "Virtual network segment functions" on page 168 - Actions and tasks for Virtual Network Segments (VNS).
- "External entities functions" on page 171 - Actions and tasks for IBM Wave managed entities.
- "Device pool functions" on page 173 - Actions and tasks for device pools.

User and Administrative Tasks

IBM Wave for z/VM contains administrative and user actions and tasks. The tasks that you can perform are dependent on your scopes and permissions.

Administrative actions

You can access the administrative menu from the **IBM Wave Main Menu**. Click **Administrative**.

Note: IBM Wave requires American English (AMENG). No other languages are supported.

For information about the **Administrative** actions, see .

User Tasks

You can access the **User Tasks** from the **IBM Wave Main Menu**. Click **User Tasks**.

© Copyright IBM Corp. 2007, 2018

View Audit Log

To view the **Audit Log Display**, from the **IBM Wave Main Menu**, click **User Tasks > View Audit Log**.

For more information, see "Audit Log Preview" on page 19. For complete details about using the **Audit Log Display**, see the "Audit Log Feature" in .

Work Units Viewer

To view the **Work Units Viewer**, from the **IBM Wave Main Menu**, click **User Tasks > Work Units Viewer**.

The **Work Units Viewer** views work units that are created by GUI, CLI, and API users for objects that are in your scope.

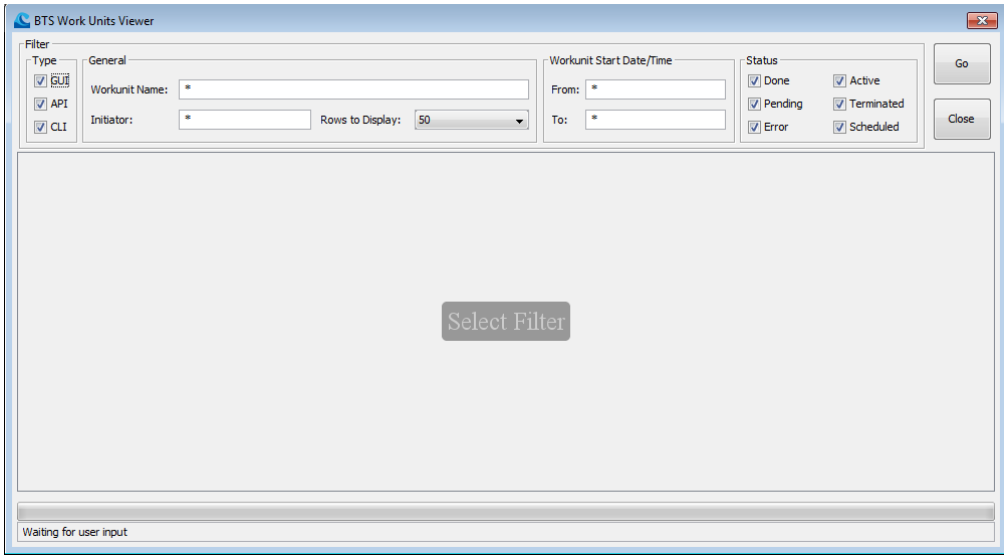

Figure 58. Work Units Viewer

The Viewer Allows you to filter the list of work units by the following criteria

- The type of user who initiated the work unit.
- The work unit name.
- The initiator of the work unit.
- The maximum number of rows to display.
- The work unit start date and time.
- The status of the work unit.

In order to view work units, set the filters in the panel on the top of the window and press the **Go** Button.

The **Close** button closes the **Work Units Viewer** window.

Change IBM Wave User Password

To change your password, use the **Change IBM Wave User Password** task.

The **Change IBM Wave User Password** task permits a regular IBM Wave User to change their password.

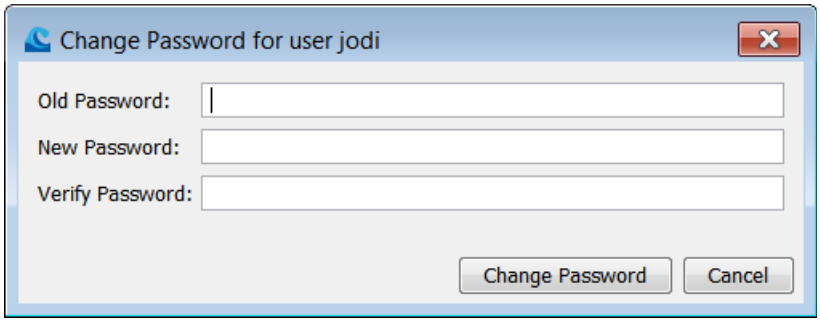

Figure 59. Change IBM Wave User Password

- **Old Password** - Enter the current IBM Wave User password for the IBM Wave User.
- **New Password** - Enter the new password.
- **Verify Password** - Enter the new password again for confirmation.

Change IBM Wave User Preferences

To change your preferences, use the **Change IBM Wave User Preferences** task.

Click **User Tasks** > **Change IBM Wave User Preferences** to change your preferences for IBM Wave.

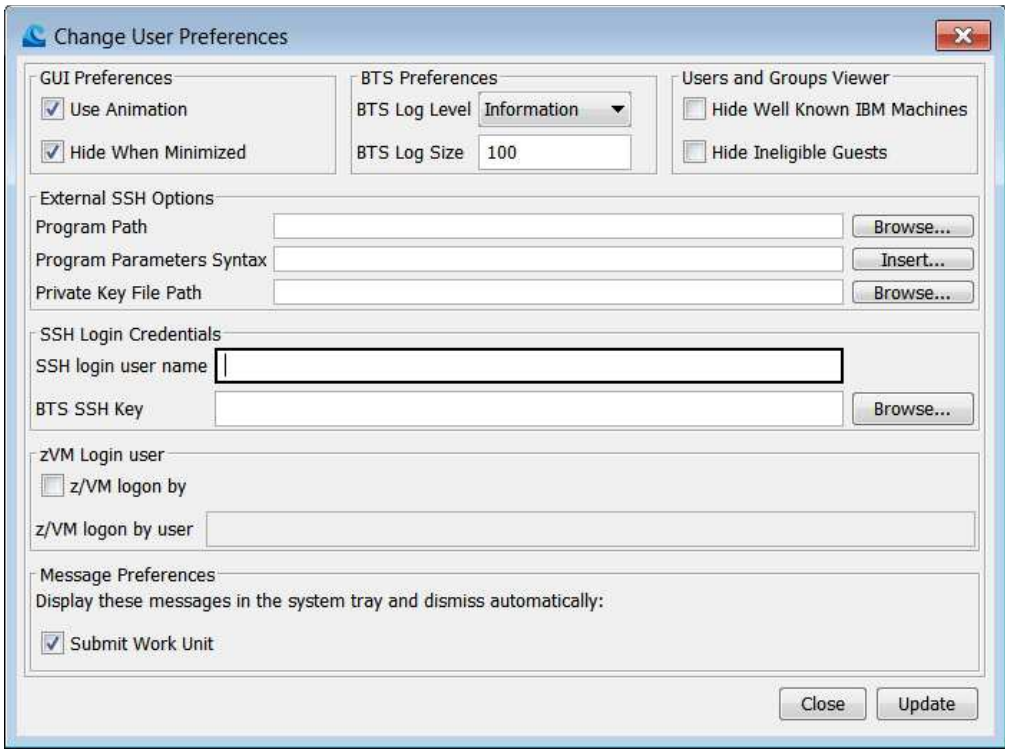

Figure 60. Change IBM Wave User Preferences

For complete information about the fields, see “Change User Preferences” in .

Display Effective Scope/Permissions for Session

To view your scopes and permissions for the session, use the **Display Effective Scope/Permissions for Session** task.

The **Display Effective Scope/Permissions for Session** task opens a window that displays the current scope and permission sets that are in effect for the current session. When LDAP, group-based scope/permission is used, the window also displays the LDAP Groups that match the user profiles.

If IBM Wave finds any conflicts during the merging of various profiles, a warning message is displayed, and you can view the permissions that are in conflict.

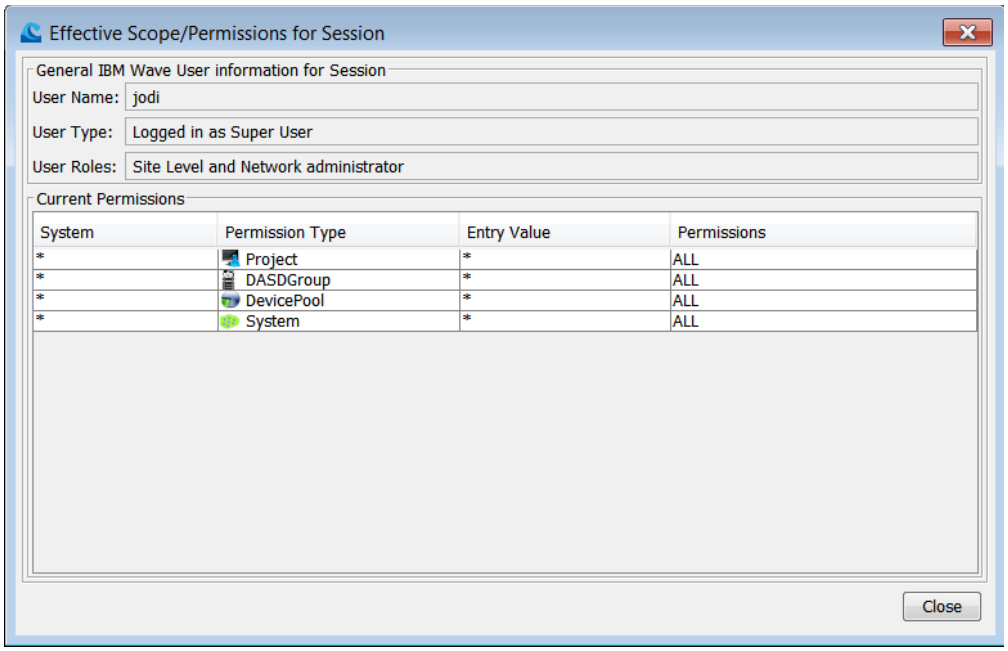

Figure 61. Effective Scope/Permissions for Session

For information about scopes and permissions, see “User Management” in .

IBM Wave User Script Manager

For complete information about the IBM Wave Script Manager, see “IBM Wave User Script Manager” on page 225.

Display projects

This option displays a list of the projects available to this IBM Wave User (Projects that are in the scope of the IBM Wave User).

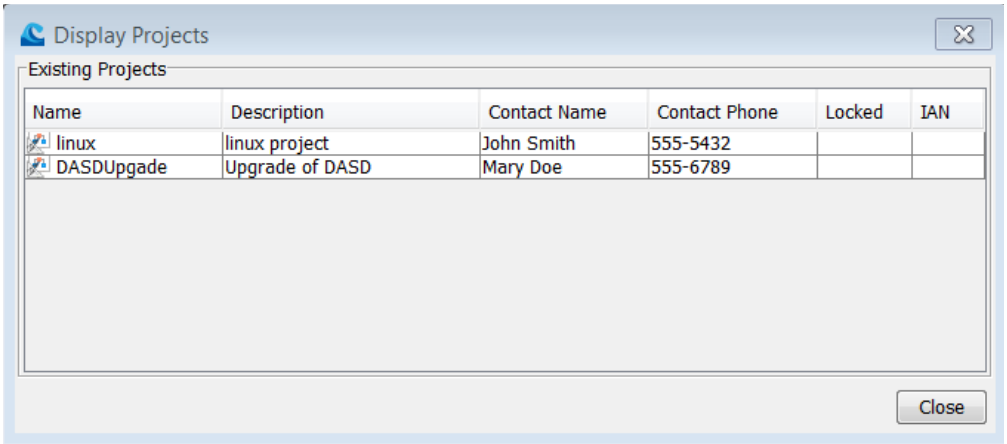

Display device pools

To view the existing device pools that are defined in your environment, display the device pools.

To display the available device pools, click **User Tasks** > **Display Device Pools**. The **Display Device Pools** window is a display-only version of the **Device Pool Manager**.

Note: Only device pools that are in your scope are displayed.

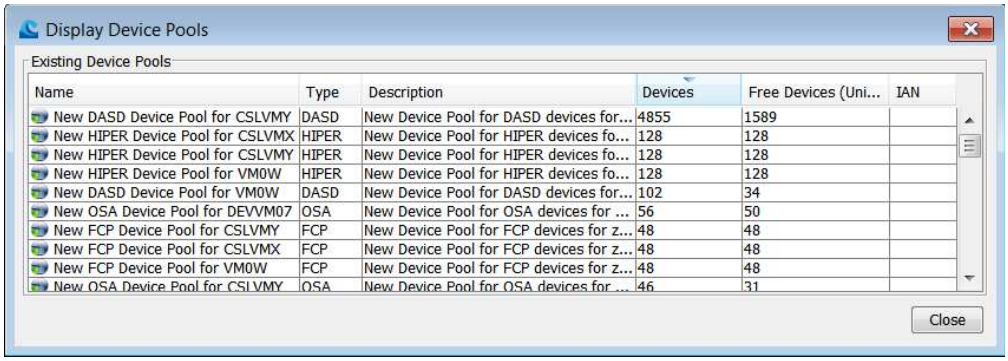

Figure 62. Display Device Pools

To access the **Device Pool Manager**, from the **IBM Wave Main Menu** click **Administrative** > **Device Pool Manager**.

- For more information about device pools, see the section about "Device pool functions" on page 173.
- For information about the **Device Pool Manager**, see "Manage Device Pools" in .

Display managed entities

The **User Tasks** > **Display IBM Wave Managed Entities** menu option displays all managed entities that are defined to IBM Wave. Because managed entities are not limited by scopes and permissions, all defined entities are shown.

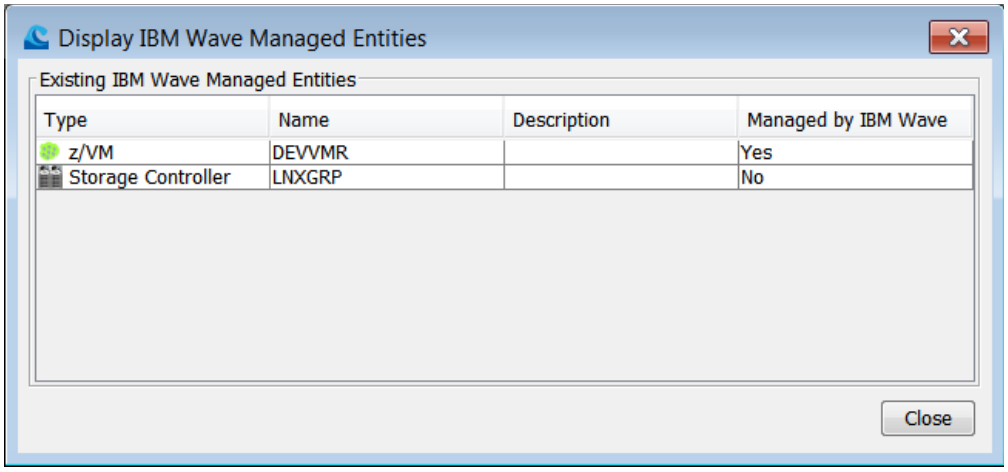

Figure 63. Display IBM Wave Managed Entities

The **User Tasks** > **Display IBM Wave Managed Entities** is a display-only version of the **Administrative** > **Site Management** > **External Entities Manager**. For information about the **External Entities Manager**, see

Display virtual network segments

This option displays all Defined Virtual Network Segments (VNS). VNS are not limited by Scopes and Permissions and therefore all defined VNS are displayed. The window is a display-only version of the Virtual Network Segment Manager, accessible from the "Administrative" menu.

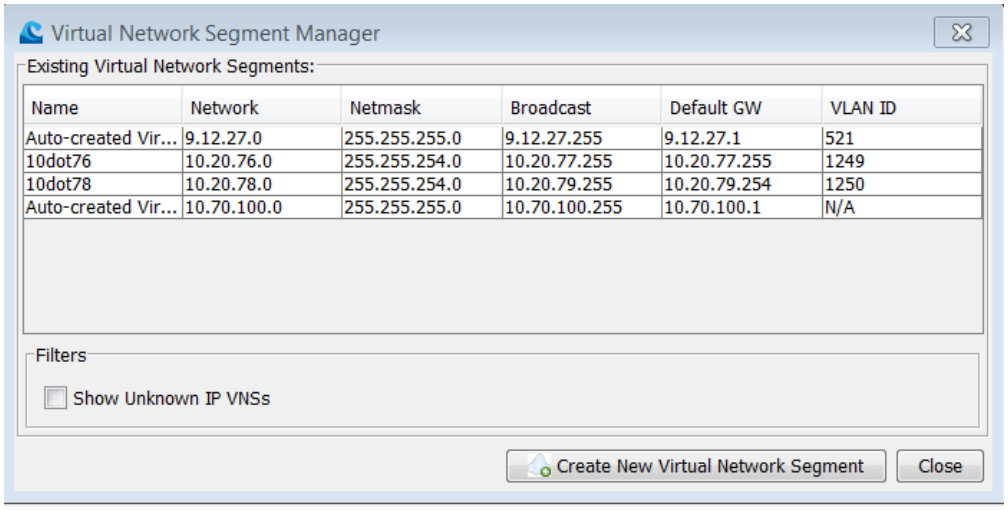

For more information about the Virtual Network Manager, see "Manage IBM Wave Virtual Network Segments" in *IBM Wave for z/VM: Administration and Customization*.

Display IBM Wave Linux repositories

The **User Tasks** > **Display IBM Wave Linux repositories** displays all of the configured IBM Wave Linux repositories.

For more information about the IBM Wave Linux Repositories, see the corresponding section in .

Display IBM Wave messages

To read a messages sent from a Linux guest, open **User Tasks** > **Display IBM Wave Messages**. The **View IBM Wave Messages** window opens.

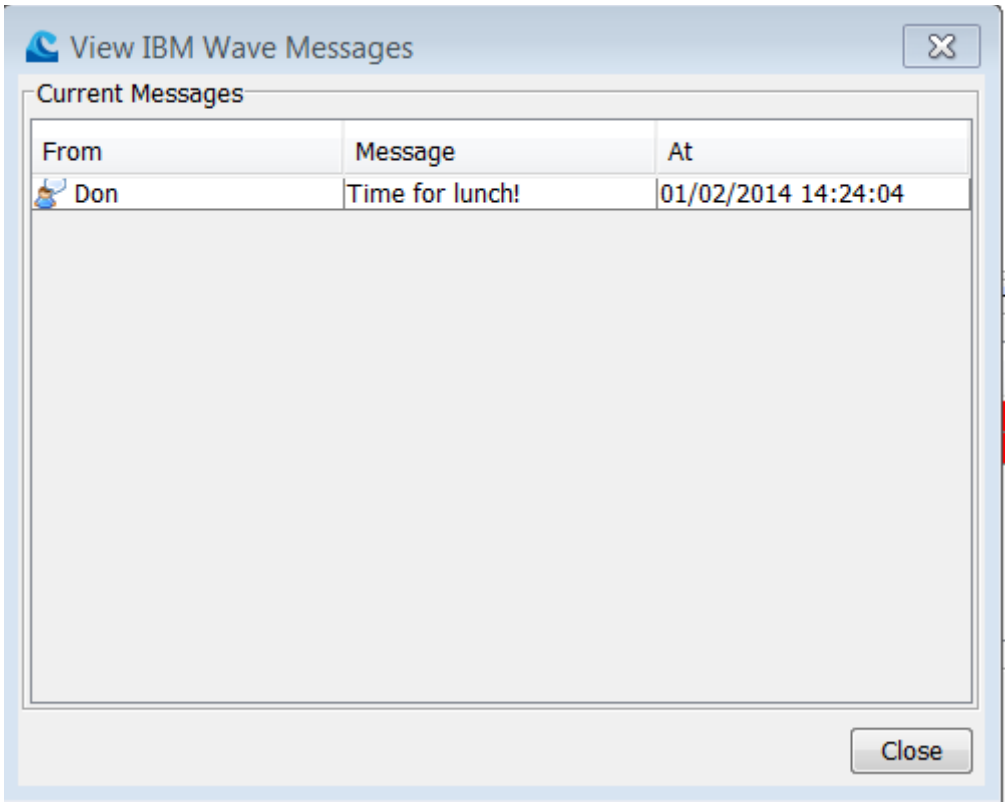

Figure 64. Display IBM Wave Messages

For more information about the IBM Wave messaging mechanism, see *IBM Wave for z/VM: Administration and Customization*.

Init Users for IBM Wave use

Use **Init Users for IBM Wave use** to configure Linux guests to be managed by IBM Wave. The following information covers both the **Init Users for IBM Wave use** interface and the process.

To manage Linux virtual servers with IBM Wave, you must use the **Init Users for IBM Wave use** task. **Init Users for IBM Wave use** applies only to active and connectable Linux servers.

To access **Init Users for IBM Wave use**, from the **IBM Wave Main Menu**, select **User Tasks** > **Init Users for IBM Wave use**. The **Initialize z/VM Users for IBM Wave use** window opens (as shown inFigure 65 on page 56).

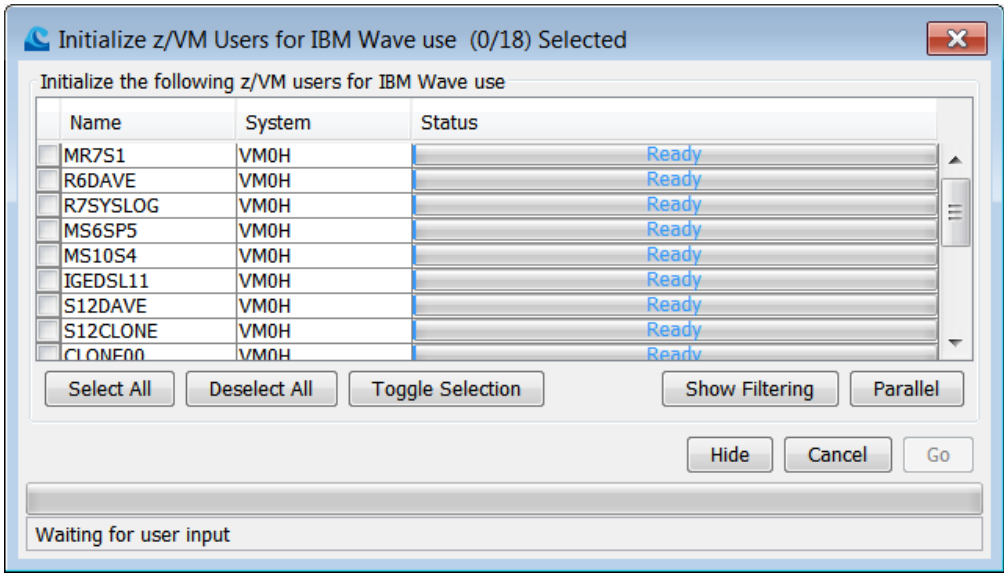

Figure 65. Initialize z/VM Users for IBM Wave use

Note: You can run **Init Users for IBM Wave use** from the **Current System Viewer** or the **Enterprise Viewer**.

Certain IBM Wave actions require that the executable files and some additional prerequisites be present in the Linux guest. The **Init for IBM Wave use** action verifies the prerequisites and places the files in a dedicated directory on the guest.

You must select the z/VM Linux guests that you want to initialize. You can use either of the following methods:

- To select all Linux guests, click **Select All**.
- To select one or more individual Linux guests, click the check box next to the guest that you want to initialize.

  Note: When a guest is already initialized, the check box is not checked by default.

When you click **Go**, IBM Wave prompts you for the authorized credentials. For information about the authorized credentials, see "z/VM guest and virtual server functions" on page 78.

Init Users for IBM Wave use process

The following information outlines the complete **Init Users for IBM Wave use** process.

For each guest, the steps for the **Init Users for IBM Wave use** process are as follows:

1. Opens an SSH tunnel into the Linux guest.
2. Checks the operating system. If the operating system is incorrect, it is automatically updated.
3. Verifies that a version of IBM Wave is not installed on the guest (the IBM Wave RPM):
   - If IBM Wave is installed on the guest, the process sets the guest as an internal IBM Wave guest, and then ends.

4. Verifies that the authorized commands (root privileges) can be attained by using the supplied credentials.
5. Verifies the Linux guest prerequisites. If any of the checks fail, a warning is added to the COR Log:
   a. CMS file system (CMSFS)
   b. VMCP
6. Verifies that the guest has the `LOGOFF ON SHUTDOWN` flag set in the `zipl.conf` file.
7. Verifies that both the `zipl.conf` and the `fstab` files adhere to the supported formats for IBM Wave.
8. Creates the "*waveuser*" user, add it to the `/etc/sudoers` file and set its password.
9. Puts the *waveuser's* public key into `/usr/wave/.ssh`.
10. Adds a link to the IBM Wave short service machine (minidisk 399).
11. Queries FCP connections for the guest. IBM Wave queries the FCP devices that are attached to this guest, the WWPNs that are accessible through the FCP devices, and the LUNs attached to the guest.
12. Install the following shell scripts under the `/usr/wave` directory:
    - `wavestrt` - Responsible for changing the host name and IP addresses of guests during the first boot after the clone process.
    - `waveload` - Loads several script and templates that are used by various IBM Wave functions
13. Creates the following empty IBM Wave exit scripts:
    - WAVECloneConfigExit.sh - The WAVECloneConfigExit.sh runs during the first boot of a cloned guest before TCP/IP initialization.
    - WAVENetConfigExit.sh - The WAVENetConfigExit.sh runs after a successful connect or disconnect action completes for the guest.
14. Adds the `waveclone` service to run on server startup.

Based on whether all the steps were successful, IBM Wave either clears or denies certain actions to be run on the guest. For example, if the **Init for IBM Wave use** process finds the `fstab` or `zipl.conf` file contains a by-UUID reference to DASD volumes, IBM Wave cannot allow the guest to be cloned from. For more information, see Appendix D, "IBM Wave by-id error resolution," on page 251.

Show Changes Log

Click **User Tasks** > **Show Changes Log** to open the Automatic Change Synchronization (ACS) window with the **ACS Log**. The ACS window lists all of the requests (Add, Delete, Update) registered by IBM Wave for the current session.

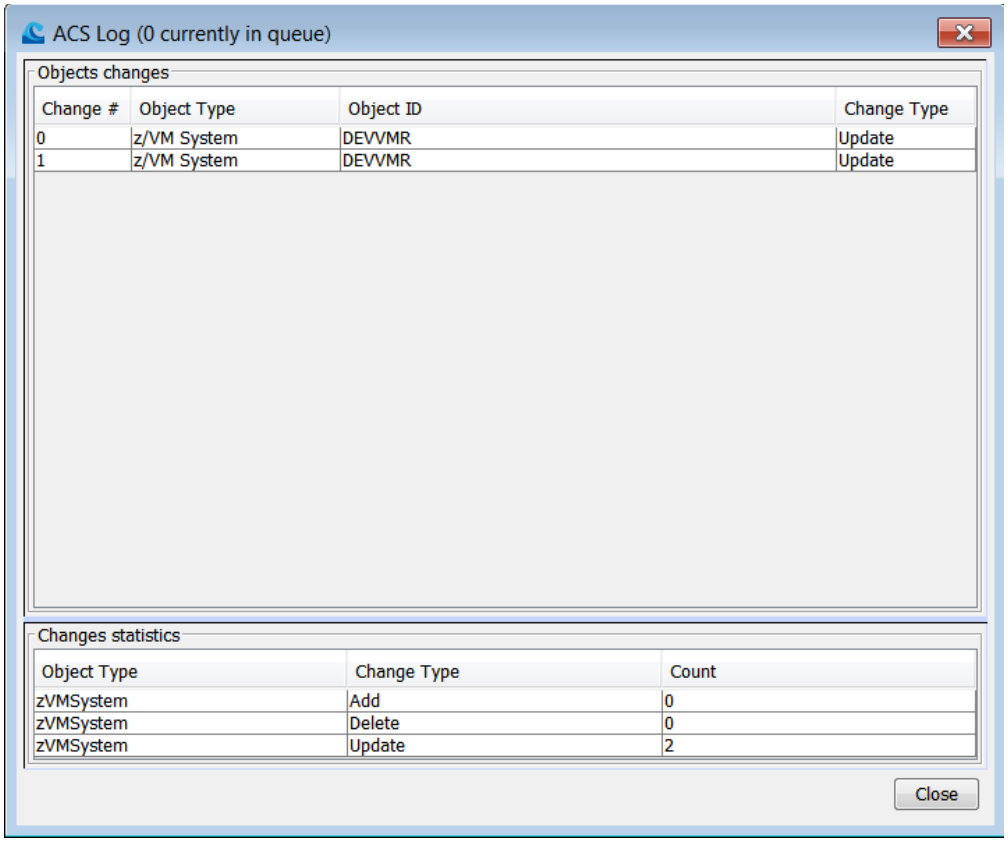

Figure 66. Show Changes Log (for ACS)

Trace GUI

This option switches the debug mode of the GUI client. Use only when directed by IBM support personnel.

Hardware functions

All hardware function that performs an update, such as an add, delete, or update, are relevant only for IBM Wave Users who have the Site Level Administrator (SLA) role and permissions.

CPC actions

Use CPC actions to work with the central processor complex (CPC) in IBM Wave.

The following central processor complex (CPC) actions are for IBM Z and IBM LinuxONE™. The CPC actions or tasks in the information matches the name of the task in the IBM Wave user interface. In some cases, the window or pane name might not match the CPC actions. For example, when you click **Add New CPC**, the **Create New CPC** window opens.

Add New CPC

Use the "Add New CPC" action to add an IBM central processor complex (CPC) to IBM Wave.

Use the **Add New CPC** action to add an IBM Z or IBM LinuxONE CPC to IBM Wave. Use one of the following options to **Add New CPC**:

- In the **Hardware Viewer**, right-click in the white space, and then click **Add New CPC**.
- From the **IBM Wave Main Menu** , click **Administrative** > **Site Management** > **Add New CPC**.

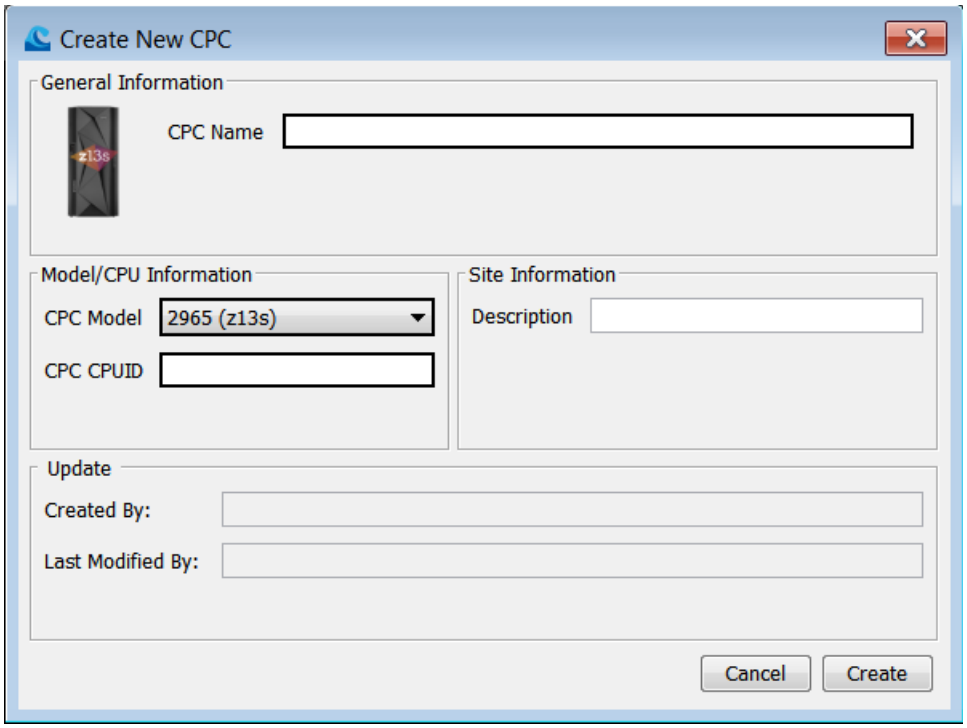

Figure 67. Add New CPC

The following fields are available:

- **CPC Name** - The CPC name. The name has no physical significance, and can be any name that you decide. Make the name applicable to your environment.
- **CPC Model** - The model of the CPC. The icon for the CPC in the **Hardware Viewer** reflects the system model.
- **CPC CPU ID** - The required CPU ID that validates that IBM Wave is licensed to run on your system. To obtain the CPU ID, run the CP command **`QUERY CPUID`**. The command returns the following response:

  ```
  CPUID = aasssssscccddddd
  ```

  The `ssssss` string (hexadecimal digits three through eight) is the CPC value that you must enter in IBM Wave.

Determining the CPC Status

Use the hover help to view the **Status** of the CPC. The CPC **Status** can be one of the following states:

- **Active** - The CPC is active. z/VM System can be added to it and managed by IBM Wave.
- **Inactive** - The CPC is not active. Any z/VM Systems added to the CPC cannot be monitored or managed by IBM Wave.
- **Restricted** - The CPC is restricted and you cannot change it.

Troubleshooting adding a new system

If you have problems adding a CPC, see .

Add New System

How to “Add New System” (z/VM system) to IBM Wave.

To add a z/VM system to IBM Wave, from the **Hardware Viewer**, right-click the z Systems CPC, and then click **More Actions** > **Add New System**. The new z/VM System must be declared as an external entity manager to be able to run on the CPC. For details, see “Create a new external entity” on page 171.

To update an existing z/VM system (LPAR), from the **Hardware Viewer**, right-click on the z/VM you want to update and click **Update Details**.

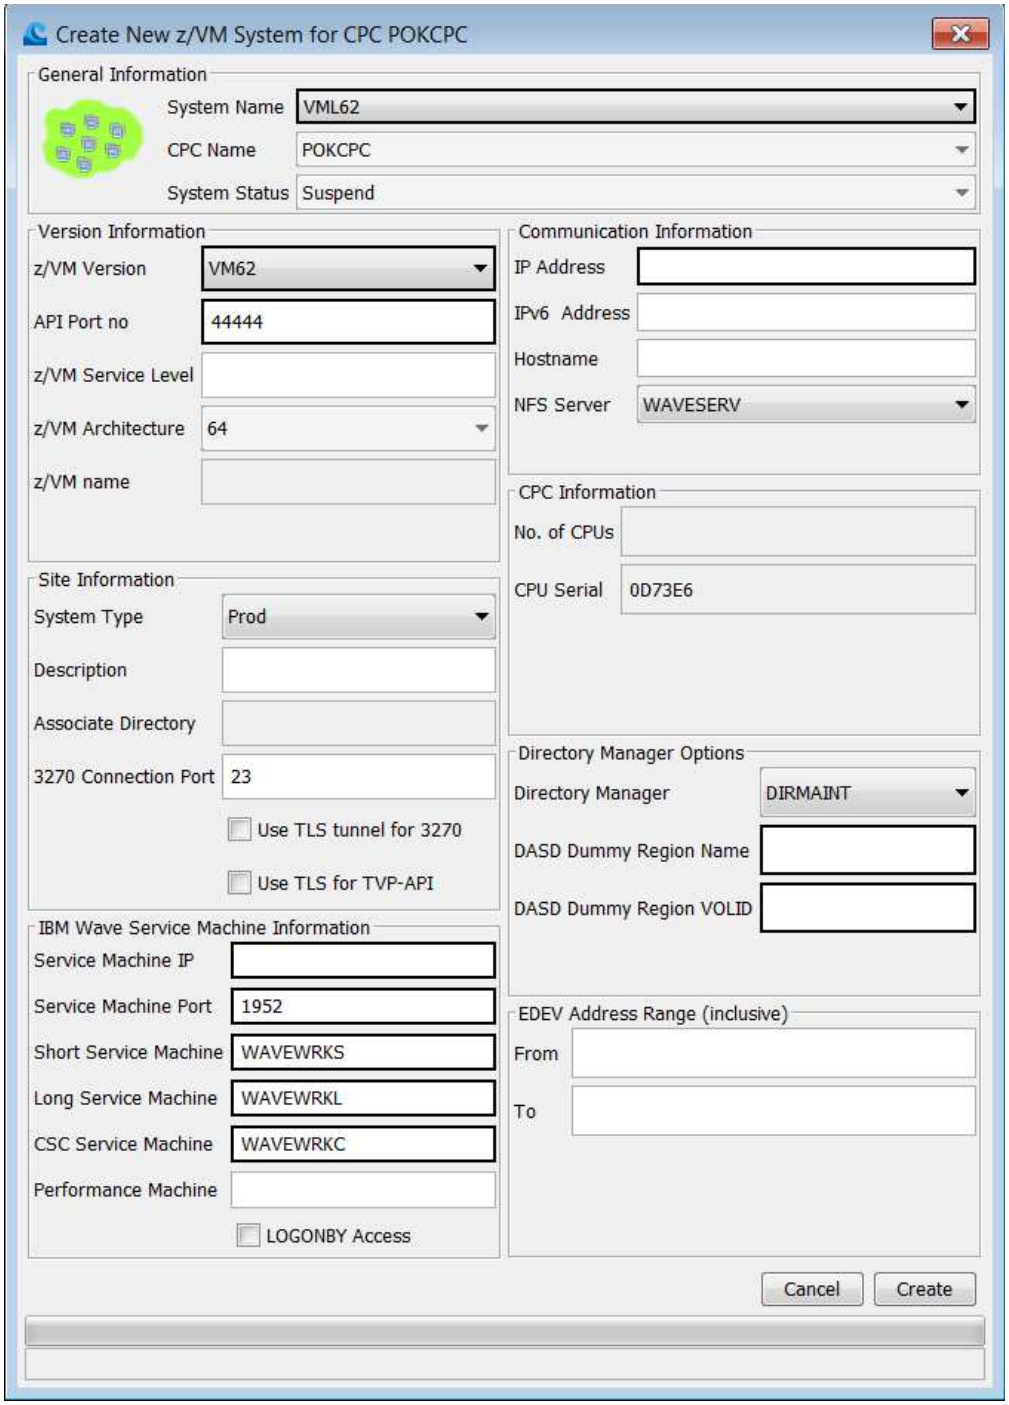

Figure 68. Add New z/VM System

The following fields apply to the **Add New System** window:

- **System Name** - The name for the new z/VM System. The name has no physical significance, but it must first be predefined as a z/VM type managed entity. For more information, see "External entities functions" on page 171.
- **CPC Name** - The name of the CPC where this z/VM system is stored.
- **System Status** - Shows the status of the system. After you add a system, the system is set to "Suspend" status. A suspended system is not automatically updated by the BTS and managed guests on that system do not receive updates.

During auto-detect, the system status changes to “Autodetect”. When auto-detection is complete, the system state changes to “Active”. The system status can also become “Not Reachable” if SMAPI is not reachable or “Suspend - Credentials Invalid” if the credentials of the API user are expired. If a user used the **Shutdown System** action, the system status changes to “Shutting Down”.

- **z/VM Version** - The z/VM operating system (OS) version that runs on the z/VM.
- **API Port No** - Specify the port on which the API server is listening of which the default is 44444.
- **z/VM Service Level** - Input the z/VM OS service level.
- **z/VM Architecture** - 64-bit.
- **z/VM Name** - The name of the z/VM System.
- **IP Address** - The IPv4 IP address for the z/VM System. The IPv4 IP must be the same IP address on which the API server listens for incoming requests.
- **IPv6 Address** - The IPv6 IP address for the z/VM System. IBM Wave currently does not support the use of IPv6 addresses. This field is for future use.
- **Hostname** - An optional host name for the z/VM System.
- **NFS Server**- The default NFS Server to use when scripts are run on guests that belong to this z/VM System.
- **System Type** - An optional indicator of the z/VM System type such as Production, QA, Debug, Test, and others.
- **Description** - An optional description for the z/VM System.
- **Associate Directory** - The description of the z/VM directory, which is completed during auto-detect.
- **3270 connection port** - The port through which 3270 session can be established for this z/VM System.
- **Use TLS/SSL Tunnel for 3270** - Controls whether 3270 must be run over an SSL/TLS tunnel. For more information about SSL/TLS support, see .
- **No. of CPUs** - The number of physical CPs or Integrated Facility for Linux (IFL) instances that are defined for this z/VM System LPAR. These fields are disabled and are periodically updated automatically by IBM Wave.
- **CPU Serial** - The serial number of the CPU.
- **Service Machine IP** - The IP address on which the WAVEWRKS Service Machine listens for incoming connections. By default, the IP is the same as the **IP address** field.
- **Service Machine Port** - The port on which the WAVEWRKS Service Machine listens for incoming requests. The default port is 1952.
- **Short/Long/CSC Service Machine** - The z/VM Guest names for the IBM Wave service machines.
- **Performance Machine** - The z/VM Performance Toolkit virtual machine name that is running on this z/VM System. It is used to retrieve performance data. If regular performance data is retrieved, by using the INDICATE command for this z/VM System, leave the field blank.
- **LOGONBY Access** - When selected, FTP access to the system's service machines uses the LOGONBY method of authentication.

  Note: LOGONBY Access is not selected by default.
- **Directory Manager** - Type of directory manager that is installed on the z/VM System. This parameter is checked later in the auto-detect process.

- **DASD Dummy Region Name** - A name for a dummy DIRMAINT region. This region is defined as one cylinder and is added to every DASD Group created by IBM Wave. This field is not relevant if VM: Secure and VM: Direct is selected as the Directory Manager. For more information about the use of the Dummy Region, see the Storage Management in .
- **DASD Dummy Region VOLID** - The VOLSER on which to define the DASD Dummy Region. The VOLSER must be an existing and real VOLSER defined to the z/VM System. This field is not relevant if VM: Secure/VM: Direct is selected as the Directory Manager. For more information about the Dummy Region, see "Storage Management" in .
- **EDEV Address Range (inclusive)** - The range of preferred addresses for the emulated device (EDEV) definitions that can be automatically used when an EDEV for SCSI storage on VM systems is defined.
- **From/To** - Starting Address Range/End Address Range. The values are four hexadecimal digits only.

When you click **Update**, a series of integrity checks are run against the z/VM System by using the Systems Management API (SMAPI). The CP Level, English language, and CPU ID are checked. After the integrity checks complete, the z/VM System is added to the IBM Wave database, and you can begin the **Auto-Detect** process.

Auto-Detect overview

The Auto-Detect process is a wizard that identifies a newly created z/VM System. It is comprised of several steps to set up the service machines, directory manager, networking, and classification. For the "Auto-Detect Wizard", see the complete process in

Display Details

Right-click a CPC, and then click **Display Details**. Use this action to display the details for a selected CPC. For descriptions of the display details, see "Add New CPC" on page 58.

Lock or Unlock a CPC

Use this task to lock or unlock an IBM Z CPC and all of the z/VM Systems that are running on it.

To lock or unlock a CPC, right-click on the IBM Z CPC and then click **More Actions** > **Lock** or **More Actions** > **Unlock**.

Locking a CPC effectively locks all entities that are owned by the selected CPC including its z/VM Systems and any guests running on the systems, Virtual Networks (VN), and others.

Note: Device Pools, Virtual Network Segments (VNS), and other managed entities, that are defined as global for IBM Wave, are not affected by the lock and unlock action.

Remove CPC

Use the **Remove CPC** action to delete an entire CPC from IBM Wave management.

Note: Remove CPC permanently deletes all database records for the selected CPC including all of the z/VM systems that are defined to the CPC.

For more information about adding or removing a CPC, see the topics about "Add New CPC " or "Remove CPC" in .

Update an IAN

Use this task to read, update, or delete an *Intelligent Action Note* (IAN) that is attached to a CPC.

Right-click on a CPC, and then click **More Actions** > **Read IAN** > |, **Update IAN**, or **Delete IAN**. An IAN attached to a CPC does not affect your use of the z Systems CPC.

Update Details

Right-click on the CPC, and then click **Update Details** to update the details. For descriptions of the update details that you can change, see "Add New CPC" on page 58.

z/VM system actions

z/VM system actions and tasks are relevant to a specific z/VM System. To access all actions, right-click on the z/VM System, and then select an option from the menu.

Activate default guests

Use the **Activate** action to activate all z/VM Guests on the currently selected z/VM System. The **Activate** action opens a multiple task window that is populated with all the z/VM Guests for the selected z/VM System.

Add Page and Add Spool

IBM Wave allows the IBM Wave User to add Page and Spool Disks to the z/VM LPAR.

The action will bring up the "Add Page" or "Add Spool" multiple action window. This window displays all known Page/Spool disks. Click “Go” to start the action. IBM Wave first ensures that the selected disk is available and contains a Page/Spool area. Then, it attaches it to the z/VM System.

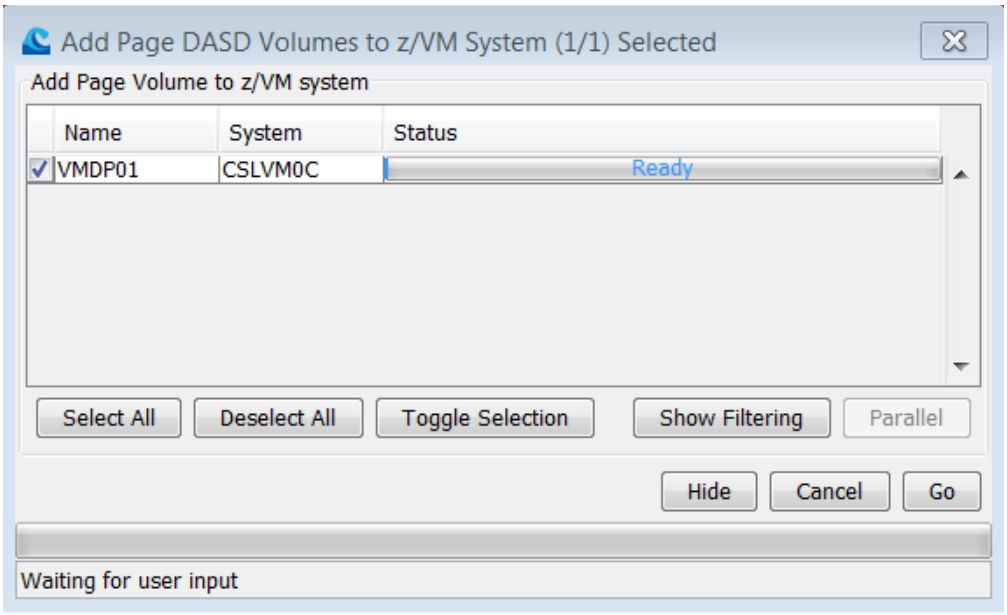

Note: IBM Wave does not make the change permanent. On the next IPL, the added Page/Spool disks are not automatically attached.

Also, if the z/VM System is a member of an SSI cluster, the DASD Volume must be properly formatted with the SSI cluster name.

For more information about Page/Spool disk handling in IBM Wave, see “z/VM Page and Spool Disks Handling” in .

Disconnect

Use the **Disconnect** option to disconnect the display and reset the **Current System Viewer**.

Use this action to disconnect the display of a selected z/VM System. The **Disconnect** resets the **Current System Viewer** to the default panel.

To disconnect a z/VM system, right-click the z/VM system, and then select **Disconnect**.

Display details

Use this action to display details on a selected z/VM System. For a complete field description, see "Add New System" on page 60.

Lock and unlock

This action will allow the IBM Wave User to lock or unlock the z/VM System. Note that locking a z/VM System will also lock all the objects which belong to it (z/VM Guests, Virtual Networks, DASD Storage Groups, DASD Volumes etc.). It will not be possible to unlock these objects as long as the z/VM System is locked. If there are specific objects that were already locked before the Lock System action was taken, these will return to be locked when the z/VM System is unlocked.

Note: Because Device Pools, VNSs and Managed Entities are defined as global for an entire installation of IBM Wave, these are not affected by the lock action.

Read, update, or delete an IAN

With this action, you can update, read, or delete the IAN attached to the z/VM System.

Reconnect

Use this action to reconnect to a z/VM System. The reconnect process rereads all the data relevant for this z/VM System from the IBM Wave database.

Remove system

Use the **Remove** option to remove a z/VM System from IBM Wave.

To remove a z/VM system, right-click the suspended z/VM system, and then select **More Actions** > **Remove**.

The z/VM System must be in **Suspended** status for the **Remove** action to be available.

Note: This action permanently deletes all records for the selected z/VM System from the IBM Wave database.

Shutdown System

Use the **Shutdown System** option to shut down a z/VM System.

To shut down a z/VM system, right-click the z/VM system, and then select **More Actions** > **Shutdown System**.

The **Shutdown System** action issues the z/VM CP command **SHUTDOWN WITHIN** through the IBM Wave Service Machine. The shutdown procedure notifies you about the following conditions:

- Does the z/VM System that is selected for shutdown contains the WAVESRV server?
- Does the z/VM System contains the WAVESRV server through which the current session is being run?

The status of the system changes to "Shutting Down" until the connection to SMAPI drops. After the connection to SMAPI drops the system status changes to "Not Reachable".

Note: If you are not the only user who is logged in to IBM Wave, IBM Wave starts a process to move to Single User Mode. For more information about Single user mode, see

Purge Spool

To open the **Spool Purger**, right-click on the z/VM system, and then select **More Actions** > **Purge Spool**.

The **Spool Purge** menu provides you with three methods to purge the z/VM System's spool and two methods to test your options and control file. The scenarios must be preconfigured by the z/VM System Programmer by using the z/VM SFPURGER utility.

You can use the following three modes to run **Spool Purger**:

- **Run** - To run **Spool Purger** with conditions.
- **Force** - To run **Spool Purger** immediately.
- **SOS** - To run **Spool Purger** the same way as **Force**, but you can use it for *emergency* spool file maintenance.

To verify that your options and control files work correctly, *without performing spool file maintenance*, you can test **Spool Purger** by using one of the following two modes:

1. **TEST** - To verify that your options and control files work correctly. The z/VM SFPURGER utility processes the SFPURGER OPTIONS file (if your z/VM System Programmer created one), and the SFPURGER CONTROL file.
2. **TESTSOS** - To verify your options and control files work correctly the same way as **TEST** mode, except **TESTSOS** processes the *emergency control file*, SOS CONTROL.

.

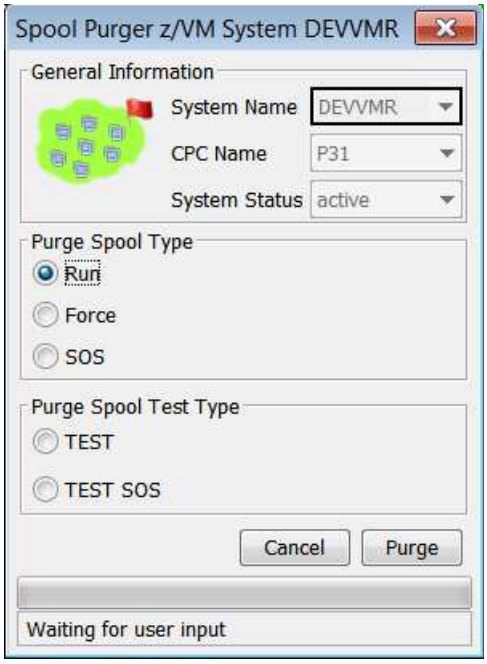

Figure 69. Purge Spool

For complete information about the SFPURGER utility, see *z/VM: CMS Commands and Utilities Reference* for the version of z/VM that is running in your environment.

Related information:

SFPURGER utility for z/VM V6R3

z/VM: CMS Commands and Utilities Reference for all z/VM releases

Update Details

Use the **Update Details** action to update the z/VM System parameters. You can update some of the system parameters. For the complete field descriptions, see "Add New System" on page 60.

If you update a system to "Suspend" state and you are not the only user who is logged in to IBM Wave, IBM Wave tries to switch to Single User Mode. For more information about Single user mode, see

To open the **Update Details** menu, right-click on the z/VM system you want to update, and then select **Update Details**.

z/VM User Groups functions

The following actions and tasks are relevant to the **Current System Viewer** > **z/VM User Groups**. Any addition, deletion, or update tasks are relevant to users who are assigned the Site Level Administrator role.

Add a Site Defined Group, Project, or Account

Use the following steps to **Create a new Site Defined Group**, **Add New Project**, or to **Add New Account**.

1. In the **Current System Viewer > z/VM User Groups** tab, use the **Group By:** drop-down menu to select "Site Defined Group", "Project", or "Account" that you want to create.
2. Right-click in the white space of the **z/VM User Groups** tab, and then select the option that matches your **Group By:** selection.
   - **Add New Site Defined Group**
   - **Add New Project**
   - **Add New Account**.

For example, if you select the **Add new Site Defined Group** menu, the **Create new Site Defined Group** menu opens (as shown in Figure 70).

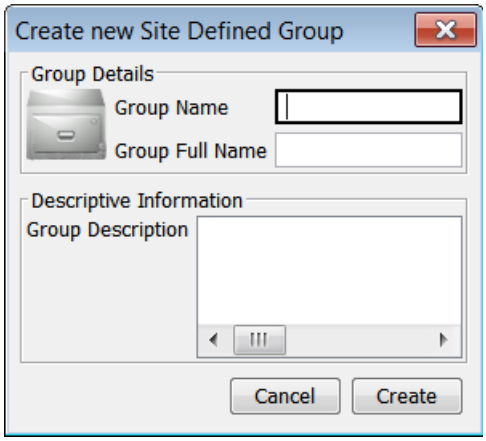

Figure 70. Add new Site Defined Group

The following fields are in the **Add new Site Defined Group** menu:

- **Group Name** - The required name of the new Site Defined Group. The name of each Site Defined Group must be unique.
- **Group Full Name** - An optional full descriptive name for the new Site Defined Group.
- **Group Description** - An optional description for the new Site Defined Group.

For complete information about the fields in the **Add new Project** and **Add new Account** menu, see the topics "z/VM Account Manager" and "Project Manager" in *IBM Wave for z/VM: Administration and Customization*.

Actions for selected z/VM User Groups

The following topics are described:

Collapse

Use this action to collapse an expanded group or groups.

Delete group

Use this action to delete a group. A group must be empty before it can be deleted. The default IBM Groups and the IBM Wave Internal Group cannot be deleted.

Note: Due to the security subsystem, a group might appear to be empty. Since Scopes and Permissions are done by project assignment, and not by Site Defined Group assignment, a certain IBM Wave User can see a Site Defined Group that appears empty, but really isn't. If this is the case, the delete action will be available, but will fail with a warning message to the IBM Wave User.

Display group information

Use this action to display information on the selected Site Defined Group. For a full description of the fields see "Add a Site Defined Group, Project, or Account" on page 68.

Note: When you use the **Group By** feature, the action changes according to the "Group-by" value. For example, if the **Group By** value is "Group By Projects" the interface shows the Projects view.

Expand

Use this action to expand a collapsed group or groups.

Lock or unlock

To lock or unlock a selected z/VM User Group, right-click on a group and select **More Actions** > **Lock User Group** | **Unlock User Group**.

Note: This action also locks any z/VM Guest that belongs to the Site Defined Group. z/VM Guests that belong to the Site Defined Group are unlocked when the

group is unlocked. The exception are any z/VM Guests that were locked prior to the group being locked, which retain their original lock status.

Read, update, or delete IAN

Right-click on the group, and then click **More Actions** > **IAN Read | IAN Update | IAN Delete** to read, update, or delete the Intelligent Active Note (IAN) attached to the z/VM Users Group.

Expand

Use this action to expand a collapsed group or groups.

Transferring z/VM guests

With IBM Wave, you can transfer z/VM Guests from one group to another. The outcome of the transfer depends on which "Group By" value you select in the **Current System Viewer**.

For example, when Project is selected as the "Group By" value, z/VM Guests are transferred between projects.

To transfer one or more z/VM Guests between groups, select the guests to transfer, and press CTRL while dragging them to the new group. To signify a valid transfer, the group you are moving the guest or guests to is highlighted with a green border.

Note: You cannot transfer z/VM Guests in or out of the following Site Defined Group (SDG):

- Default IBM Wave SDG (labeled "IBM-").
- IBM Wave internal SDG (labeled "WAVE-INTERNAL").

For more information, see "Transferring z/VM guests between groups" on page 25 and "Special considerations for site defined groups" on page 25.

Update group information

Use this action to update a group's information. For a full description of the fields see "Add a Site Defined Group, Project, or Account" on page 68.

Note: If the "Group By" feature is used, this action will change according to the group-by value. For example, if the Group By value is "Group By Projects" then this action will show the Update Project Details. The same applies to the Display Group Information action described above.

z/VM User Actions for Groups

You can apply **z/VM User Actions** to all z/VM users in the **Current System Viewer** > **z/VM User Groups** viewer. The actions have the same effect as the traditional z/VM Guest actions, but are accessed from the **z/VM User Groups** viewer.

Right-click on the z/VM User Group, and then click **z/VM User Actions** > ***selected action***.

Where *selected action* is one of the following actions:

- "Activate" on page 72
- "Assign distribution" on page 72
- "Assign to Project" on page 72

Activate

To activate all z/VM Guests inside a group, right-click in the group and then select **z/VM User Actions** > **Activate**. For more information about the activate function, see "Activate" on page 110.

Assign distribution

Use this action to assign an operating system distribution to all z/VM Guests inside a group. For complete information about **Assign Distribution**, see "Assign Distribution" on page 116.

Note: This option is only available when all the z/VM Guests inside the z/VM Guest Group are Linux Virtual Servers.

Assign to Project

Use the **Assign to Project** action to assign all the z/VM Guests in a group to a **Project**. The **Assign to Project** option is available only when all of the z/VM Guests inside the z/VM Guest Group are Linux Virtual Servers. For more information, see "Assign Guests to Project" on page 117.

Create New z/VM Guest

Use the following information to create a new z/VM Guest.

You can open the **Create New z/VM Guest** window from a Site Defined Group (SDG) or a Project defined in the **Hardware Viewer**. Right-click inside the SDG or Project, and then select **z/VM User Actions** > **Create New z/VM User**.

When you use the SDG method to **Create New z/VM Guest**, the "Group Name" field is completed. When you use the Projects view to **Create New z/VM Guest**, the "Project" field is completed. You can manually specify the "Group Name".

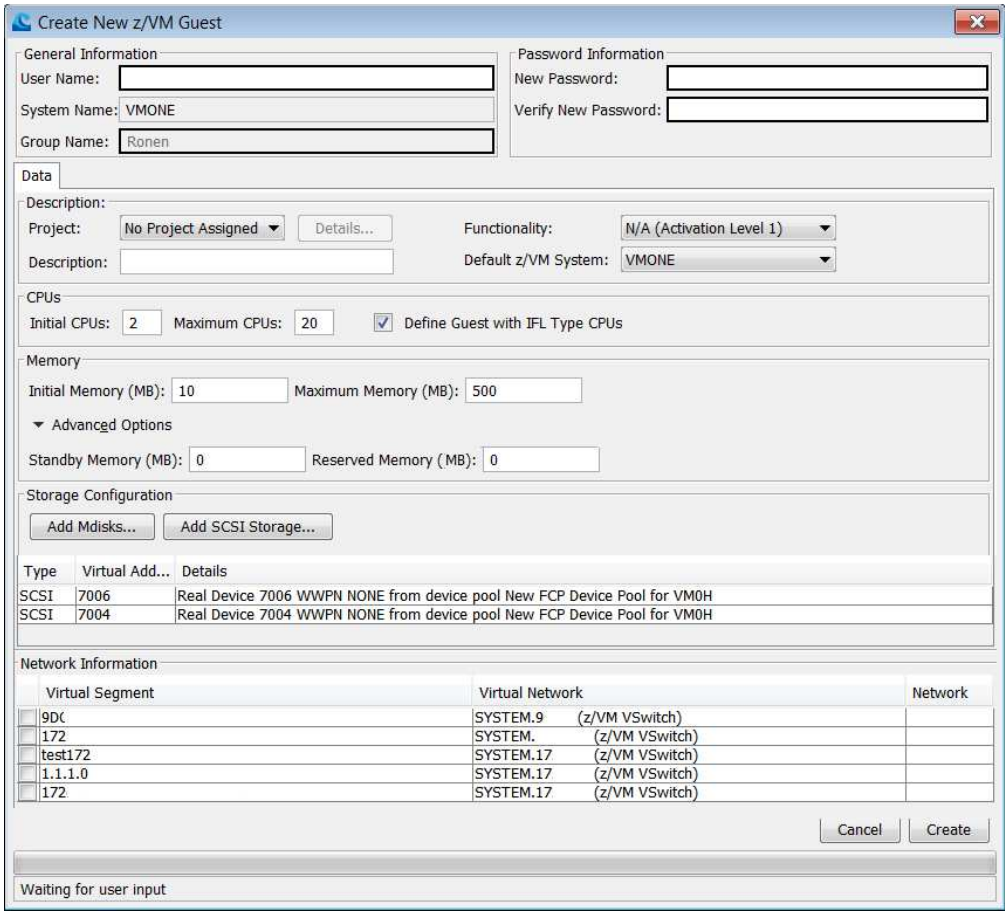

Figure 71. Create New z/VM Guest

The following options are available in the **Create New z/VM Guest** window:

General Information

User Name
The new z/VM Guest name, which is limited to 8 characters. The name must be unique within each z/VM System. If the user name exists, the system issues an error message.

System Name
The name of the z/VM system on which the z/VM Guest is defined.

Group Name
The name of the site defined group (SDG) to which the z/VM Guest belongs. When you **Create New z/VM Guest** from the Projects view, you can use the drop-down menu to complete the **Group Name** field.

Password Information

New Password
A password for the new z/VM Guest. The password field is 8 characters.

Verify New Password
Reenter the 8 character password for confirmation.

Description

Project
An optional project name you want assigned to the z/VM Guest. The **Project** field is predefined when you create the guest from the project view. For more information, see "Assign to Project" on page 72.

Functionality
An optional function field where you can specify the activation level. For more information, see "Activate" on page 110.

Description
An optional description field for the z/VM Guest.

Default z/VM system
The default z/VM System to which you assign the new z/VM Guest.

CPUs

Initial CPUs
The initial number of CPUs assigned to the z/VM guest.

Maximum CPUs
The maximum number of CPUs that are allowed to be assigned to the z/VM guest.

Define Guest with IFL Type CPUs
The check box is selected by default to ensure the proper definitions are created for an Integrated Facility for Linux (IFL).[1].

Memory

Initial Memory (MB)
The initial amount of memory, specified in MB, that is allocated to the z/VM Guest.

Maximum Memory (MB)
The maximum amount of memory, specified in MB, that can be allocated to the z/VM Guest.

Advanced Options (for Memory) - controls adding memory dynamically. To add memory dynamically, you must set Standby Memory or Reserved Memory. The options are collapsed by default, but shown in Figure 71 on page 73 for illustrative purposes.

Standby Memory
The amount of standby memory.

Reserved Memory
The amount of reserved Memory

Important: Dynamic Memory reconfiguration is only supported when the initial memory size for the guest is an exact multiple of the memory block size. To understand the calculation of memory block size, see the following topics:

- "Overview of estimating memory and CPU requirements" at www.ibm.com/support/knowledgecenter/SSB27U_6.2.0/com.ibm.zvm.v620.hcpl0/memcpuov.htm.

1. An Integrated Facility for Linux (IFL) is a processor that is dedicated to Linux workloads on IBM z Systems. An IFL configuration is supported by z/VM, the IBM Wave for z/VM software, and the Linux operating system.

- "Steps for estimating memory and CPU requirements" at http://www.ibm.com/support/knowledgecenter/SSB27U_5.4.0/com.ibm.zvm.v54.hcpl0/hcsx0b3030.htm.

Storage Configuration

Add Mdisks

To **Add Mdisks** (minidisk) storage, you must specify the amount of DASD storage you want to add in MB or GB in the **Size** field. You must also select the **Storage Group** from which to obtain the storage for the new z/VM Guest. The choices available for **Storage Group** are only the groups in your scope.

Only one entry of DASD storage can be present. You must remove the DASD Storage entry before defining a new entry. IBM Wave might split the amount into more than one mdisk statement while creating the guest to allocate the requested storage. To remove, right-click on the table entry, and then click **Remove**.

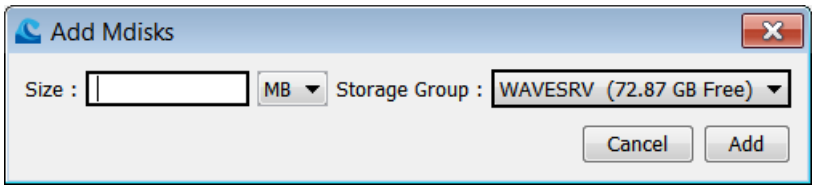

Figure 72. Add Mdisks

Add SCSI Storage

The **Add SCSI Storage** button is enabled when the system has FCP device pools that contain FCP devices. The **Add SCSI Storage** button is disabled when no FCP devices are defined.

From the **Storage Configuration** pane, click **Add SCSI Storage**. Select a Device Pool from the "Devices in Device Pool" menu to see the list of available device pools to add to the new guest.

Important: If there are no FCP devices available, work with your Linux Administrator and z/VM Storage Administrator to define storage and ensure the support is in place. For more information, see "Direct-attached SCSI Devices" on page 166.

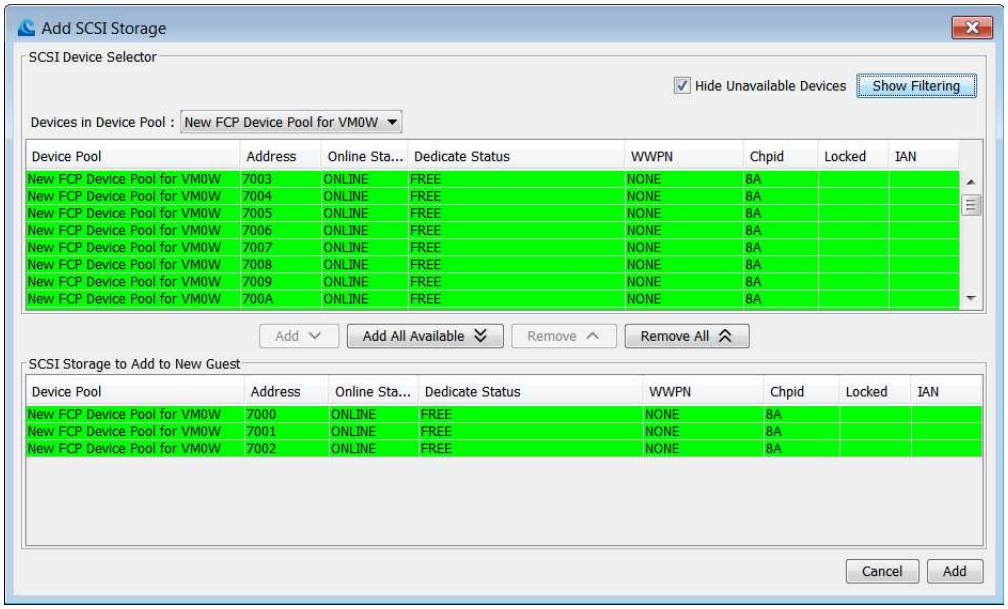

Figure 73. Add SCSI storage

Network Information

A multiple selection table that contains network interface information for **Virtual Segment**, **Virtual Network**, and the **Network** with its associated IP address. Select the check box next to the **Virtual Network** to choose the connection for the new z/VM Guest.

Important: To view the menu of available VSwitch options, check the row in the **Network Information** multiple selection table, and then the cell in the **Virtual Network** column.

Deactivate

Use **Deactivate** to deactivate all z/VM Guests inside a group. For a complete description of the deactivate function, see “Deactivate z/VM Users” on page 118.

Delete Linux guests

Use this action to delete all z/VM Guests inside a group. For more information about the **Delete** guests multiple action, see “Delete z/VM Guests” on page 120.

Note: This option is available when all the z/VM Guests inside the group are Linux Virtual Servers.

Execute script

Use this action to execute a script on all z/VM Guests inside a group. For more information about the "Execute Script" action, see “Execute script” on page 121.

Note: This option will only be available if all the z/VM Guests inside the z/VM Guest Group are Linux Virtual Servers.

Execute REXX

With the proper scopes and permissions, you can run a REXX EXEC for every virtual object that IBM Wave manages.

A REXX EXEC can run for each object that is selected. You can specify any parameter to REXX that is relevant to the group of objects selected. Only groups of

objects from one type can be selected for this action. For example, either a group of z/VM Guests or a group of DASD volumes, but you cannot select both z/VM guests and DASD volumes.

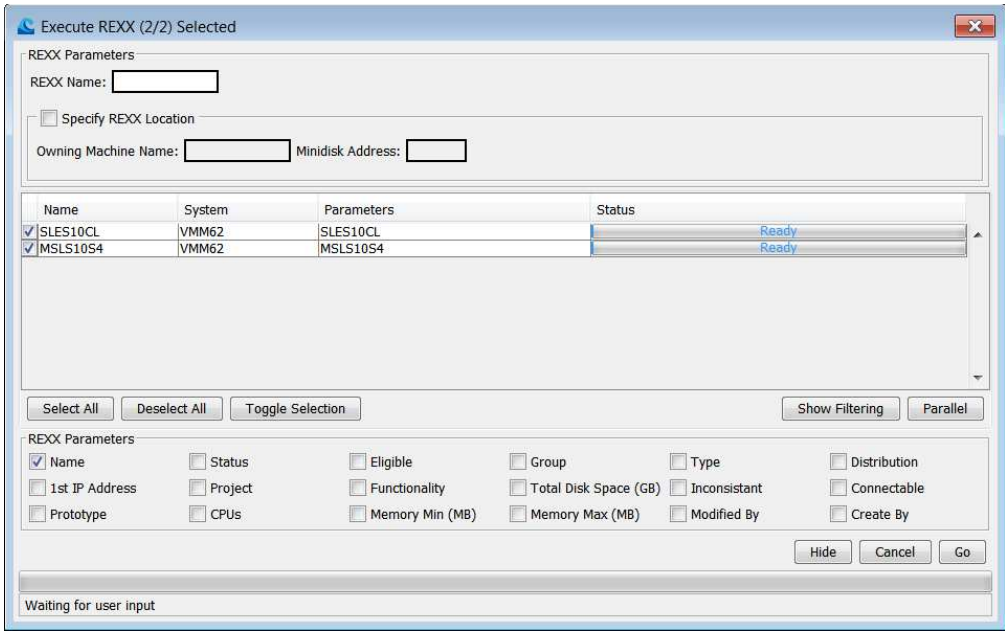

Figure 74. Execute REXX group action

In Figure 74, two z/VM Guests are selected. You can specify the location of the REXX EXEC. The default location (unless the **Specify REXX Location** check box is selected) is the WAVEWRKS 399 minidisk. The selected REXX EXEC runs for each selected object (twice in Figure 74).

By default, the parameter that is given to the REXX EXEC is the object name. To change the default, select the check box properties from the list in the **REXX Parameters** pane. While no limitation is given to the number of parameters specified, the total length of the parameter field cannot exceed 240 characters.

The parameter values in the parameter list that IBM Wave passes to the REXX exec are separated by commas. For example, if you selected **Name**, **Status**, and **Group** parameters for a guest that is called `SC0` that is inactive, and is in the `USER-LOCAL` site defined group, the resulting parameter list is the following string:

```
SC0,INACTIVE,USER-LOCAL
```

When you click **Go**, a Background Task Scheduler (BTS) work unit is generated and submitted to the BTS. The work unit contains Execute REXX requests for each object included in the action. The request uses the WAVEWRKS service machine to run the REXX. If the REXX location is specified, the service machine links to the specified location and runs the REXX. To view the results and output of the REXX EXEC, drill down into the BTS work unit, select the request for the particular object, and then switch to the **REXX Output COR** tab.

To capture the output of the CP or CMS commands that REXX runs, use the CP and or CMS pipeline stages. For example, to capture the output of the **`CMS ID`** command and the **`CP D G`** command, the REXX can output code similar to the example shown in Figure 75 on page 78. If the `Say` output is not captured, you can use the output pipe subcommand instead by replacing "**`Say`**" in the code example

with "**output**":

```
Address Command Pipe CMS ID '|' STEM ID.
Do I = 1 to ID.0
   Say ID.I
End
Address Command Pipe CP D G '|' STEM DG.
Do D = 1 to DG.0
   Say DG.D
End
```

Figure 75. Output Example for Pipe Subcommand

Note: The REXX EXEC is run by the WAVEWRKS service machine, which is also used by IBM Wave for other functions. Because it can delay other IBM Wave functions, it is important that the REXX EXEC is not a long-running task.

Generate disk storage map

Use this action to generate a Disk storage map for all the z/VM Guests inside a group, a project or a distribution group. For more information about this action, see "Generate Disk Storage Map" on page 123.

Recycle

Use the **Recycle** action to recycle all z/VM Guests inside a group. For complete details about the recycle action, see "Recycle z/VM users" on page 127.

Note: This option is only available when all the z/VM Guests inside the "z/VM Guest Group" are Linux Virtual Servers.

Send message

Use this action to send a message to all z/VM Guests inside a group. For more information about the "Send Message" action, see "Send message" on page 131.

Note: This option is only available if all the z/VM Guests inside the z/VM Guest Group are Linux Virtual Servers.

z/VM guest and virtual server functions

Certain actions can run on one system while other actions can run against multiple systems.

The following actions are relevant to z/VM Guests and z/VM virtual servers. Even though a virtual server is a type of z/VM Guest, many descriptions refer only to z/VM Guests. Some actions require a direct connection to the z/VM Guest, and other actions require direct interaction with the Linux operating system (OS) running on the z/VM Guest. In general, the following two action types apply to z/VM Guests and virtual servers:

- "Single selection actions" on page 79 are actions that can operate on only one z/VM Guest at a time.
- "Multiple task actions" on page 105 are actions that can operate on one or more z/VM Guests simultaneously.

Depending on the parameter settings for the authentication method (**Administrative** > **IBM Wave Parameters** > **Security**), each action that requires direct access to the Linux OS, requires you to input the authorized SSH credentials, root password, or both.

When credentials are required, IBM Wave automatically prompts you to enter them, as shown in Figure 76.

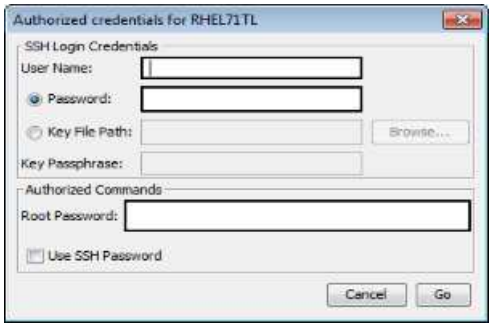

Figure 76. Authorized credentials for Linux guest

For more information about the SSH credentials, see "Linux Login Security Options" and "Changing user preferences" in *IBM Wave for z/VM: Administration and Customization*.

Single selection actions

Single Selection Actions are actions that you can run on only one object at a time. (For example, you right-click to select a z/VM Guest or a z Systems CPC.)

CLC access

Use the IBM Wave Communication-Less Connection (CLC) in an emergency when there is no IP connectivity available for a z/VM Guest. Using CLC, it is possible to make the changes in the z/VM virtual server, and edit files with the **CLCEDIT** command. CLC uses the IBM Wave resource serialization (WRS) mechanism, which means there can be only one concurrent CLC session per z/VM guest.

Note: The vi editor is not available under CLC.

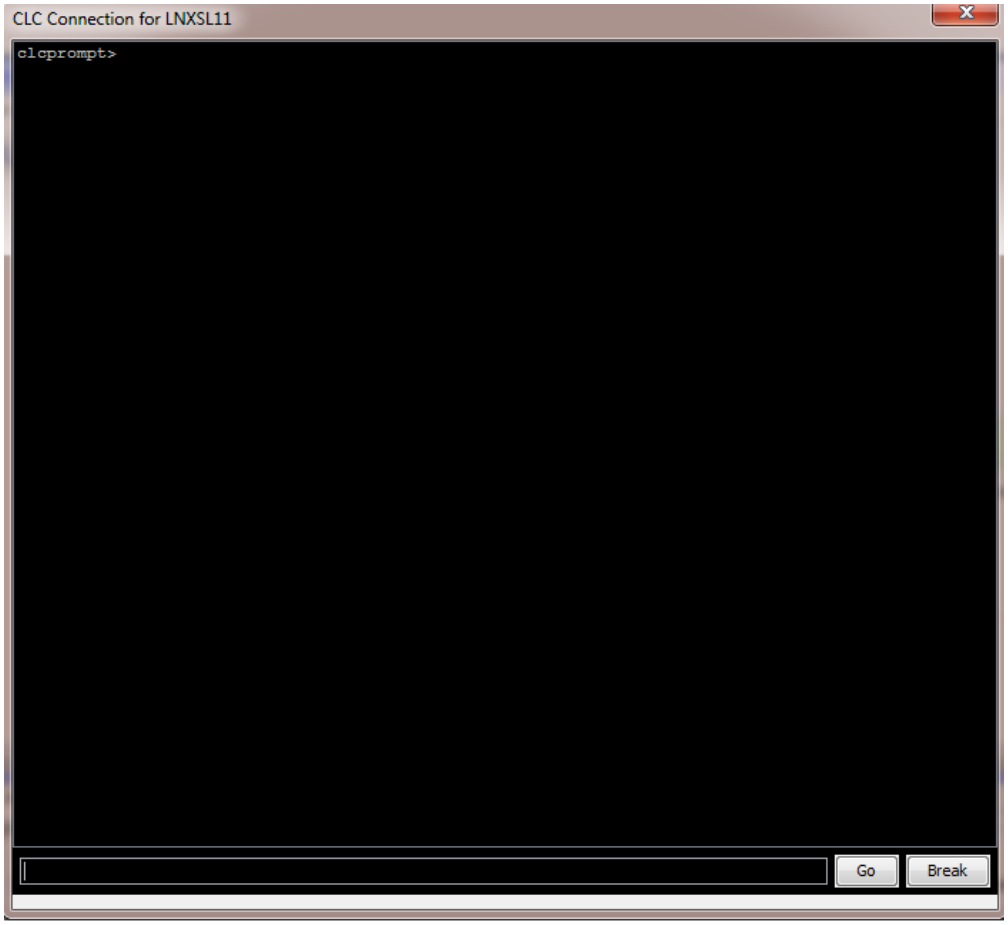

Figure 77. CLC editor

The main text area is used to input text. The editor is a normal (OS-dependent) text input area.

- To send a command to the virtual server:
  - Press "Enter" on a keyboard.
  - Press "**Go**", which is located in the lower-right corner of the CLC window.
- To enter a continuous command (such as **`ping`**), press "**Break**" to return to the command line.
- To edit files in the virtual server, use the **`clcedit`** command (or the **`ce`** command alias).
- To save a file, press "**Save**".
- To close CLC, press "**X**", which is in the upper-right corner of the screen.
- To copy and paste, you can use the normal copy-paste methods.

Note: If you enter the **`Shutdown`** or **`init`** command in the CLC session, the status of the z/VM Guest changes to "Shutting down".

For an overview of the CLC technology, see the topic about "CLC access" in *IBM Wave for z/VM: Administration and Customization*.

Clone

To clone a z/VM Guest, right-click on an inactive z/VM Guest, and then click **Cloning** > **Clone**.

Only the cloned users' IP address and hostname is changed. The clone contains all of the data from the selected z/VM Guest.

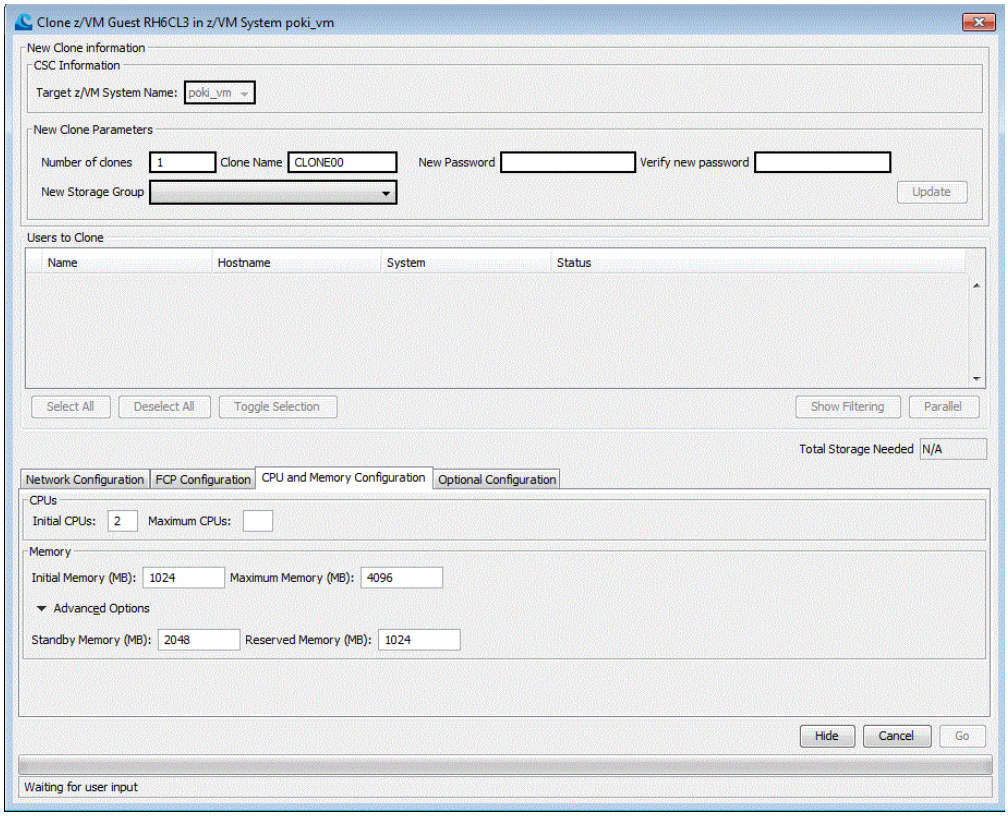

Figure 78. Clone z/VM Guest

The cloning process creates a temporary prototype from which to clone. The prototype is erased at the end of the process. For the information about the fields, see "Clone from a Prototype" on page 141.

Convert to prototype

Use this action to convert a z/VM Guest into a prototype. The action will create a directory manager prototype with the same name as the z/VM Guest. It will also disable logon to the z/VM Guest (for example, set its password to NOLOG). After the action is complete, the z/VM Guest will no longer appear on the z/VM Guests and Groups or the Network Viewers. A new prototype object will be created on the Prototype viewer, from which more z/VM Guests can be cloned.

Before converting a guest to a prototype, a guest should meet all of the following:

1. Assigned the appropriate OS distribution
2. Initialized for IBM Wave use
3. Deactivated

Disconnect from virtual network

Use the **Diconnect from virtual network** action to disconnect a z/VM Guest from a specific Virtual Network. For more information, see "Disconnecting z/VM guests from virtual networks" on page 138.

Display Information

To display information about a z/VM or Linux Guest, right-click on the guest, and then click **Display Information**.

You can use the **Display Information** window to view the details about the guest. **Display Information** uses both API and SSH connections to obtain information from the z/VM Guest (or CLC when the z/VM Guest is an active Linux server). The initial view is set to display information that is stored in the IBM Wave database.

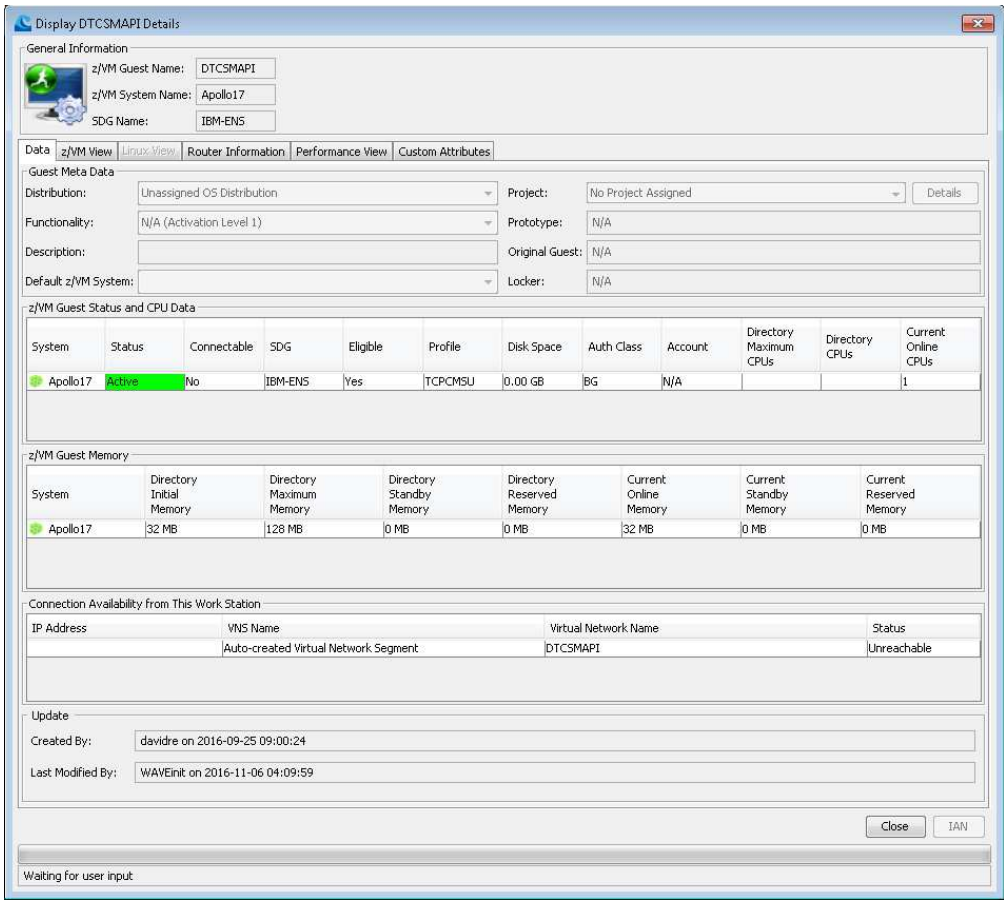

Figure 79. Display information

The following panes and fields are in the **Display** ***guest name*** **Details** window:

Metadata

The **Metadata** tab contains the following information from the IBM Wave database:

- **Distribution** - The operating distribution that is assigned to the z/VM Guest.
- **Functionality** - The activation level that is assigned to the z/VM Guest.
- **Default z/VM System** - The default z/VM system.

- **Project** - The project to which this z/VM Guest is assigned. The field states when the guest has "No Project Assigned".
- **Prototype** - If the z/VM Guest was created from a prototype, the name of the prototype.
- **Original Guest** - The name of the original guest.
- **Locker** - The name of guest user or action that is blocking the enqueue.

The **z/VM Guest Status and CPU Data** table contains the following columns.

- **z/VM System** - The system that the guest is a member.
- **Status** - The status of the z/VM Guest.
- **Connectable** - Indicates whether at least one of the network interfaces is up and running.
- **SDG** - The site-defined group (SDG) to which the z/VM Guest is assigned.
- **Eligible** - Indicate whether the guest is eligible for logon.
- **Disk Space** - The total disk space in GB that is allocated for the z/VM Guest.
- **Auth Class** - The authorization class for the guest.
- **Account** - The Account name to which the guest is assigned if applicable.
- **Directory Maximum CPUs** - The maximum number of CPUs that are allowed to be assigned to the z/VM guest.
- **Directory CPUs** - The number of CPUs currently assigned to the z/VM guest.
- **Current Online CPUs** - The number of CPUs that are currently online.

The **z/VM Guest Memory** table contains the following columns.

- **System** - The z/VM system to which the guest is defined.
- **Directory Initial Memory** - The initial amount of memory in MB that is allocated to the z/VM Guest.
- **Directory Maximum Memory** - The maximum amount of memory in MB that can be allocated to the z/VM Guest.
- **Directory Standby Memory** - The amount of standby memory in MB that is allocated to the z/VM Guest.
- **Directory Reserved Memory** - The amount of reserved memory in MB that is allocated to the z/VM Guest.
- **Current Online Memory** - The amount of memory in MB that is online in the guest.
- **Current Standby Memory** - The amount of standby memory in MB that is available for the guest.
- **Current Reserved Memory** - The amount of reserved memory in MB that is set for the guest.

The **Connection Availablility** table contains the IP address, Virtual Network Segment name, Virtual Network name, and its status.

z/VM View

The **z/VM View** tab displays the directory entry for the z/VM Guest as it appears in z/VM.

Linux View

The **Linux View** tab displays the following Linux information. (If the z/VM Guest is not a Linux Server, the **Linux View** is disabled.)

- **File System information** - The current occupancy of the file system as it appears in the Linux **df -ah** command. If the occupancy exceeds the threshold that is specified in the IBM Wave parameters, the bar is red or orange, depending on the warning and error thresholds defined.
- **Network Information** - The status of all network interfaces. Active interfaces are green. Inactive interfaces are white.
- **Routing Information** - The routes that are defined to this Linux virtual server. Routes that point to specific hosts as gateways are checked for connectivity. When the gateway is reachable, the row is green. When the gateway is not reachable, the row is red.
- **Uptime** - The number of days that the Linux server is active.

Performance View®

The **Performance View** shows performance information for the z/VM Guest. When the z/VM Guest is inactive, the tab is not available. For more information about the fields in the **Performance View** tab, see "Status or Performance View" on page 99.

Custom Attributes

The **Custom Attributes** tab displays all the values that are assigned to the guest for each Custom Attribute defined. For more information and defining custom attributes, see "Assign custom attribute" on page 114 and "Creating custom attributes" on page 114.

Duplicate z/VM User Definition

You can open the **Duplicate z/VM User Definition** window from the z/VM Users Group. Right-click on the guest you want to duplicate, and select **Cloning** > **Duplicate z/VM User Definition**. You can now create a new z/VM Guest from the duplicate.

When you complete the fields and click **Create**, IBM Wave creates a dummy Directory Manager prototype, and then creates the new guest that uses the prototype. The dummy prototype is erased at the end of the process.

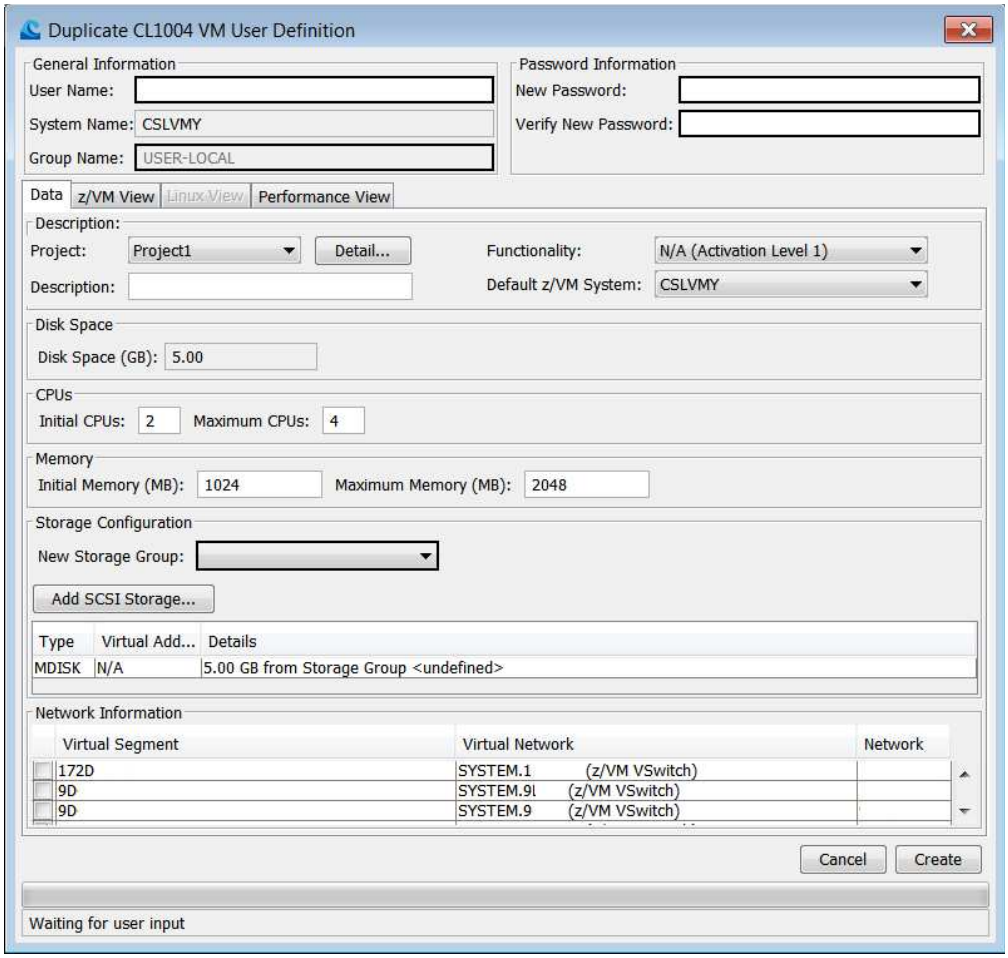

Figure 80. Duplicate z/VM User Definition

The following fields are in the **Duplicate** ***guest_name*** **z/VM User Definition**.

General Information

User Name
The new z/VM user name, which is limited to 8 characters. The name must be unique. If the user name exists, the system issues an error message.

System Name
The name of the z/VM system from which you cloned the z/VM user definition.

Group Name
The name of the site defined group (SDG) to which the original z/VM user belonged.

Password Information

New Password
A password for the new z/VM user. The password field is 8 characters.

Verify New Password
Reenter the password for confirmation.

Disk Space

Disk Space (GB)
Indicates the total disk space in GB.

CPUs

Initial CPUs
The initial number of CPUs assigned to the z/VM Guest.

Maximum CPUs
The maximum number of CPUs that are allowed to be assigned to the z/VM Guest.

Memory

Initial Memory (MB)
The initial amount of memory, specified in MB, that is allocated to the z/VM Guest.

Maximum Memory (MB)
The maximum amount of memory, specified in MB, that can be allocated to the z/VM Guest.

Storage Configuration - The Storage configuration table is populated with new storage configuration based on the source guest. If the source guest had SCSI devices, IBM Wave suggests allocating the same number of devices from the same device pool that were used for the source guest. If the device pool does not contain enough available devices, a warning message is posted at the top of the window.

New Storage Group
If you want to use storage from a different storage group, select the new storage group.

Add SCSI Storage
Add SCSI Storage is enabled when the system has FCP device pools that contain FCP devices. **Add SCSI Storage** is disabled when no FCP devices are defined.

To **Add SCSI Storage**, select a Device Pool from "Devices in Device Pool" menu to see the list of devices that are available to add to the new guest. Select one or more devices and click **Add** to add the selected device or devices to the "SCSI Storage to Add to New Guests" table. To add all of the available devices without selecting each one, click **Add All Available**.

Note: For more information about adding and removing SCSI devices, see "Adding SCSI Storage" on page 166.

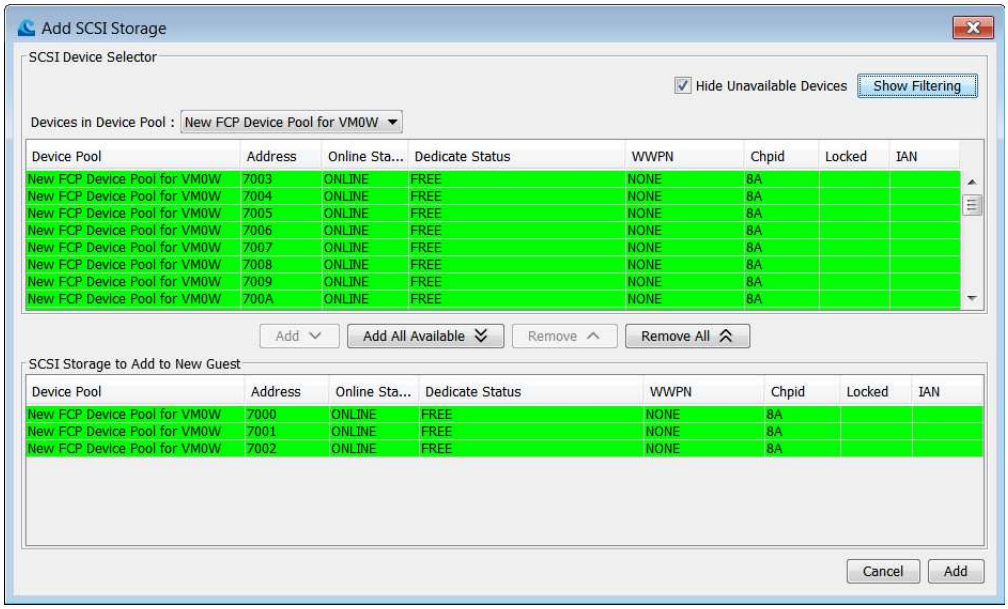

Figure 81. Add SCSI Storage

Network Information

A multiple selection table that contains network interface information for **Virtual Segment**, **Virtual Network**, and the **Network** with its associated IP address. Select the check box to view a list of **Virtual Network** choices that are available for the new user.

Important: To view the menu with its available VSwitch options, select the row in the **Network Information** multiple selection table, and then the table cell for the **Virtual Network**.

Related concepts:

"Multipath support for FCP devices" on page 97

Related information:

"Direct-attached SCSI Devices" on page 166

Small Computer System Interface (SCSI) is an ANSI standard electronic interface that allows systems to communicate with peripheral hardware, such as storage devices with faster speed and more flexibly than other interfaces. After you work with your Linux Administrator and z/VM Storage Administrator, IBM Wave can use the serial SCSI command protocol to enable you to add storage for new z/VM guests or duplicate z/VM definitions.

Installing Linux with the BMI wizard

Use the **Bare-Metal Installation (BMI) wizard** to interactively install a Linux distribution on a z/VM Guest.

Before you begin

The following prerequisites are required to use certain features of the **IBM Wave Bare-Metal Installation (BMI) Wizard**:

1. The z/VM Guest must not be active, and have at least 1 GB of storage defined.
2. A Linux repository must be defined in the **IBM Wave Linux Media Repository Manager**.
3. WWPN and LUN ID discovery requires z/VM V6R2 and above.
4. For z/VM Guests with FCP devices defined, **Step 4 - Configure SCSI Devices** is enabled for Red Hat Enterprise Linux (RHEL). You can use this step to select the SCSI paths to be visible during the installation.

   Note: SUSE Linux Enterprise Server (SLES) and Ubuntu Server dynamically discover FCP devices, so it is not necessary for you to "Configure SCSI Devices" for the SLES or Ubuntu servers.

For more information, see the following topics:

- "Display IBM Wave Linux repositories" on page 54
- "Direct-attached SCSI Devices" on page 166 and "Manipulating FCP paths" on page 95
- "Relocate to..." on page 128.

Procedure

1. Right-click an inactive Linux Guest with at least 1 GB of storage, and select **Install** > **Launch Linux Installation**. The interactive wizard for the **Step1 - Welcome to the IBM Wave Bare-Metal Installation Wizard** opens as shown in Figure 82.

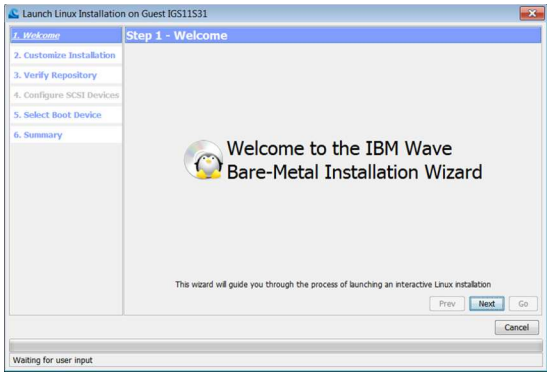

Figure 82. Step 1 - Welcome

2. Click **Next** to load the **Step 2 - Customize Installation** screen as shown in Figure 83 on page 89. Complete the mandatory fields for **Hostname** and **Service Machine password**.

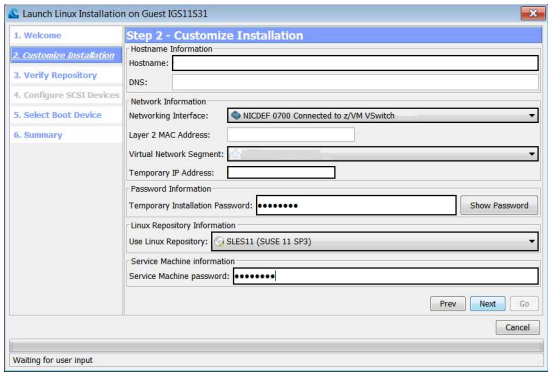

Figure 83. Step 2 - Customize Installation

IBM Wave verifies the location of the Linux media to install on the guest. You can customize the installation for the specific Linux distribution. Some of the following fields are predefined based on your choice of the z/VM Guest for the install.

- **Hostname** - Input the host name for the Linux installation. The name can be modified during the actual installation program flow. By default, IBM Wave completes the z/VM Guest name as the host name. If a default domain is defined in the IBM Wave Parameters, it is concatenated at the end of the z/VM Guest name.
- **DNS** - The domain name server (DNS) to use during the installation.
  - Red Hat Enterprise Linux (RHEL) requires the **DNS**.
  - SUSE Linux Enterprise Server (SLES) and Ubuntu Server can optionally use the **DNS**.
- **Networking interface** - The NIC address and Virtual Network name to use for the installation system.

  Note: When you select the **Networking interface**, the Virtual Network Segment (VNS) field is populated with all the VNSs that are connected to the selected virtual network. When you select a Virtual Network Segment, the **Temporary IP address** field is automatically populated with the next available IP address on that segment. If the z/VM Guest is already connected to the selected **Networking interface**, and IBM Wave knows the IP address of the guest on that connection. The values for the **Virtual Network Segment** and the **Temporary IP address** fields are automatically input.
- **Virtual Network Segment** - VNS to use for the installation. The VNS dictates some of the networking parameters such as the network, broadcast address, netmask, and default gateway.
- **Temporary IP Address** - The IP address for the installation program to use.
- **Linux Repository** - The Linux Media Repository to use for this installation. You can open the **Linux Repository** to review the settings.
- **Temporary Installation Password** - The temporary password for the installation program.
- **Service Machine password** - The service machine password.

Click **Next**. IBM Wave submits a Background Task Scheduler (BTS) request to verify the Linux repository and the media installation.

3. Use **Step 3 - Verify Repository**, as shown in Figure 84 on page 90, to confirm the existence and integrity of the files that are stored in the dedicated minidisk of the IBM Wave Service Machine.

a. Select the storage group from the **Select Storage Group** field from the **Specify Minidisk Allocation Parameters** pane.

   Note: If it is your first time working with the Linux Repository, IBM Wave must know the storage group to use for the minidisk.

b. The BTS verification request queries the Short Service Machine on the z/VM System on which the Guest is running. IBM Wave verifies that a dedicated minidisk is created for the Linux Repository.

- When the minidisk is found, the BTS request compares the size and time stamp of the Linux installation program files that are stored on the minidisk to the files in the repository. If the BTS verification finds an inconsistency with the program files, the files are sent through FTP to the minidisk.
- When the minidisk is not found, you are prompted to provide a minidisk address (automatically generated), and to supply the DASD group from which to provide the storage for the minidisk. The minidisk size is 100 cylinders.

Click **Next** to configure SCSI devices for Red Hat Enterprise Linux (RHEL).

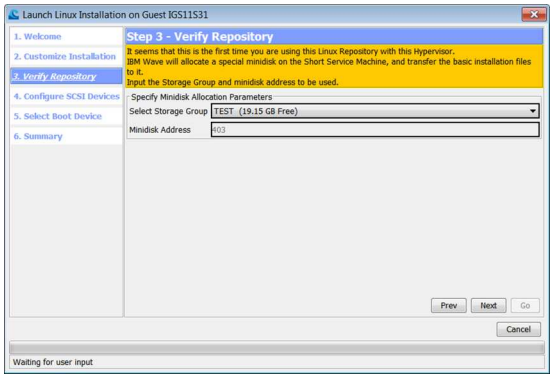

Figure 84. Step 3 - Verify Repository

4. Predefine the SCSI paths to be visible during installation by using **Step 4 - Configure SCSI Devices**, as shown in Figure 85. You must use the step when you install Red Hat Enterprise Linux (RHEL) on a guest with FCP devices attached. SLES and Ubuntu automatically discover SCSI paths, so there is no need to use **Step 4 - Configure SCSI Devices** when you install SLES or Ubuntu, or when a guest does not use FCP devices.

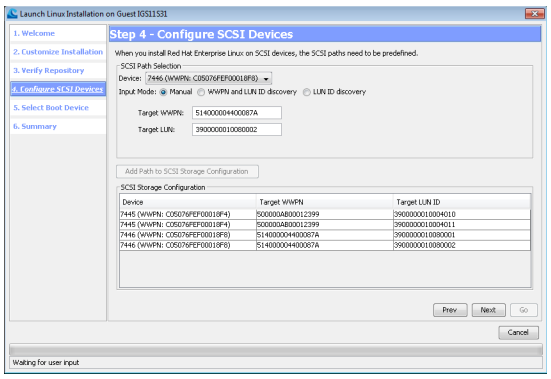

Figure 85. Step 4 - Configure SCSI Devices RHEL

a. From the **SCSI Path Selection** pane, select a "Device" and "Input Mode".

- **Manual** - When you know the target WWPN and LUN IDs, select the **Manual** option and type in the WWPN and LUN IDs.
- **WWPN and LUN ID discovery** - When you want IBM Wave to discover the WWPN and LUN ID, use the **WWPN and LUN ID discovery** option. Click **Discover WWPN and LUN IDs**.
- **LUN ID discovery** - When you know the WWPN and want to discover the LUN ID, use the **LUN ID discovery** option. Type in the WWPN, and then click **Discover LUN IDs**.

After the Target WWPN and Target LUN fields are populated, the **Add Path to SCSI Storage Configuration** is enabled.

b. Click **Add Path to SCSI Storage Configuration** to populate the **SCSI Storage Configuration** table with the storage selections.

c. Repeat **steps a** and **b** as needed to populate the **SCSI Storage Configuration** table with the SCSI paths you want to use.

Notes:

- WWPN and LUN ID discovery requires a minimum V6R2 z/VM system.
- For instructions about how to specify SCSI devices, see "Adding SCSI Storage" on page 166.

Click **Next** to **Select Boot Device**.

5. Select the boot device in **Step 5 - Select Boot Device**. Click **Use current boot device** if no changes are needed for the boot device.
   - **RHEL Configuration** - If you configured SCSI for use with RHEL, **Step 5 - Select Boot Device**, as shown in Figure 86, is populated with the selections you made in **Step 4 - Configure SCSI Devices**.
   - **SLES and Ubuntu Configuration** - If you are using SLES or Ubuntu, you can optionally configure the boot device for use with SCSI, as shown in Figure 87 on page 92.

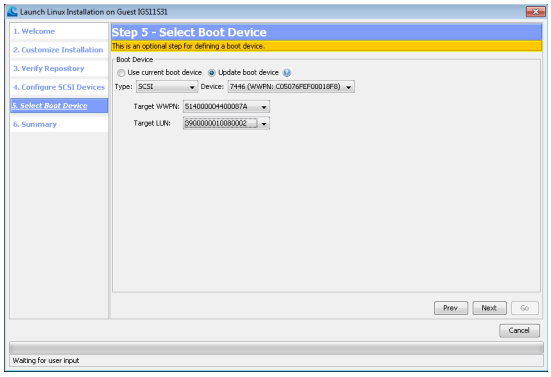

Figure 86. Step 5 - Select Boot Device RHEL

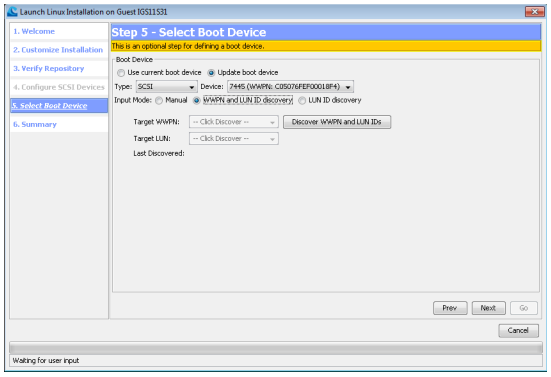

Figure 87. Step 5 - Select SCSI Boot Device: SLES and Ubuntu

You completed the boot device selection.

6. Use **Step 6 - Summary**, as shown in Figure 88, to review the choices you made in the previous steps. You can go back and modify one or more of the previous steps. The **Summary** includes information about the Linux installation, processes for creating and verifying the dedicated minidisk on the IBM Wave Service Machine, the Boot Device, and the Linux PARMFILE.
   When you are satisfied with the information in **Step 6 - Summary**, click **Go** to

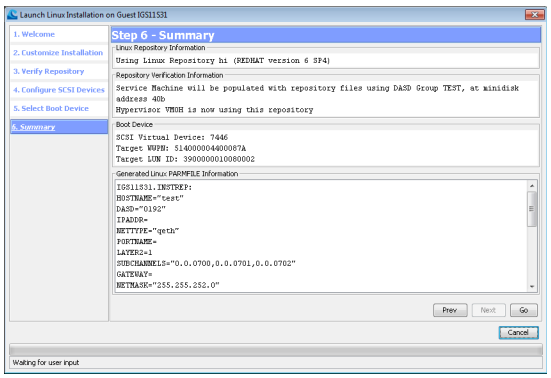

Figure 88. Step 6 - Summary

submit the BTS work units.

Results

After the BTS completes the request, the Linux Guest is ready to be initialized.

What to do next

Initialize the new Linux guest to be managed by the IBM Wave environment. See "Init Users for IBM Wave use" on page 55.

Related tasks:

"Adding SCSI Storage" on page 166
Use the following procedure to add SCSI storage.

Related information:

"Direct-attached SCSI Devices" on page 166
Small Computer System Interface (SCSI) is an ANSI standard electronic interface that allows systems to communicate with peripheral hardware, such as storage devices with faster speed and more flexibly than other interfaces. After you work with your Linux Administrator and z/VM Storage Administrator, IBM Wave can use the serial SCSI command protocol to enable you to add storage for new z/VM guests or duplicate z/VM definitions.

Manage Storage

Use the Manage Storage option to add more storage to a z/VM Linux Guest.

To **Manage Storage**, right-click a Guest, and select **More Actions** > **Manage Storage**. Use the **Manage Storage** option to add more storage to a z/VM Linux Guest. IBM Wave automatically accesses the virtual server and queries the active file systems that are mounted.

The file systems are shown in table format with the following fields:

- Linux device that is associated with the file system
- Disk space usage
- Type of the partition

  STD Standard

  LVM Logical Volume Manager
- Storage Type describes the storage medium such as Extended Count Key Data (ECKD) or DASD. For example, as shown in Figure 89, Fibre Channel Protocol (FCP) and DASD are the storage types in use. The type of an LVM partition is determined by the type of the first disk in the VG.
- Mount point indicates the directory path. IBM Wave tries to retrieve the `/etc/fstab` file. If IBM Wave cannot retrieve the file, an error message appears, and if you try to create a new partition, the "Add new file system to /etc/fstab" check box is disabled.
- Status bar that indicates the percentage of capacity in use.

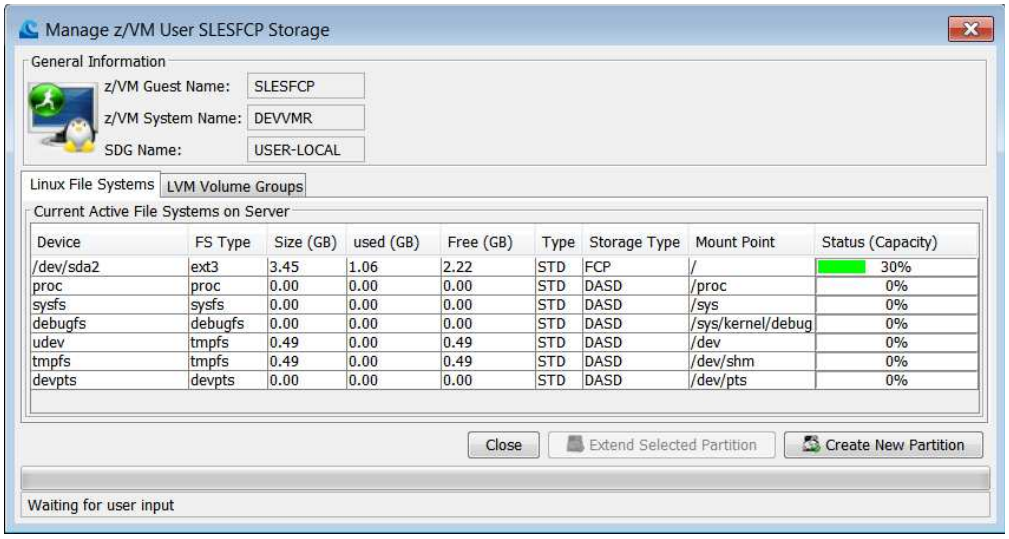

Figure 89. Manage Storage

After any configuration changes, a background task scheduler (BTS) work unit is submitted to the BTS with the relevant requests.

Create new partition

To create a new partition, click **Create new partition**. If FCP devices are available in the z/VM System, IBM Wave prompts for a selection for the medium, FCP or DASD, on which to create the new partition.

Extend selected partition

When the selected partition is type Logical Volume Manager (LVM), IBM Wave can extend the partition and the file system that is defined on it. To extend a partition, right-click on a valid entry in the **Manage Storage** table and click **Extend Selected Volume Group**.

The type of disks to add is determined by the Storage Type that was identified in the Storage Type column.

Create or extend a DASD partition

The **Create New DASD Partition** window, as shown in Figure 90, contains the following panes:

Unit Specifications
Contains the fields for specifying the amount and unit of storage to allocate and the Storage Group from which to take the storage.

Partition Definitions
Contains fields to specify the type of partition, Standard or LVM, and the file system to define on the new partition.

File System Options
Contains fields for specifying the file system type and format.

Note: If IBM Wave cannot retrieve the `/etc/fstab` file, you cannot add a mount point to the file. After you select the file system type, notice that the check box “Add new file system to /etc/fstab” is disabled.

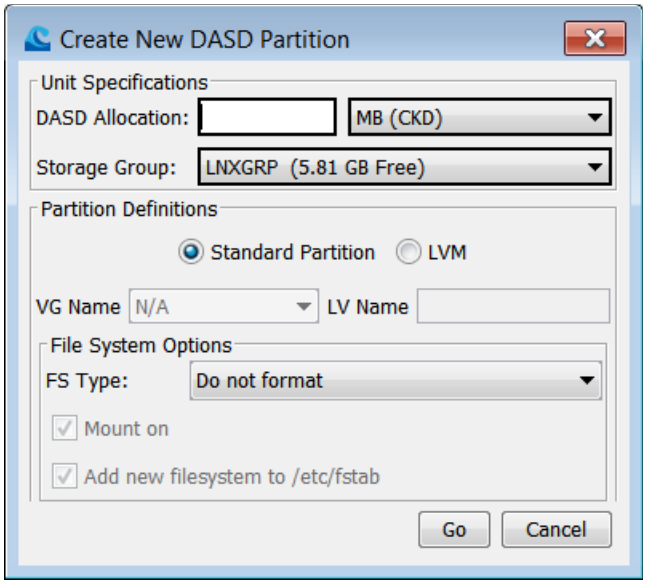

Figure 90. Create New DASD Partition

Create or extend an FCP partition

The **Create New FCP Partition** window contains several sections:

FCP Device Paths
The table contains a list of FCP paths for the new storage. To add paths, click **Add Path**. After paths are defined, they can be manipulated by right clicking on the path entry.

Partition Definitions
The pane that contains fields for specifying the type of partition, Standard or LVM, and the file system to define on the new partition.

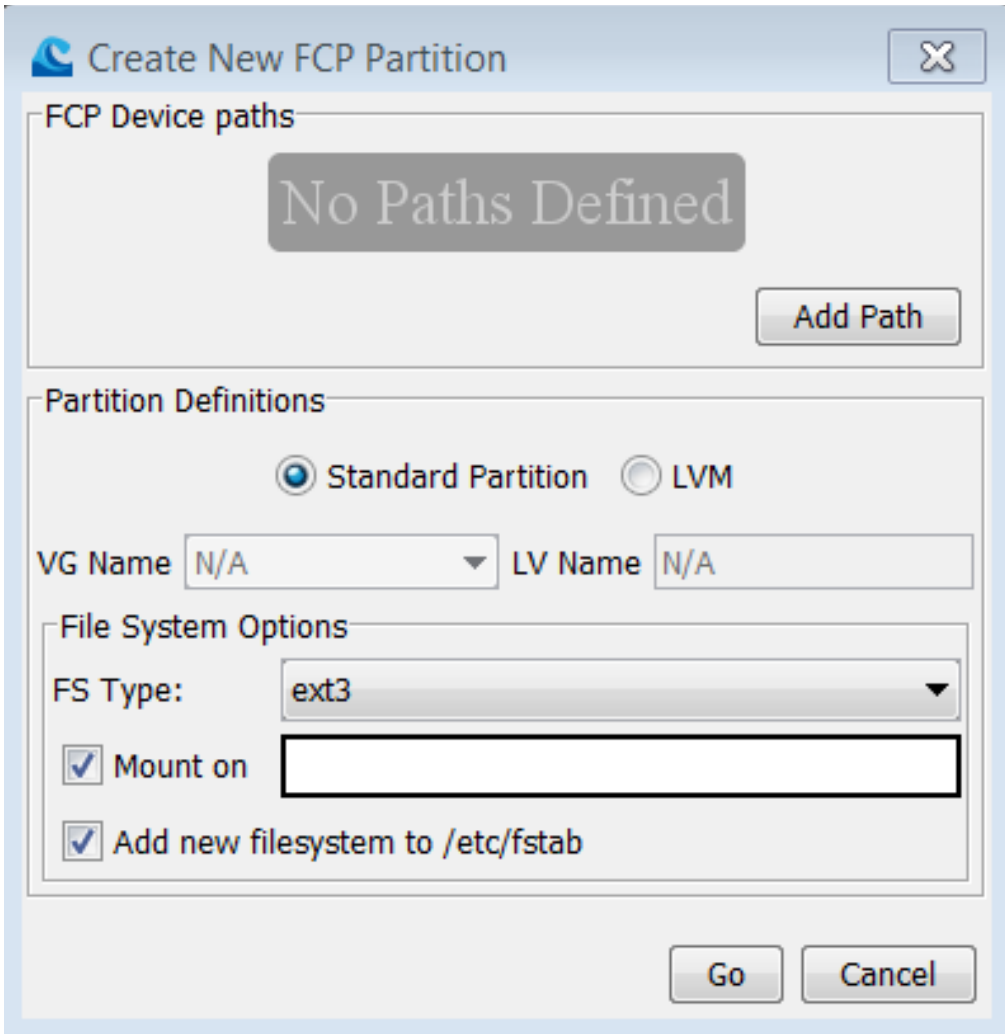

Figure 91. Create FCP Partition

Manipulating FCP paths: An FCP Path is comprised of an FCP Device, a Target WWPN and a Target LUN ID.

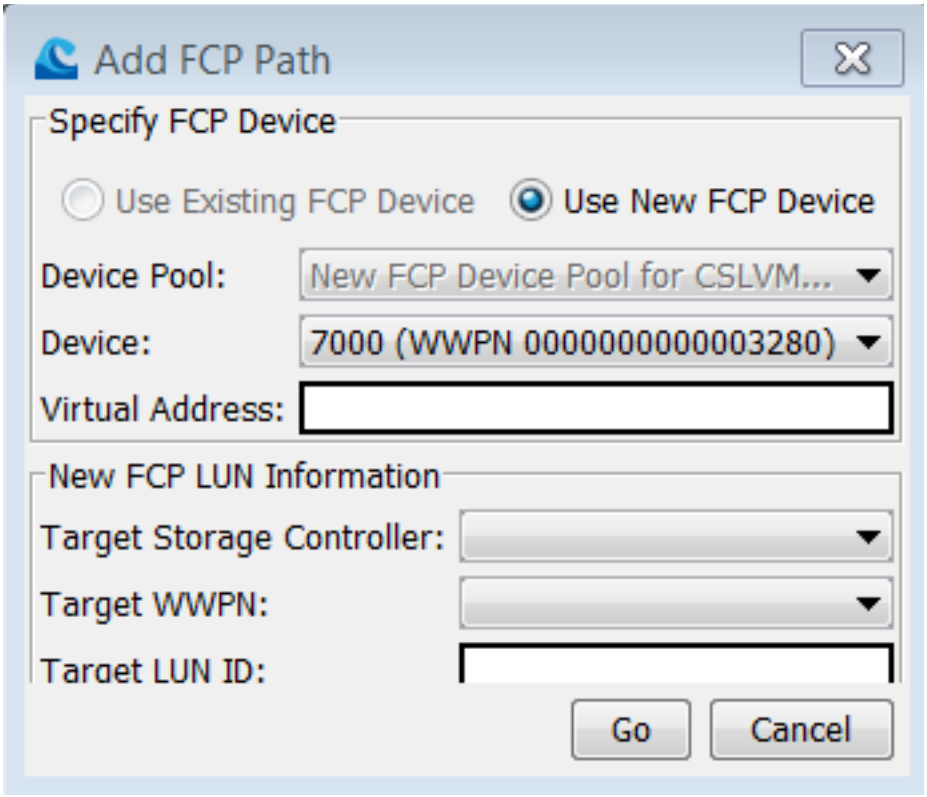

When adding/Updating FCP Paths, it is possible to select:

1. **FCP Device** - Specify whether an existing FCP device or a new FCP device should be used for a path. When specifying a new FCP Device, it is possible to select the FCP Device Pool from which to carve out the device. IBM Wave automatically suggests free FCP devices from the selected Device Pool. You can also select a different free device. For example, you require a specific source WWPN.
2. **Target Storage Controller** - Select the storage controller on which the LUN is located. Selecting a value in this field populates the Target WWPN box with values.
3. **Target WWPN** - Select the Target WWPN for the path.
4. **Target LUN** - Specify a 16 digit LUN ID for the new storage allocation.

Note:

- The Partition Definition Area is disabled when you extend an existing partition.
- IBM Wave checks if LVM is supported on the server. If not, the LVM option is disabled.

Working logic

After a new configuration is built, IBM Wave builds a Background Task Service (BTS) work unit with the appropriate BTS Requests and sends the work unit to the BTS. Depending on the configuration that is selected, the work unit can be comprised of the following requests:

Add <*Amount*><*Unit*> to Guest
: This request adds the specified amount of DASD storage to the z/VM Guest. The request handles only the z/VM side (allocates the Minidisk/Full Volume, adds it to the guest directory).

Dedicate Device <*Real Device*> to Guest
This request dedicates the specified real device to the guest. Use this request when you want to add storage type Fibre Channel Protocol (FCP) and when you want to add paths with new FCP devices.

Create New CKD Partition
This request runs after the new storage is allocated to the guest directory. The request links the new storage to the guest, brings it online, prepares a low-level format of the storage, creates a partition table, and then updates the `zipl` file when necessary.

Create New FCP Partition
This request is run as part of adding FCP storage to a guest. If a new FCP device is selected, the request attaches it to the guest, and brings it online. The request adds the Target WWPN port to the device, if necessary, and adds the LUN and updates any necessary configuration files to make the changes permanent. Depending on the distribution of Linux, these files can be the `udev.rules` files, `zfcp.conf` files, and others.

Create or Extend Logical Volume Manager (LVM) Volume Group
The request to create a new LVM volume group or extend a new one. For each new storage device added (minidisk, full volume, FCP, LUN, and others), the request creates an LVM physical volume (PV). If an existing LVM volume group (VG) is being extended, the request adds the physical volume (PV), or volumes to the VG. If a new VG is being created, it is created with the new PV.

Note: Before you use managed storage and LVM, use the filter parameter in `/etc/lvm/lvm.conf` to ensure that the devices names are not filtered out by LVM. The device names that IBM Wave uses to create LVM physical volumes depends on the device-type that is being managed:

- For multipath SCSI devices: `/dev/mapper/`.
- For DASD or non-multipath SCSI devices: `/dev/disk/by-path/`.

Create or Extend LVM Logical Volume
The request to create or extend an LVM logical volume. If you create or extend a new logical volume, extra storage is allocated for it regardless of how much storage is available in the volume group. When there is not enough space in the volume group, do not add new minidisks to the LVM.

Create New Filesystem
This request creates a new file system on the new storage device that is introduced to the guest. When relevant, the request creates a partition table on the new storage, and a low-level format is done. When the storage is ready for Linux use, the request creates the selected file system, and then, depending on the user's selection, mounts the file system and modifies the `/etc/fstab` file.

Resize Filesystem
This request is for resizing an existing file system after more storage is added to the specified medium. Some distributions of Linux do not support online resizing of file systems. When an online resize is not supported, the request attempts to bring the file system offline to resize, and then remounts the file system.

Multipath support for FCP devices

IBM Wave provides multipath support. To recognize FCP storage, the Linux administrator must configure the multipath daemon to include SCSI devices, and to start on the manage guest. The configuration ensures that when you manipulate

FCP storage, IBM Wave recognizes that a multipath configuration exists. IBM Wave then uses the multipath device name for all operations, such as addressing the storage and adding it to the `/etc/fstab` file system.

Storage restrictions

IBM Wave attempts to satisfy an IBM Wave user request, but there are cases when it is not possible:

1. When a user requests an amount of storage that requires IBM Wave to allocate more than one minidisk, and the standard partition check is selected. IBM Wave prompts the user with a warning and offers to define a new LVM group and assign the newly defined storage to that LVM Group. LVM logical volumes can then be defined on the LVM group.
2. When a user requests to extend the root file system. The extend FS action is unavailable for the root file system (a file system that IBM Wave detects as mounted on "/"). The btrfs sub volume, which is a child of the root file system, is also unavailable for the extend FS action.
3. When an IBM Wave User asked to extend a file system that cannot be extended online. IBM Wave attempts to unmount, resize, and then mount the partition. If the partition is in use and cannot be unmounted, the extend file system action fails.

Notes: The number of minidisks added to Linux depends on the following factors:

1. The amount of storage requested.
2. The physical size of the DASD disks that are in the DASD group.
3. The maximum contiguous span of available DASD.

Examples:

1. An IBM Wave user requests 4 GB of storage. The DASD is model 3 (~2.67 GB disk), with a DASD group that contains 2 physical DASD disks that each have a 2 GB contiguous spans. IBM Wave then allocates the two minidisks with 4 GB of storage.
2. An IBM Wave user requests 1 GB of storage, and the DASD is model 3 (~2.67 GB disk) with a DASD group that contains 2 physical DASD disks, but neither disk has a 1 GB contiguous span. IBM Wave attempts to *assemble* 1 GB storage out of the maximum contiguous spans of the two disks. Disk A has 500 MB of contiguous space, and disk B has 650 MB of contiguous space. IBM Wave takes the 650 MB from disk B and the remaining 350 MB from disk A. IBM Wave then allocates the two minidisks with 1 GB of storage.

Read, Update, or Remove an IAN

Use the Read, Update, or Remove action on the IAN attached to the z/VM Guest.

Remove from DB

This action applies for inconsistent z/VM Guests only. Use this action to delete the z/VM user entry from the IBM Wave database.

SSH Access

IBM Wave contains an internal Secure Shell (SSH) engine to provide standard SSH access to z/VM Linux Users with an SSH server that listens on the default SSH port. The SSH interface is standard with bash auto complete capabilities. You can also define an external SSH application to be used with IBM Wave.

To open an SSH connection to a Linux user, right-click on the user that you want to connect, and then click **Access...** > **SSH Access**.

Figure 92. SSH Linux user access

The SSH access window, as shown Figure 92, contains the following panes:

SSH Login Credentials

After entering the **User Name**, you can authenticate in one of two ways:

1. By using a **Password**.
2. By specifying the location of the **Key File Path**, and **Passphrase**, if needed.

Connection Options

Interface - Select the network interface IP address.

Set this connection as default - Select the check box to set the interface as the default connection for the Linux user.

Notes:

1. If the private key file contains an encrypted private key, the external SSH application is responsible for decrypting the private key with a passphrase.
2. Only users who have “Update” permissions can change the guest connection priority.

For more information about the external SSH application, see “Security parameters” and “Linux Login Security options” in .

Status or Performance View

There are multiple ways to see the Performance View for any type of guest.

To open the **Performance View** information for a Linux Guest, select the Linux Guest, right-click and select **Status**. For all guests, double-click a guest to open the **Display Details** view. The **Display Details** window contains the **Performance View** tab.

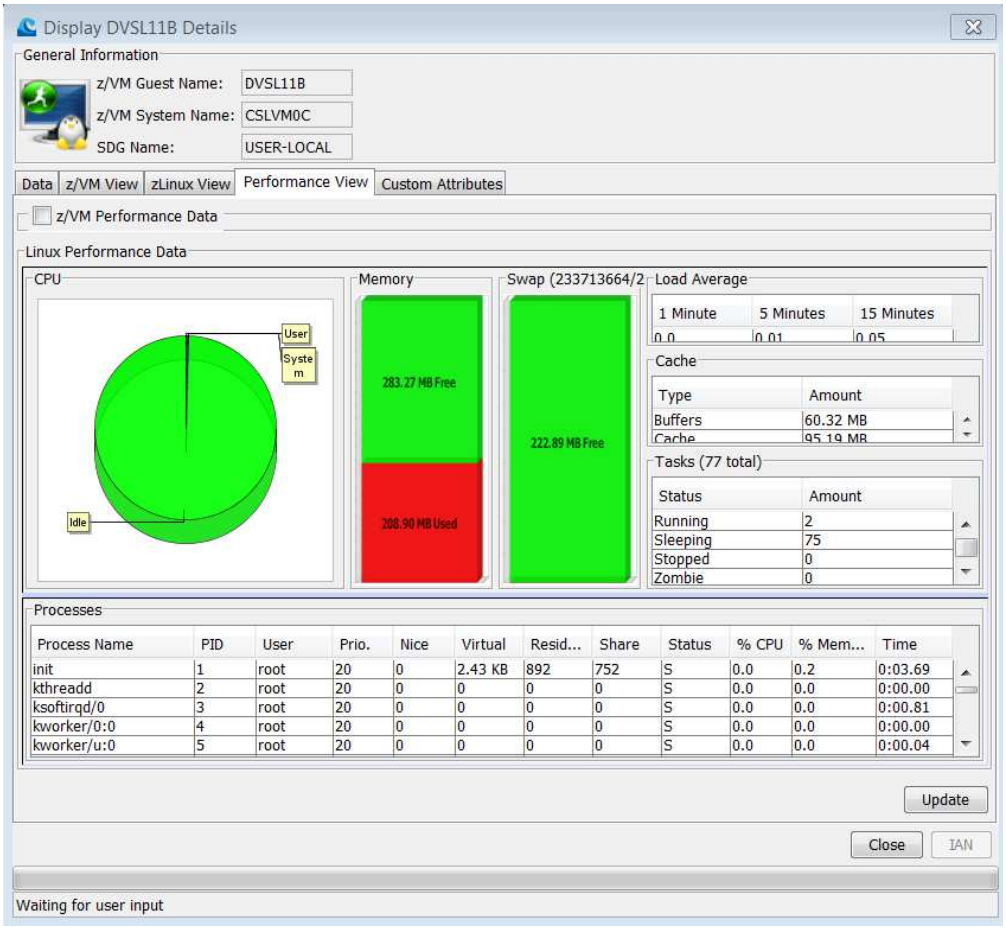

Figure 93. Performance View: Linux Guest

For any z/VM Guest that is active, z/VM performance data (retrieved from the INDICATE command) is displayed in the top portion of the window. For active z/VM Guests that run Linux and are inited-for-IBM Wave, the **Performance View** also contains performance information from Linux.

Fields in the z/VM Performance portion of the window:

- **Machine Type** - z/VM machine type.
- **IPL device** - V. Address of the IPL device.
- **Storage Type** - Mapped to Virtual/Real.
- **Configured Storage** - Available to the z/VM Guest.
- **XStore** - Amount of extended memory.
- **RES** - Number of resident pages in memory.
- **WS** - Working set size.
- **RESVD** - Number of reserved pages in memory.
- **LOCK** - Number of locked pages in memory.
- **PREF** - Compatibility value (always 0).
- **NPREF** - Number of private address space pages owned by this user that are currently residing on paging space.
- **READS** - Number of reads from expanded storage to real storage.
- **WRITES** - Number of writes to expanded storage from real storage The Linux view displays a "top" output.

Click **Refresh** to refresh the current tab's display.

Update z/VM User

To update a z/VM guest right-click on the guest, and then click **Update** > **Update Information**.

When you update a CMS z/VM Guest, the "Router Information tab" on page 104 becomes visible.

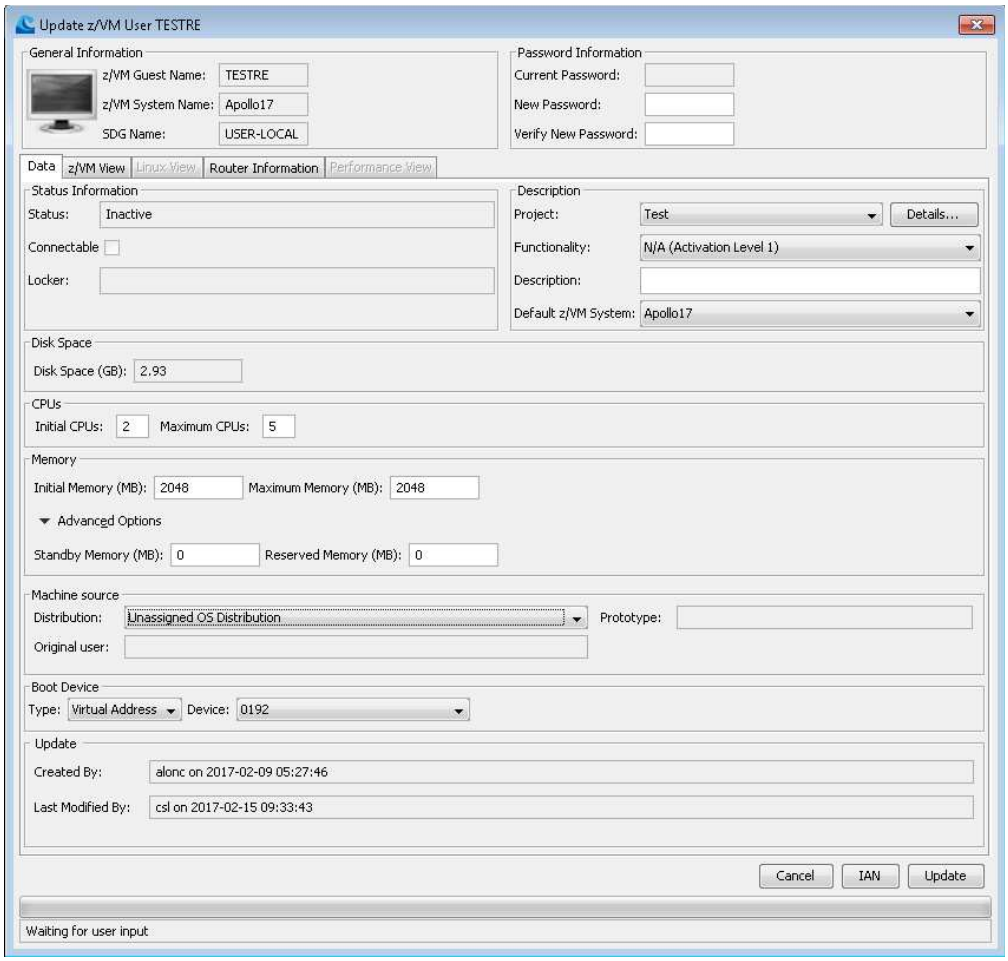

Figure 94. Update z/VM User

The **Update z/VM User** window contains a subset of fields and tables that can be updated.

New Password
The password to update z/VM user. The password field is 8 characters.

Verify New Password
Reenter the password for confirmation.

The **Data** tab contains the following panes:

Status Information

Status
Indicates if the guest is **Active** or **Inactive**.

Connectable
Indicates when at least one network interface is connectable.

Locker
Indicates if there is an enqueue (lock) held for the user.

Description

Project
The **Project** to which you assign the updated user. For more information, see "Assign to Project" on page 72.

Functionality
The optional activation level of the user. For more information, see "Activate" on page 110.

Description
The optional description field for the user.

Default z/VM system
The default z/VM System to assign to the updated guest.

Disk Space

Disk Space
The total amount of disk space available in GB.

CPUs

Initial CPUs
The initial number of CPUs assigned to the z/VM guest.

Maximum CPUs
The maximum number of CPUs that can be assigned to the z/VM guest.

Memory

Initial Memory (MB)
The initial amount of memory, specified in MB, that is allocated to the z/VM guest.

Maximum Memory (MB)
The maximum amount of memory, specified in MB, that can be allocated to the z/VM Guest.

Advanced Options (for Memory) - controls adding memory dynamically. To add memory dynamically, you must set Standby Memory or Reserved Memory. The options are collapsed by default, but shown in Figure 94 on page 101 for illustrative purposes.

Standby Memory
The amount of standby memory.

Reserved Memory
The amount of reserved Memory

Important: Dynamic Memory reconfiguration is only supported when the initial memory size for the guest is an exact multiple of the memory block size. To understand the calculation of memory block size, see the following topics:

- "Overview of estimating memory and CPU requirements" at www.ibm.com/support/knowledgecenter/SSB27U_6.2.0/com.ibm.zvm.v620.hcpl0/memcpuov.htm.
- "Steps for estimating memory and CPU requirements" at http://www.ibm.com/support/knowledgecenter/SSB27U_5.4.0/com.ibm.zvm.v54.hcpl0/hcsx0b3030.htm.

Machine Source

Distribution
The operating system that the z/VM guest is running.

Prototype
The name of the prototype if applicable.

Original Source
The name of the original guest if applicable.

The **Boot Device** that is selected depends on the storage configuration for the guest you are updating. For example, as shown in Figure 95, the **Boot Device** is SCSI.

Figure 95. SCSI Boot Device

The following **Boot Device Type** options are available:

None No boot device is configured.

CMS The boot device is CMS. The directory entry automatically contains an `IPL CMS` statement, which means that the guest automatically boots into the CMS operating system when it is activated.

Virtual Device
The boot device is a Virtual Device. The device address (four hexadecimal digits) specifies the virtual device to use or to boot from.

SCSI The boot device is SCSI. After you select the virtual device, the following options are available:

- **Manual**: When you know the target WWPN and LUN IDs, select the **Manual** option and type in the WWPN and LUN IDs.
- **WWPN and LUN ID discovery**: When you want IBM Wave to discover the WWPN and LUN ID, use the **WWPN and LUN ID discovery** option. Click **Discover WWPN and LUN IDs**.
- **LUN ID discovery**: When you know the WWPN and want to discover only the LUN ID, use the **LUN ID discovery** option. Type in the WWPN, and then click **Discover LUN IDs**.

Note: WWPN and LUN ID discovery requires:

- z/VM V6R2 and above.
- The guest must be inactive.

When you are done updating the user, click **Update** to make the changes.

Router Information tab

When you update a CMS z/VM Guest, the **Router Information** tab becomes visible with the following options:

Use this user as router
Select the check box to indicate that IBM Wave can offer the guest as a router user when creating Guest LANs.

Router Configuration File Machine Name
The z/VM User name on which the TCPIP.PROFILE file is stored.

Router Configuration File Minidisk Address
The minidisk address on which the PROFILE file is stored.

Router Configuration File Name
The PROFILE file name.

When you are done updating the routing information, click **Update**.

For more information, see the following topics:

- "Adding SCSI Storage" on page 166.
- For additional field descriptions, see "Create New z/VM Guest" on page 72.
- To duplicate an existing user, see "Duplicate z/VM User Definition" on page 85.

IBM Wave 3270 Linux console

IBM Wave contains in integrated 3270 interface for you to interact with the z/VM Guest on a basic 3270 level.

To access a z/VM Guest through a 3270 console, select the guest and then click **Access > IBM Wave 3270 Linux console**.

For the IBM Wave service machines (WAVESRV, WAVEWRKS, WAVEWRKL), the 3270 window does not provide a command line.

There cannot be more than one 3270 session per a z/VM Guest, so the 3270 access uses the Wave resource serialization (WRS) mechanism. When a 3270 session is already open to the same z/VM Guest by another IBM Wave User, an error message displays.

IBM Wave also checks whether there are other 3270 sessions open for the z/VM Guest (outside of IBM Wave). If the z/VM Guest is already in a session, an error message displays.

Tip: When you use the 3270 console, you can recall commands that were previously entered by using the up and down arrows.

IBM Wave 3270 Linux display only console

This option is the same as "IBM Wave 3270 Linux console," except that the Linux Console does not allow commands to be sent.

Multiple task actions

Multiple task actions are actions that can be performed on one or more z/VM Guests.

Overview of multiple task actions

Multiple task actions are actions that you can run on one or more objects simultaneously.

Multiple task actions are presented to you with uniform appearance and function. Multiple task actions are dispatched in parallel on selected parallel virtual servers, according to the parallel level defined in the window. The default parallel level is the number of objects you select to run the action against. If only one object is selected in the action, **Parallel** is disabled. Parallel execution depends on the following functions:

- z/VM API level.
- Number of directory manager service machines configured for your installation.
- Number of API Worker machines configured for your installation.
- Multitasking capability of the client workstation.
- The API server must be a multitasking server.

z/VM User Activation, in Figure 96, is an example of a multiple task action.

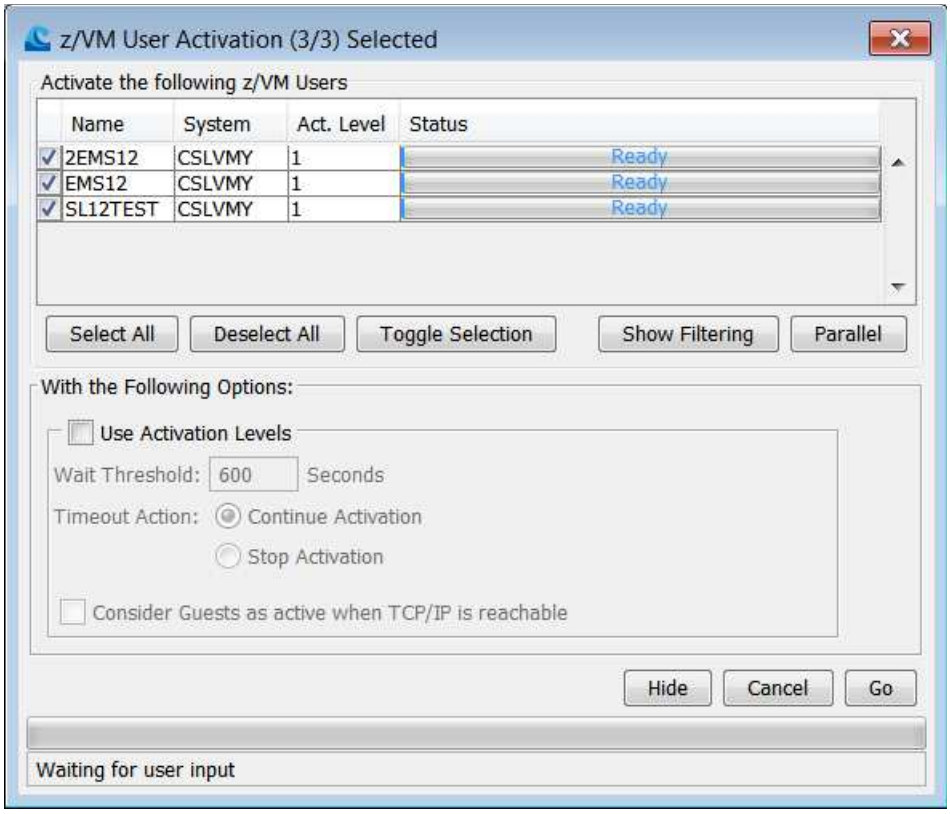

Figure 96. Example of a multiple task action

All the selected z/VM Guests on which the action is run are listed in the main pane. The action is dispatched on the z/VM Guests that are selected. You can use

the check box to select or clear z/VM Guests from the list. The status field indicates the status and progress of the action for each individual guest. The following messages are some of the possible status indicators:

- **Ready** - Indicates that the action did not happen yet. IBM Wave is waiting for you to click **Go**.
- **<Current Action>** - Indicates that the task is running the specified action.
- **Done** - Indicates that the actions finished successfully for a specific z/VM Guest.
- **Skipped + <reason>** - Indicates that the z/VM Guest was skipped with the **<reason>**, and the action is not completed.
- **Error + <Error>** - Indicates that the action ended in error for the z/VM Guest.
- **Pending Available Executor** - Indicates the parallel level that is defined for the action is lower than the number of objects on which to run the action. The message also indicates that all of the executors are processing other objects and the action for the object is pending an available executor.

Use the following options on the bottom of the pane to ease the selection process:

- **Select All** - Use to switch the check box to check for all z/VM Guests in the action.
- **Deselect All** - Use to switch the check box to cleared for all z/VM Guests in the action.
- **Toggle Selection** - Use to toggle between selecting and clearing the check boxes.
- **Show Filtering** - Use to open the filtering pane below the main table. Using the filters, you can narrow down the selected objects.
- **Parallel** - Use to open up the parallel level input menu. Valid values for the parallel level are **1 -<number of objects>**.

Close the window to exit the action. Click **Hide** to minimize the task to the **Session Tasks Viewer** window. The status and progress of the multiple task action can be monitored on the **Session Tasks Viewer**. To display the action, double-click the corresponding entry in the **Session Tasks Viewer** or right-click and select **Expand**.

For more information, see “Session Tasks Viewer” on page 40.

To completely remove the multiple-task action from the **Session Tasks Viewer**, click **Cancel** in the action window.

Add CPUs Dynamically

Use **Add CPUs Dynamically** to increase CPUs for one or more z/VM guests without recycling the guest.

Before you begin

To add CPUs dynamically, check the following conditions:

- Right-click on the guest, and then click **Update**. If the **Update** > **Add CPUs Dynamically** task is enabled, you can follow the procedure. If you cannot access **Update** or **Add CPUs Dynamically**, ensure that you have the proper permissions.
- To **Add CPUs Dynamically**, the **Directory Maximum CPUs** value (the persistent value for the MACHINE statement in the z/VM User Directory) must be set. Right-click on the guest, and the click **Update** > **Update Information** to verify the value is set.

- If you plan to save the updated configuration state as "Persistent Across Guest Restart," you must have **Update** authority to the z/VM user directory for each of the selected guests.

About this task

Use this task to dynamically increase CPUs for one or more z/VM guests without recycling the guest.

Note: There is hover help in the **Add CPUs Dynamically** pane. Move your cursor over the cell to get help.

Procedure

1. Right-click on one or more active Linux guests, and then select **Update** > **Add CPUs Dynamically**. The **Add CPUs Dynamically** window opens as shown in Figure 97.

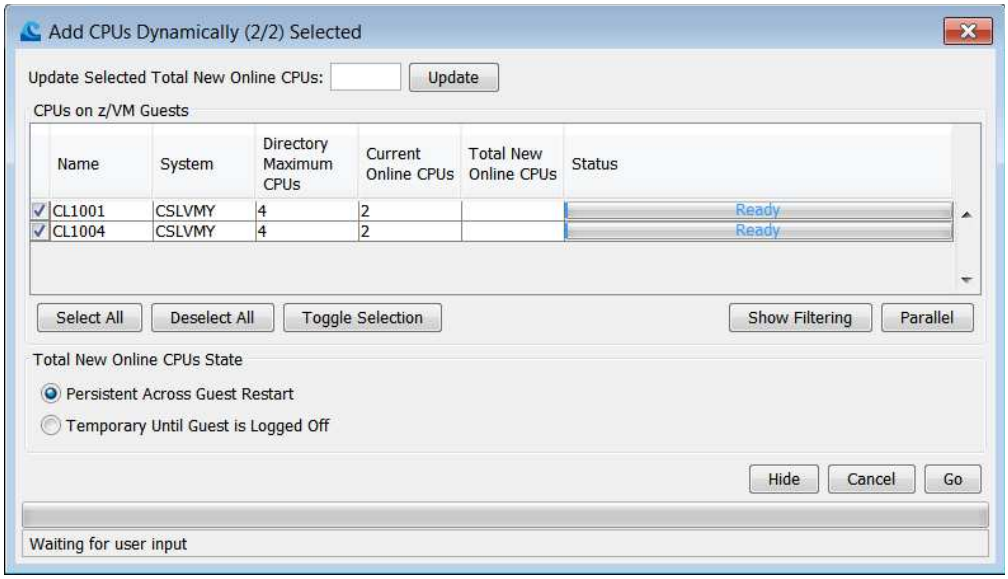

Figure 97. Add CPUs Dynamically

   If any problems exist with the statements that are defined in the z/VM directory, an error message displays in the **Status** column (where **Ready** is displayed in Figure 97). The following common errors can occur:
   - If IBM Wave detects one or more guests lack the **Update** authority to save to the z/VM Directory, no CPU increase, persistent or temporary, is saved for those guests. After you click **Go**, Wave issues an error message failing the persistent increase for the guests that lack authority.
   - If the z/VM MACHINE statement does not exist, which is possible when the guest is created or updated outside of IBM Wave, the **Directory Maximum CPUs** cell is blank.
   - If IBM Wave detects that the guest is in use, you must either unlock the user, stop the process that is holding the enqueue, or wait for the action to complete.
2. Specify the number for the **Total New Online CPUs** for all of the selected guests. In the **Update Selected Total New Online CPUs** field, enter the number for the **Total New Online CPUs** and then click **Update**.

- The number you enter must be less than or equal to the **Directory Maximum CPUs**, but greater than the **Current Online CPUs**.
- The **Total New Online CPUs** is an editable cell for each guest.

For example, you want ten guests to have 3 **Total New Online CPUs**, but "GuestZ" needs 4 CPUs. Enter **3** in the **Update Selected Total New Online CPUs** field, and then click **Update**. After the **Total New Online CPUs** cells are populated, double-click to change the cell for "GuestZ" to have **4** CPUs.

3. Click **Update** to populate the **Total New Online CPUs** column. If a problem exists with the values, IBM Wave issues an error message explaining the problem. You must either correct the problem, or clear the guest's check box to remove it from the process.
4. The radio buttons for the **Total New Online CPUs State** are based on your permissions.
   - If you have both **Add CPUs Dynamically** and **Update** authority for all of the selected guests, the default state for **Total New Online CPUs State** is "Persistent Across Guest Restart".
   - If you do not have **Update** authority, the default state is "Temporary Until Guest is Logged Off".
   - If your permissions are mixed, you can clear the guest's check box to remove it from the process. For example, you do not have **Update** authority for "GuestZ", but you have **Add CPUs Dynamically** and **Update** authority for all of the other guests. Clear the check box for "GuestZ" to successfully update the other guests.
5. When you are satisfied with the **Total New Online CPUs** values, the appropriate state is selected, and any guests with errors are cleared from the process, click **Go**.

Results

The request is added to the Background Task Scheduler (BTS). After the request completes, you are finished increasing the CPUs dynamically (without any need to recycle the guest).

Adding Memory Dynamically

Use the **Add Memory Dynamically** action to increase the memory for one or more active z/VM guests without recycling the guest.

Before you begin

To **Add Memory Dynamically**, the **Standby Memory** or the **Reserved Memory** values must be set up. To check, right-click on the guest, and then click **Update** > **Update Information** to verify that the values are set. If not, update the **Thresholds and Defaults** parameters.

Important: Dynamic Memory reconfiguration is only supported when the initial memory size for the guest is an exact multiple of the memory block size. To understand the calculation of memory block sizes, see the following topics for more information:

- The "**DEFINE STORAGE**" command at www.ibm.com/support/knowledgecenter/SSB27U_6.3.0/com.ibm.zvm.v630.hcpb7/defstor.htm
- "Overview of estimating memory and CPU requirements" at www.ibm.com/support/knowledgecenter/SSB27U_6.2.0/com.ibm.zvm.v620.hcpl0/memcpuov.htm.

- "Steps for estimating memory and CPU requirements" at www.ibm.com/support/knowledgecenter/SSB27U_5.4.0/com.ibm.zvm.v54.hcpl0/hcsx0b3030.htm.

About this task

Use this task to dynamically increase memory for one or more z/VM guests without recycling the guest.

Procedure

1. Right-click on one or more active Linux guests, and then select **Update** > **Add Memory Dynamically**. The **Add Memory Dynamically** window opens as shown in Figure 98.

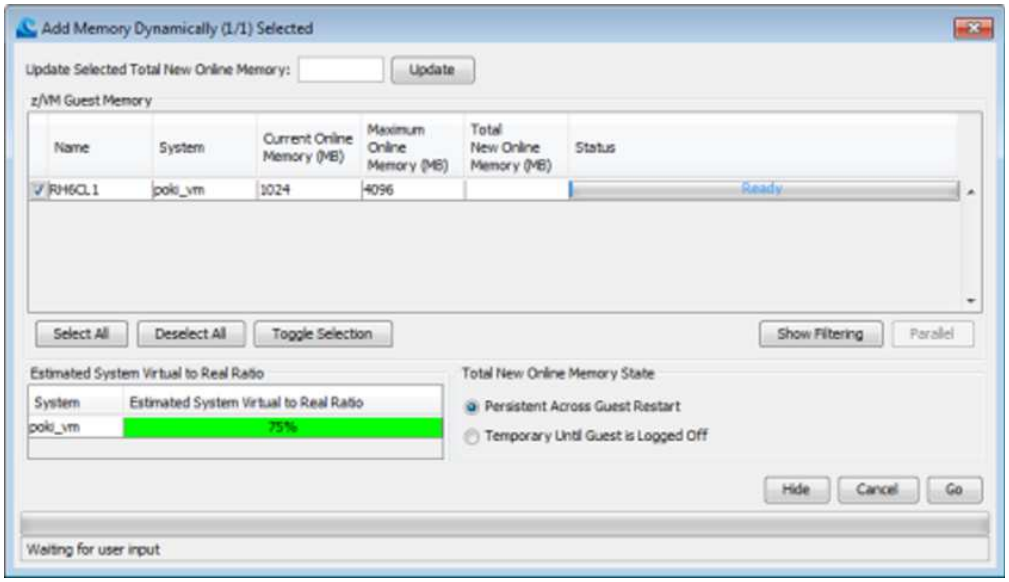

Figure 98. Add Memory Dynamically

   In the **Update Selected Total New Online Memory** field, enter a value that is greater than the Current Online Memory (1024 in the example), but less than Maximum Online Memory (4096 in the example).
2. Configure all of the selected guests at one time, by entering the value in the **Update selected New Online Memory** field.
   - The **Total New Online Memory** is an editable cell for each guest. You can also specify the exact value that you want directly in the **Total New Online Memory (MB)** field for one or more selected of the selected guests.
   - The number that you enter must be less than or equal to the **Directory Maximum Memory**, but greater than the **Current Online Memory**.
3. Click **Update** to populate the **Total New Online Memory** column. If a problem exists with the values, IBM Wave issues an error message that explains the problem. You must either correct the problem, or clear the guest's check box to remove it from the process.
4. The radio buttons for the **Total New Online Memory State** are based on your permissions.
   - If you have both **Add Memory Dynamically** and **Update** authority for all of the selected guests, the default state for **Total New Online Memory State** is "Persistent Across Guest Restart".

- If you do not have **Update** authority, the default state is “Temporary Until Guest is Logged Off”.
- If your permissions are mixed, you can clear the guest's check box to remove it from the process. For example, you do not have **Update** authority for “RH6CLI”, but you do have **Add Memory Dynamically** and **Update** authority for other guests. Clear the check box for “RH6CLI” and the you can successfully proceed to update the other guests.

5. When you are satisfied with the **Total New Online Memory** values, the appropriate state is selected, and any guests with errors are cleared from the process, click **Go**.

Results

The request is added to the Background Task Scheduler (BTS). After the request completes, you are finished dynamically increasing the memory (without the need to recycle the guest).

Activate

z/VM User Activation can start one or more guests or multiple guests in a specific order.

Use the **z/VM User Activation** process to activate one or more z/VM Guests in an orderly manner.

Note: Activation levels are processed in numeric order beginning with 1. For example, activation level 1 is processed before activation level 2.

Important: IBM Wave issues a warning message if you try to activate a guest on a z/VM system that is not its default.

The **z/VM User Activation** process is equivalent to turning on the power to a physical server. The activation task uses an API call to the z/VM XAUTOLOG facility to activate the z/VM Guest.

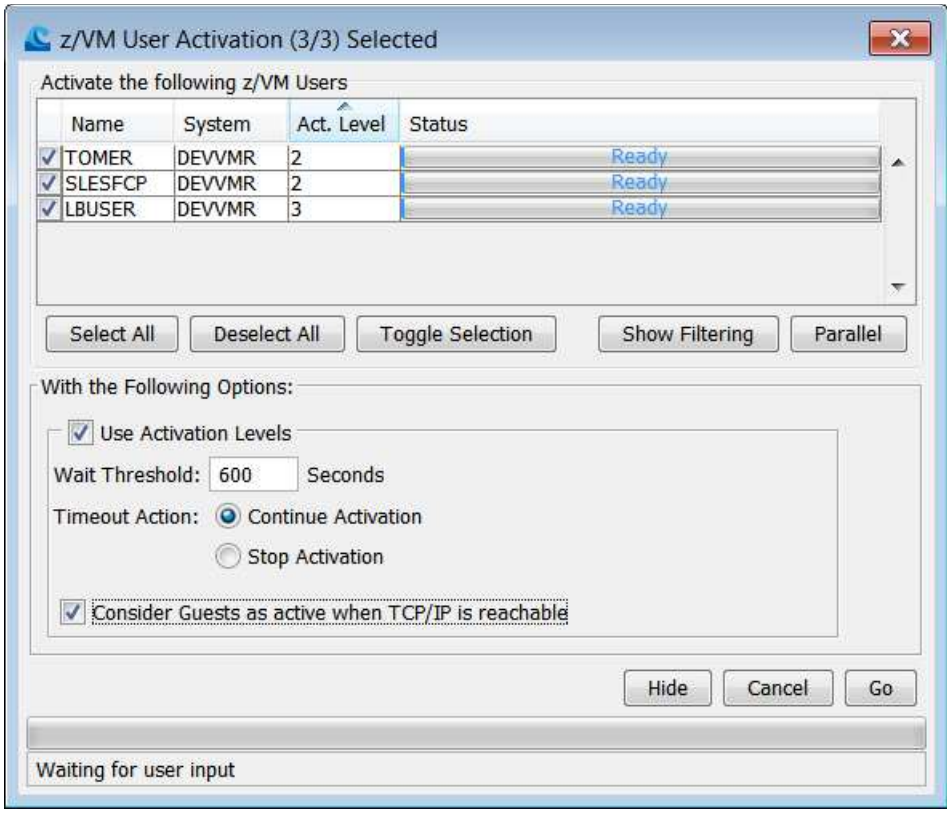

Figure 99. z/VM User Activation window

If the **Activation Level** column specifies an Activation Level, it is associated with the Functionality Type assigned to the z/VM Guest.

Note: Using **Activation Level** does not necessarily mean that the operating system under the z/VM Guest can boot. The boot process depends on the IPL statement of the z/VM Guest, and the contents of the `PROFILE EXEC` file.

When you click **Go**, IBM Wave tries to activate all of the selected z/VM Guests in the table. You can clear the check box to remove a guest from activation.

With the Following Options pane: When you check **Use Activation Levels**, you can specify the following options:

1. The "Wait Threshold " for the activation. The "Wait Threshold " parameter specifies the amount of time (in seconds) that IBM Wave must wait for the guest to become active. Under the "Wait Threshold " conditions, a z/VM Guest is considered active when one of the following actions occur:
   a. An activation signal is sent from the z/VM Guest to IBM Wave.
   b. When "Consider Guests as active when TCP/IP is reachable" is selected, the z/VM Guest is considered active when IBM Wave connects to one of its TCP/IP interfaces.

      Note: A z/VM Guest is considered active when condition a or b is met regardless of the order.
2. Use the "Timeout action" parameter to instruct IBM Wave to take action when the timeout threshold is reached. The following actions are possible:
   - Stop the entire activation process.

- Continue the entire activation process.

When you check “Use Activation Levels”, IBM Wave activates the z/VM Guests in the list according to the activation level order. In the following example, the activation process occurs with the default “600 seconds” and “Continue activation”. As shown in, Figure 99 on page 111, the guests are assigned the following activation levels:

- **Tomer** uses Activation Level 2
- **SLESFCP** uses Activation Level 2
- **LBUSER** uses Activation Level 3

1. Because z/VM Guest name **Tomer** and **SLESFCP** have the same activation level, they are activated as a group in parallel. IBM Wave waits 10 minutes (600 seconds), or until “Activation Done” signals are sent from both z/VM Guests, whichever comes first.
2. The z/VM Guest **LBUSER** is activated. IBM Wave waits for 10 minutes (600 seconds), or until an “Activation Done” signal is sent from the z/VM Guest, whichever comes first.

When multiple z/VM Guests are assigned the same activation level, they are activated as a group in parallel. IBM Wave waits for all of the z/VM Guests to be active before it moves to the next activation level.

A BTS work unit is generated for the entire process. The work unit contains six requests (an activation request, and an activation wait request for each of the z/VM guests). The BTS work unit is marked as done only after all three of the z/VM Guests are activated and considered active by IBM Wave.

For information with examples about working activation levels and functionality types, see the topic about “Functionality types, activation levels, and activation done signaling” in .

Assign account

Use this action to assign a z/VM Account to one or more z/VM Guests.

The top section of the window allows the user to select the target z/VM Account to assign.

Click **Go** to submit a BTS work unit and the request that assigns the account to the guests.

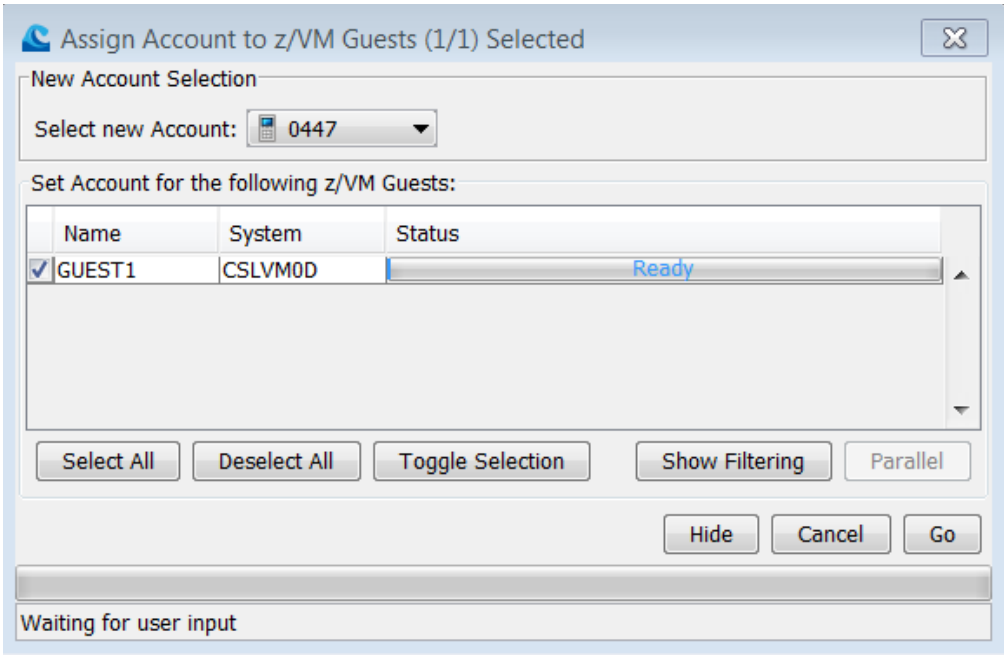

Note: It is only possible to assign predefined z/VM Accounts, which are managed through the **z/VM Account Manager**. For more information, see **Administrative** > **Site Management** > **z/VM Account Manager** in .

Assign custom attribute

Use this action to assign one or more z/VM Guests a value for a particular Custom Attribute. Select the attribute name from the **Attribute Name** menu, and then the value for the attribute to be assigned to the guests from the **Attribute Value** menu.

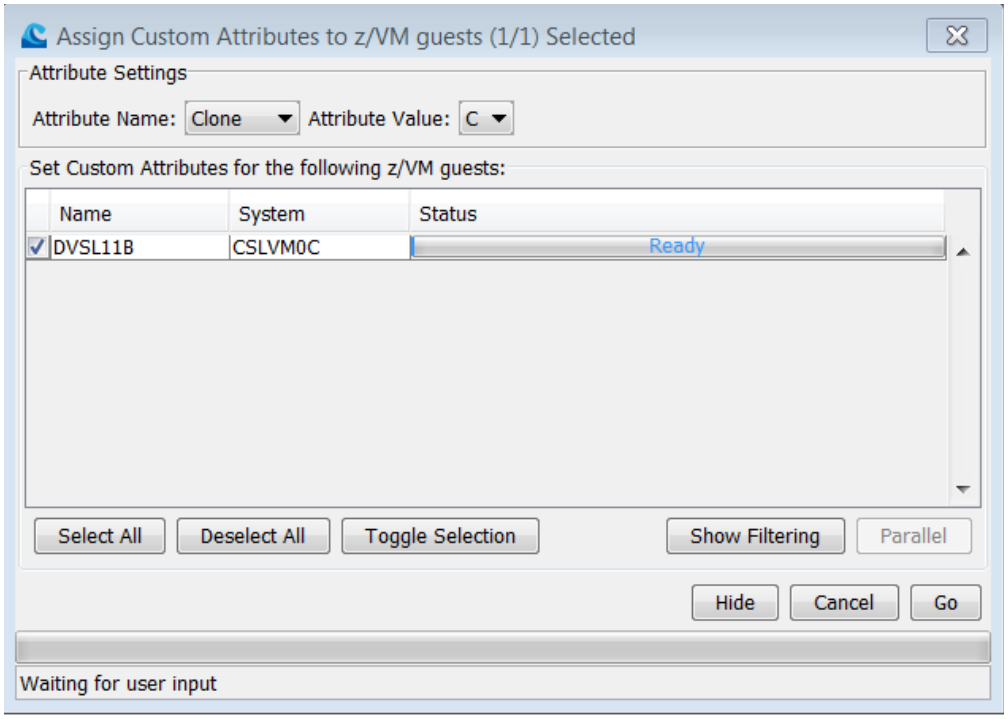

Figure 100. Assign custom attribute

Press **Go** to generate a BTS work unit that is submitted to the BTS. The work unit contains **Assign Custom Attribute** requests for each guest included in the action. For more details, see the topic about "Creating custom attributes."

Creating custom attributes

Create custom attributes for use in your IBM Wave environment.

About this task

Use the Custom Attribute Manager to define custom attributes that are specific to the objects in your environment. You can use the defined attributes to classify z/VM guests. To open the **Custom Attribute Manager**, from the main IBM Wave menu select **Administrative > Site Management > Custom Attribute Manager**.

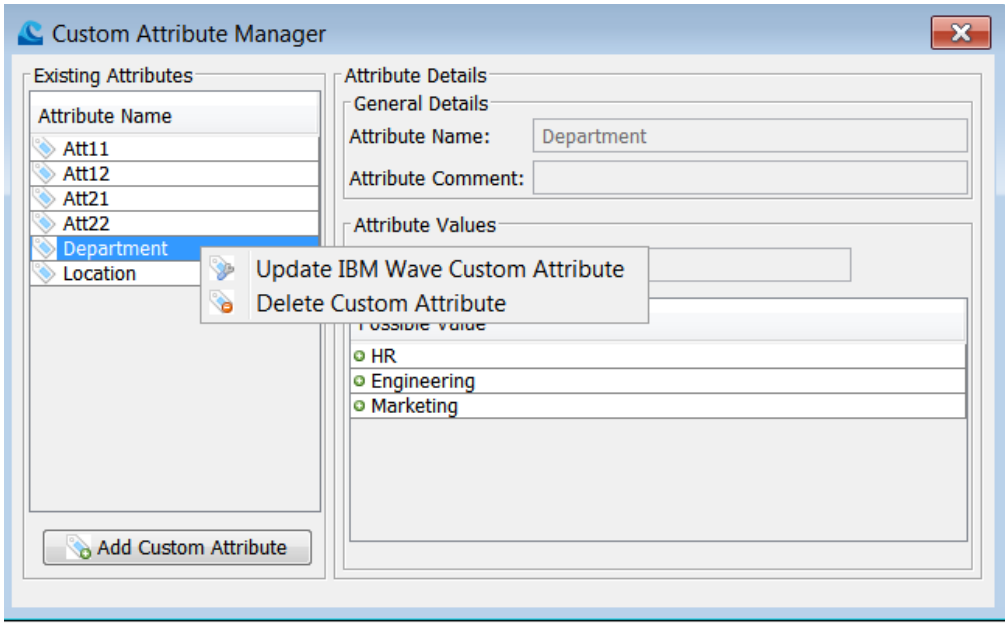

Figure 101. IBM Wave: Custom Attribute Manager

Procedure

Use the following steps to define custom attributes in IBM Wave:

1. To open the **Custom Attribute Manager**, select **Administrative > Site Management > Custom Attribute Manager**.
2. To create a new attribute, click "Add Custom Attribute".
3. Enter the “Attribute Name”, the “Default Value”, and any “Attribute Comments”. If you do not enter a default value, IBM Wave use the default value, which is “Not Assigned”.
4. To assign values to your attributes, click "Add Possible Values". An input field opens to accept the new values. The information that you supplied in step 2 and step 3 now populates the “Existing Attributes” pane and the “Attribute Details” pane in the **Custom Attribute Manager**.
5. When you are done adding attribute values, click “**Go**.” The work unit is submitted to BTS, and the new or modified attributes appear in the “Existing Attributes” pane.
   a. To delete an attribute, right-click on the attribute in the “Existing Attributes” table and click “Delete Custom Attribute”. After you delete the attribute, it is permanently deleted from the IBM Wave.
   b. To update an attribute, right-click on the attribute in the “Existing Attributes” table and click “Update IBM Wave Custom Attribute”. After you change the attribute, the updated attribute appears in IBM Wave **Custom Attribute Manager**.

Results

You are ready to use the **Custom Attribute Manager** to classify guests for the z/VM Guests and Groups Viewer.

What to do next

For more information about generating a report of the custom attribute values that are defined in your environment, see the following topics:

- "IBM Wave Custom Attribute Report" on page 219.
- Additionally, for an overview of the Custom Attribute Manager that includes an example for working with custom attributes, see the topic about "Custom Attribute Manager" in .

Custom Attribute Manager Behaviors:

- When a custom attribute is removed from the **Custom Attribute Manager**, the custom attribute is no longer valid and can no longer be included in z/VM reports. If the z/VM report is open at the time a custom attribute is deleted, IBM Wave issues a warning message that indicates a custom attribute was removed. If you save the report, the removed custom attribute is not saved. When a saved z/VM report has one or more custom attributes removed, IBM Wave issues a warning message that indicates the names of the attributes that were removed. When you open and save the report again, the removed custom attribute is removed from the report. For more information, see "Custom attribute report behavior" on page 220.
- If you create two custom attributes with the same name but use a different style case, the custom attribute is not created, and IBM Wave does not issue an error message. For example, you have an existing custom attribute "Marketing Department", and you create a new custom attribute "marketing department" (or any combination of the same "Marketing Department" upper and lowercase letters). IBM Wave does not create "marketing department" nor does it issue an error message.
- If you create two custom attributes with the same name and add spaces, the custom attribute is created and IBM Wave does not issue an error message. For example, you have an existing custom attribute "Marketing Department", and you create a new custom attribute "Marketing (with two or more spaces here) Department". IBM Wave creates "Marketing (with the number of spaces you added here) Department".

Assign Distribution

To assign an operating system distribution, in the **Current System View > z/VM Users** tab, select all of the guests, right-click, and select **Update > Assign Distribution**.

The **Assign Distribution** ***Operating System*** **to z/VM Users** window opens, as shown in Figure 102 on page 117.

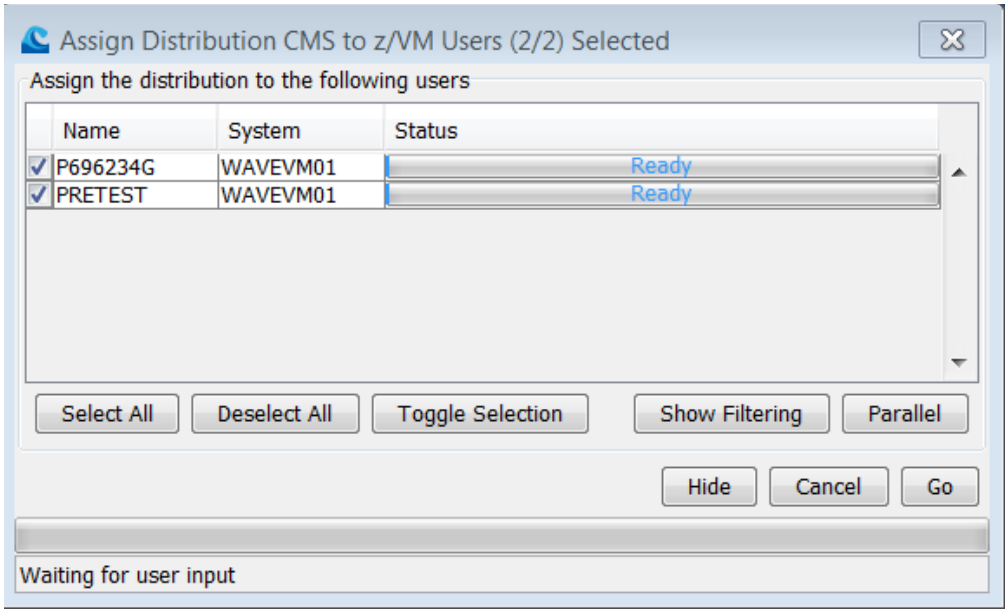

Figure 102. Assign Distribution

When you click **Go**, IBM Wave attempts to assign the operating system (OS) distribution to all of the selected (checked) z/VM Guests.

Note: The assignment is purely metadata information. The actual OS running on the virtual server is verified when either of the following actions occur:

- An OS-dependant action is done against the virtual server.
- A connection is made to the virtual server.

Assign Guests to Project

Use the **Assign Project** option to assign a z/VM guest to a project, or move a group of Linux guests into a project.

To assign one or more z/VM Guests to a project, right-click on the z/VM guest or group of guests, and then select **Update** > **Assign Project**.

Important: When Automatic Guest Classification (AGC) is active and you assign guests to projects, all of the projects must use the "Bidirectional" rule. For more information, see "Run Automatic Guest Classification" on page 131.

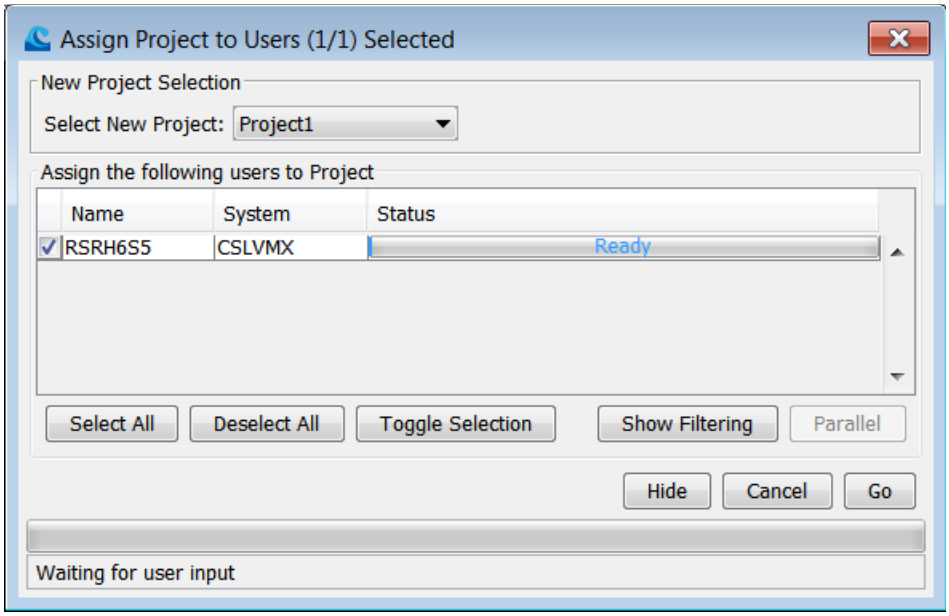

Figure 103. Assign Project to Users

Use the **Select New Project** field to select the project you want the guest or guests that are assigned to, and then click **Go**. IBM Wave attempts to change the "Project" metadata for the selected guest or guests to the assigned "Project".

Note: The project information is metadata that is saved to the IBM Wave database. No information about the project is saved to the z/VM System.

Deactivate z/VM Users

To **Deactivate** a z/VM Guest, right-click on one or more z/VM Guests, and select **Deactivate**.

Deactivation is a process in which the operating system (OS) that is running on the z/VM Guest is shut down in an orderly manner, and then the z/VM Guest is logged off.

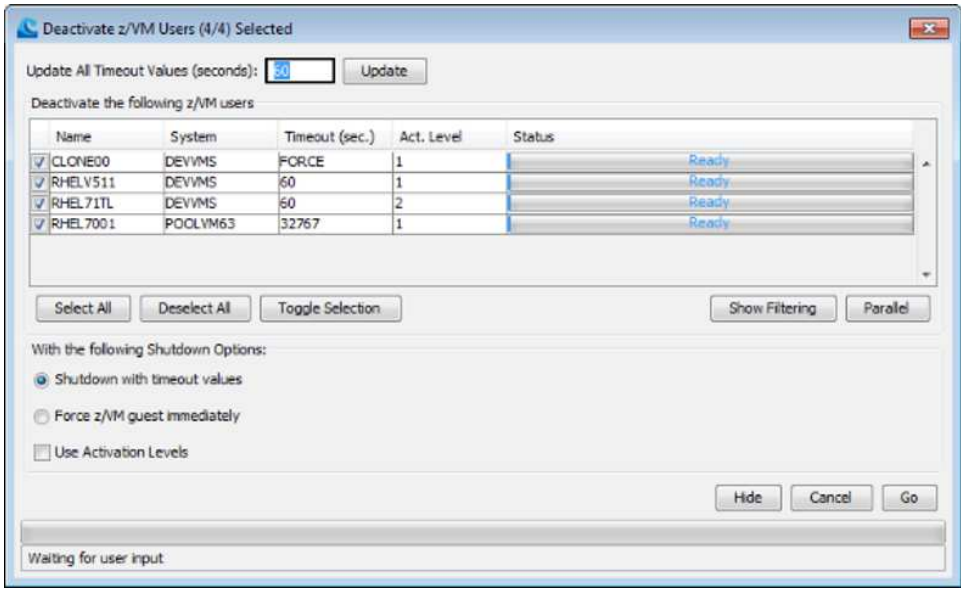

Figure 104. Deactivate z/VM Users

IBM Wave generates a BTS Work unit for the entire deactivation process. When multiple z/VM Guests are selected for deactivation, multiple BTS deactivate requests are included in the BTS Work unit. You can track the BTS Work unit in the **BTS Work unit Viewer**.

IBM Wave supports two methods of deactivating a z/VM Guest:

1. Using the **Shutdown with timeout values** option issues the **`SIGNAL SHUTDOWN`** command to the z/VM for Linux Guests and the **`FORCE IMMEDIATE`** command for guests that are not Linux. Linux guests are shut down in an orderly manner. After the OS completes the shutdown process, the z/VM Guest are logged off. The timeout value is specified in seconds. The timeout value is automatically set for each guest with the default value for the signal-shutdown timeout that is specified in its z/VM system. If no default is set in z/VM, IBM Wave uses a default value of 32767.

   You can use the following two methods to override the default z/VM timeout value:

   - **Update All Timeout Values (seconds)**, which is at the top of the **Deactivate z/VM Users** menu: Select the timeout value to use, and then click update. All Linux guests in the table are populated with the new timeout value.
   - Edit the **Timeout (sec.)** field in the table: Double-click the timeout value in the table that you want to change, and enter the new value.

   Notes:

   - Overriding the timeout value during the deactivate process is not persistent. The default value is repopulated during any subsequent deactivate.
   - If the OS running on the z/VM Guest does not shut down within the allotted timeout value, the z/VM Guest is **`FORCED`**.
   - If none of the guests that are selected for deactivation are Linux guests, the **Shutdown with timeout values** option is disabled.
2. Using the **Force z/VM guest immediately** option issues the **`FORCE IMMEDIATE`** z/VM command to all the selected guests to log off the z/VM Guest.

Note: Only use the **Force z/VM guest immediately** option for z/VM Guests that run workloads in which an interruption has no affect. Using this method to deactivate z/VM Guests that run Linux is highly discouraged.

When the "Use Activation Levels" check box is selected, IBM Wave deactivates the z/VM Guests by using the selected method in descending order of the Activation Levels. The generated BTS work unit also contains a "wait deactivation" BTS request for each z/VM Guest. IBM Wave cannot start the deactivation of a z/VM Guest until the previous z/VM Guest (in terms of Activation Level) finishes the deactivation process.

As shown in Figure 104 on page 119, the following actions occur:

1. The z/VM Guest RHEL71TL is deactivated first.
2. IBM Wave waits until z/VM Guest RHEL71TL is no longer logged on.
3. The z/VM Guests CLONE00, RHELV511, RHEL7001 are deactivated together in parallel.
4. IBM Wave waits until the z/VM Guests CLONE00, RHELV511, RHEL7001 are no longer logged on.

When multiple z/VM Guests are assigned the same activation level, they are deactivated as a group, in parallel. IBM Wave waits for all of the z/VM Guests in the group to be deactivated before it moves on to the next Activation Level. For example, z/VM Guest RHEL71TL has an activation level of 2, so it is deactivated first. After RHEL71TL is logged off, all guests with activation level 1 are deactivated and then logged off.

Delete z/VM Guests

To delete one or more z/VM Guests from a z/VM System, select the guests, and then right-click and select **More Actions** > **Delete**.

Important: The **Delete** action is irreversible. It causes all data on the z/VM Guests and minidisks to become unavailable.
Additionally, the delete action does not remove the data from the minidisk, but instead marks the space that was occupied on the minidisk as free. If the z/VM Guest was connected to a permanent Virtual Network, the GRANT commands are deleted from the AUTOLOG parameter files.

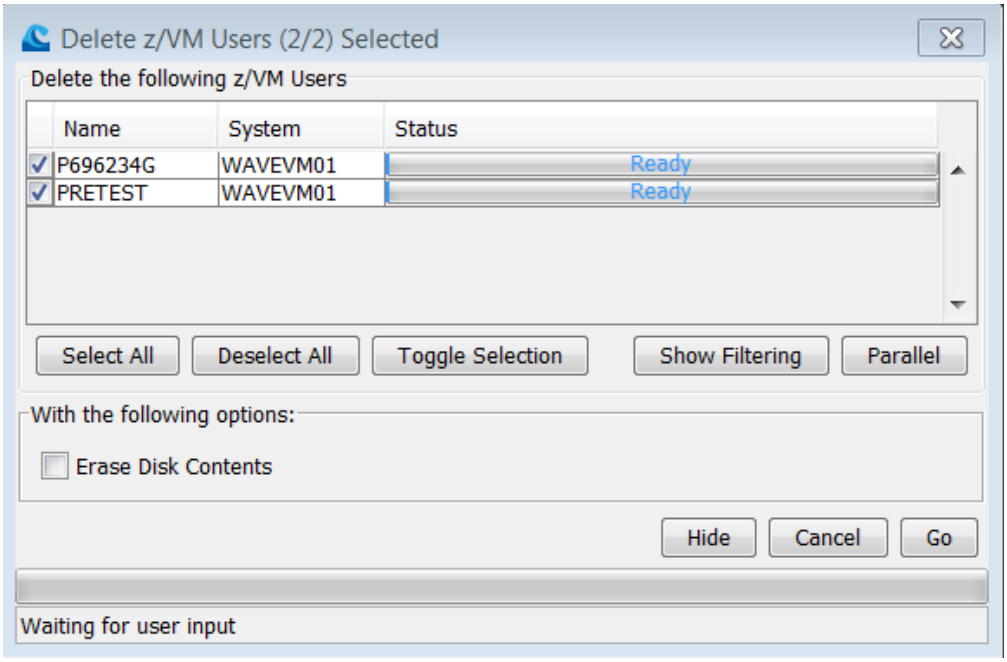

Figure 105. Delete z/VM Guests

You can optionally select the **Erase Disk Contents** check box to erase the minidisk that is associated with the z/VM Users. By default, the disks are not erased.

Press **Go** to delete all of the (checked) z/VM Users.

Execute script

Use this action to execute a script on one or more Linux virtual servers. In order to select a script, press the **Browse...** button. The Script Manager window opens. Select a script from the Script Manager window by double clicking the corresponding entry. The script name, description, and creation or modification information is displayed in the relevant fields in the bottom section of the window.

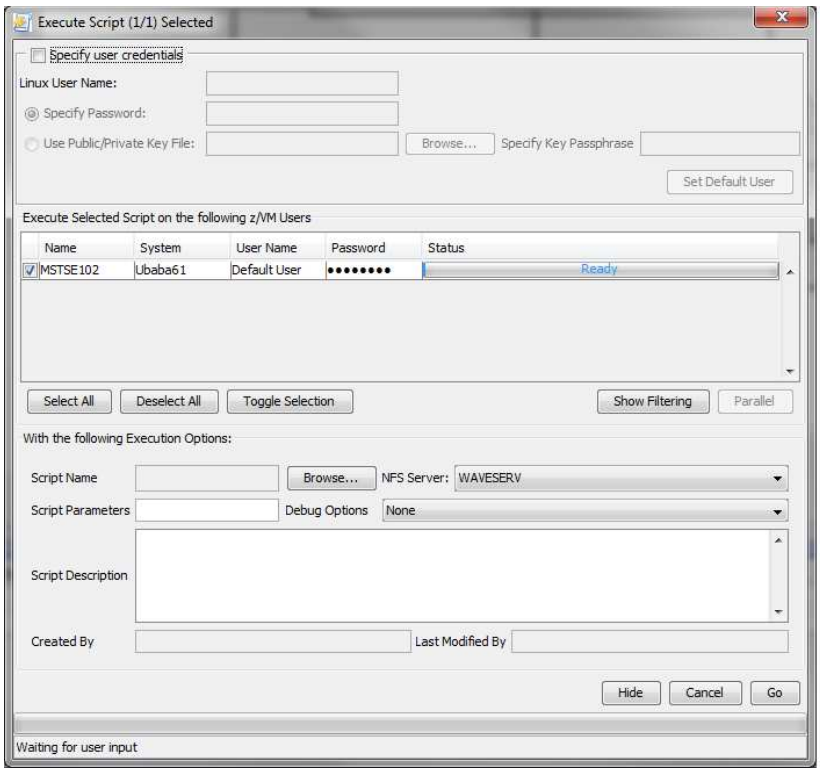

The NFS Server dropdown box will by default contain the lowest common denominator for all guests included in the action. If all guests included in the action belong to a z/VM System which has a default NFS Server defined to it, that server will appear as the default for the action. Otherwise, the IBM WAVESRV server will appear as the default. If additional NFS Servers are defined, it is possible to select a particular NFS Server from the dropdown box. The dropdown box will always contain at least 2 options:

1. **WAVESERV** - Indicates that the IBM WAVESRV will be utilized as the NFS Server for the action.
2. **Default for guest z/VM System** - Indicates that the default NFS Server defined to the z/VM System to which the guest belongs will be used as the NFS Server for the action.

Note: If the selected script is not synced to the selected NFS server, an error message will appear. It is not possible to select a script and NFS Server combination where the script has not been synced to the NFS Server.

It is possible to select input parameters for the script and a debug level.

It is also possible to specify a Linux user under which this script will be executed. By default, IBM Wave will use the normal Authentication Method and user to run the script on the Linux Guests. However, under some circumstances, the IBM Wave User might need to run the script under a different Linux User. Authentication of the "Execution Phase" Linux user can be done either by providing a password, or specifying a location of an openSSH format private key file. There are 2 ways to specify an "execution Linux user":

1. In the top portion of the window, specify the Linux users' credentials. Once the "Set Default User" button is pressed, IBM Wave will change the execution user for all Linux Guests involved in this action.

2. Specify an execution user per-server. This can be done by manually editing the "User Name" and "Password"/"Key File Location" fields of the table.

Once the "Go" button is pressed, a BTS work unit will be generated and submitted to the BTS. The work unit will contain "Execute Script" requests for each guest included in the action. For each Linux Guest, the request is divided into 3 phases:

- **The Preparation Phase** - In this phase, the Linux Guest is accessed using the default Linux User (depending on the selected Authentication Method), some integrity checks are performed and the script directory on the IBM WAVESRV virtual server will be mounted on a temporary directory.
- **The Execution Phase** - The selected script will be executed. The selected script is executed by a script executor, which provides access to various variable, debug levels etc. By default, this phase is run under the default Linux user (The same user used for the Preparation Phase), unless a different, execution user was specified.

  Note: If the private/public key authentication is used, and the private key is encrypted with a passphrase, IBM Wave prompts the IBM Wave User for the passphrase.

  Depending on the value of the NFS Server selected for the action, the script directory is mounted and the script is executed from there.
- **The Cleanup Phase** - After the script has executed, the script directory is unmounted. This phase is also executed under the default User, regardless of the execution user specified.

The BTS Work unit can be tracked from the BTS Work unit Viewer. Provided that the script was executed successfully, the results can be viewed by drilling down into the BTS Work unit, selecting the request and switching over to the “Script” COR tab. All stdout and stderr output of the script will be directed into the script COR entry for the request.

Note: Although an execution user can be defined, the preparation phase and the cleanup phase will always run under the default Linux user (Depending on the selected Authentication Method). If a different execution user is specified, the execution phase will run under that Linux user.

Generate Disk Storage Map

You can generate a visual representation of the disk storage map associated with a selected z/VM guest. From the guest, select **z/VM User Actions > More Actions > Generate Disk Storage Map**. To include the IBM Service Machines in the generated map, you can optionally select “Include IBM Service Machines”.

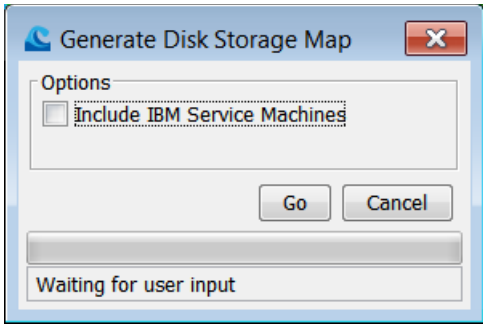

Figure 106. Generate Disk Storage Map

Click **Go** to open the **Disk Storage Map View**. The **Disk Storage Map View** is layer-based (similar to the **Network Viewer**). The following layers are available:

- DASD Volumes
- EDEV Volumes
- FCP Volumes
- Storage Controllers
- Device Pools
- Guests

As shown in Figure 107, the **Disk Storage Map View** contains several "Predefined Views" for the predefined layer combinations. To customize the view, you can toggle an individual layer on or off in the "By Layers" pane.

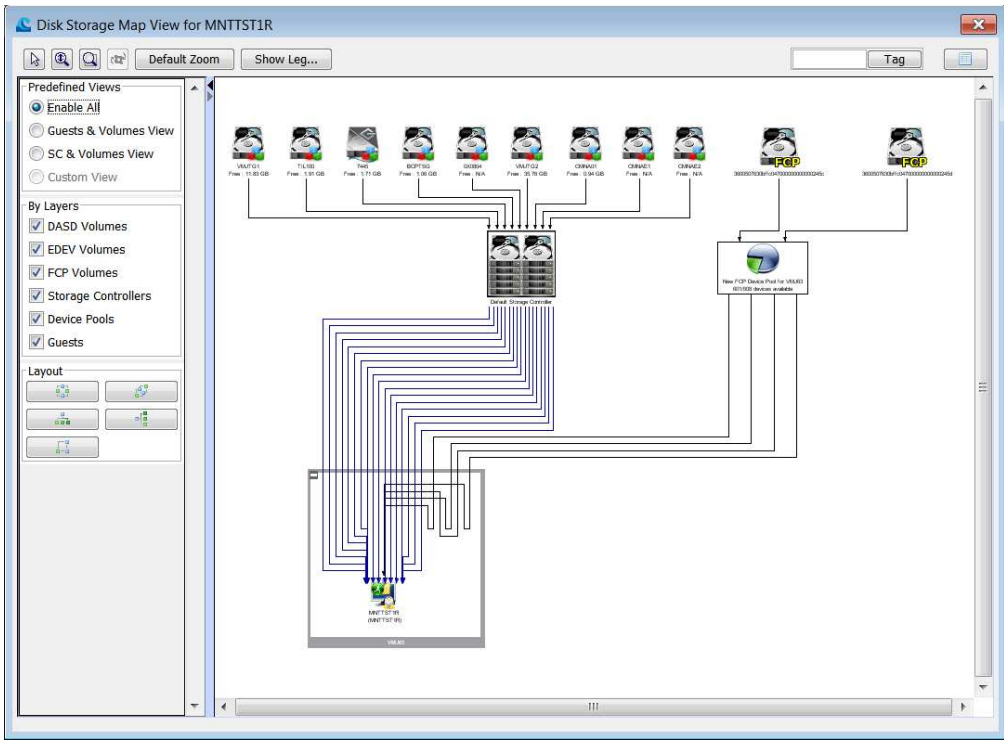

Figure 107. Disk Storage Map View

The generated **Disk Storage Map View** illustrates the type of connection between the z/VM Guests and the Volumes. IBM Wave displays dedicated DASD, and DASD that is associated to the z/VM Guest by using the profile that is included in

the z/VM Guest directory. IBM Wave also shows linked DASD Volumes without limitation to the link-level. (For example, a z/VM Guest that is linked to a device in another z/VM Guest that is itself linked to another z/VM Guest.)

The same applies to FCP storage. IBM Wave displays all of the FCP storage controllers and Target WWPN and LUN IDs that are attached to the selected z/VM Guest. For FCP attached storage, FCP devices that are attached to z/VM Guests appear as connector lines (similar to the manner that NICs are displayed in the **Network Viewer**).

The **Disk Storage Map View** contains an alternative table view, in the upper right portion of the screen, for FCP storage. The **FCP Connections** table contains the FCP Device and the Target WWPN and LUN IDs for every FCP LUN attached to the guests. You can also view the other guest fields in the **FCP Connections** table.

Init users for IBM Wave use

See "Init Users for IBM Wave use" on page 55.

Lock or unlock z/VM Users

To lock or unlock z/VM Users, right-click of the guests and click **More Actions** > **Lock or Unlock**. z/VM Users that are already in the target state are skipped.

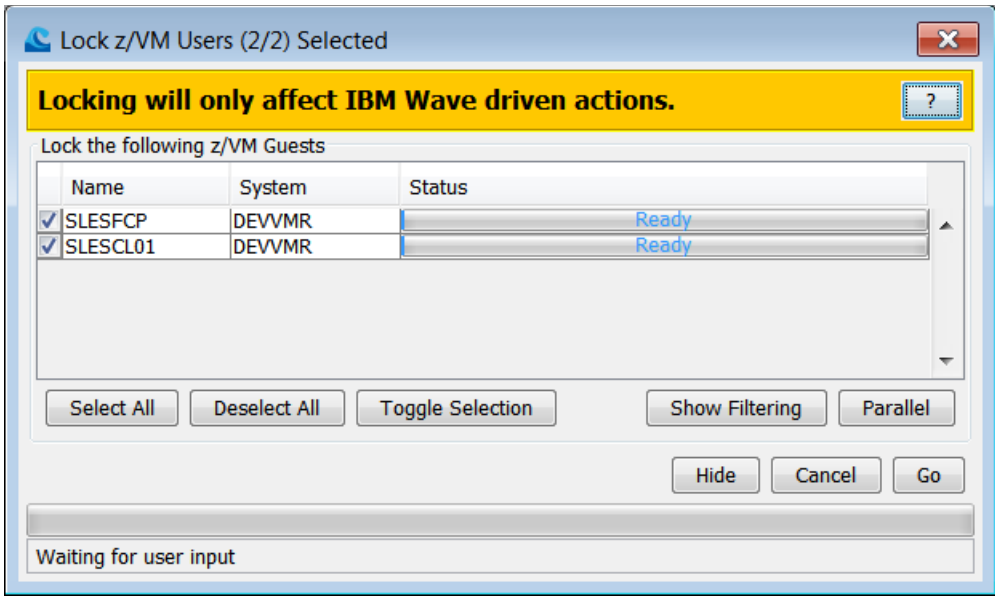

Figure 108. Lock z/VM Users

Mark as IBM Wave service machine

Use this action to mark a group of z/VM Guests as internal IBM Wave service machines.

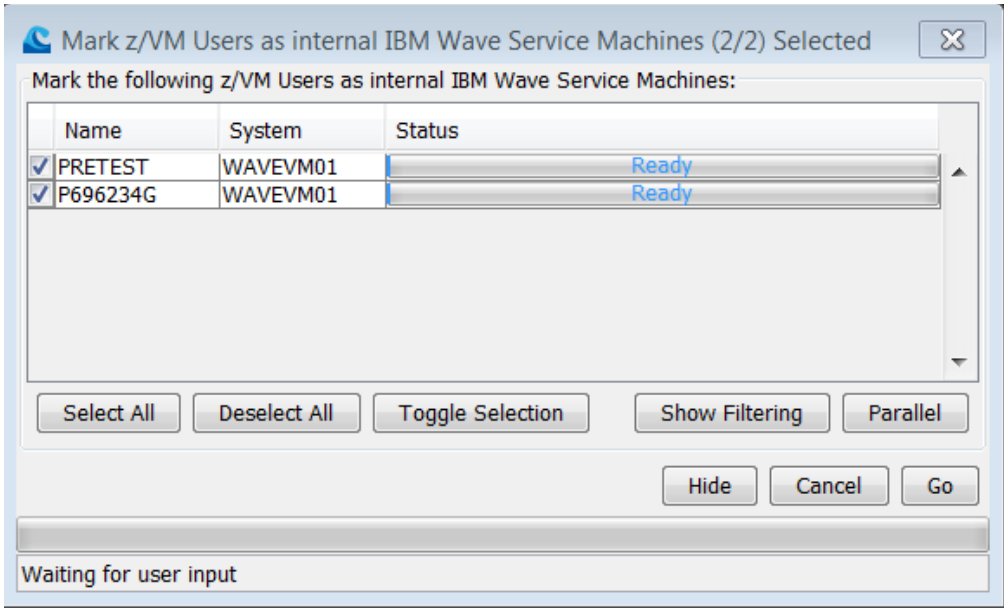

Important: This action must only be used when the z/VM Guest Directory is shared between several z/VM Systems that are managed by IBM Wave and IBM Wave has failed to recognize the other z/VM Systems' service machines as such.

Mark as initialized for IBM Wave

Use this action to indicate to IBM Wave that the selected z/VM Guest has already been initialized for IBM Wave.

Note: This action should only be used when one or more of the following conditions are met:

1. A certain z/VM Guest resides on shared storage between more than one z/VM System, and was already initialized for IBM Wave on the other z/VM System.
2. A certain z/VM System has been added to IBM Wave management, the z/VM Guest has been initialized for IBM Wave, and then the z/VM System was removed from IBM Wave management and re-added at a later time.

Pause

To suspend the CPU of one or more z/VM Guests, right-click on the guest or guests, and select **Pause**. The IBM Wave service machines issue the **CP STOP** command to put the z/VM guests in the **CP READ** state.

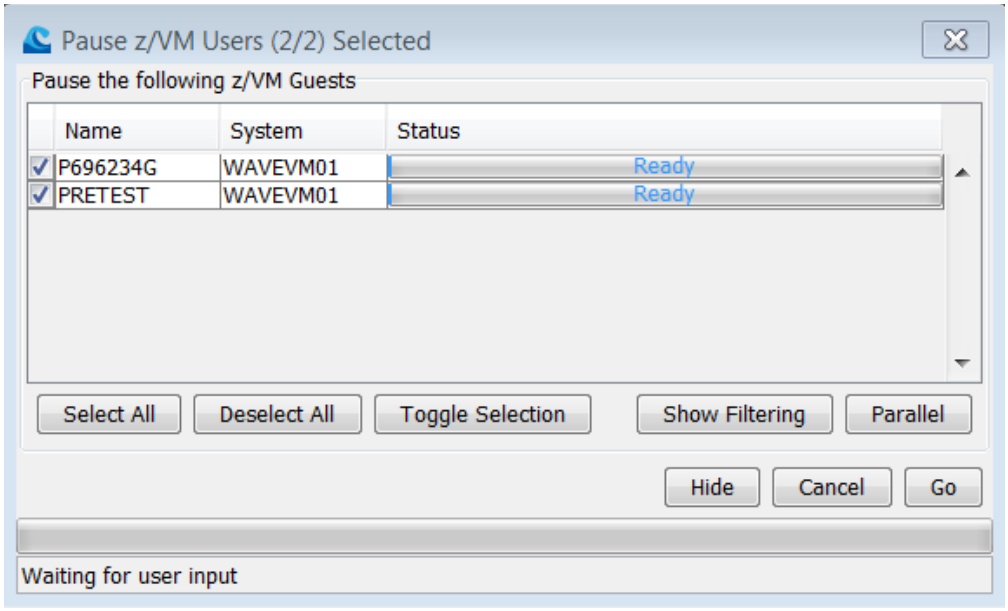

Figure 109. Pause z/VM Users

To **Resume**, right-click on the guest or guests, and select **Resume**.

Recycle z/VM users

Use the **Recycle** action to deactivate and then activate one or more z/VM users. When you select "Shutdown with timeout values" during the deactivation, the process is done in the same manner as a typical deactivate (for more information, see "Deactivate z/VM Users" on page 118). During the recycle, after a successful deactivation, the z/VM user is automatically reactivated.

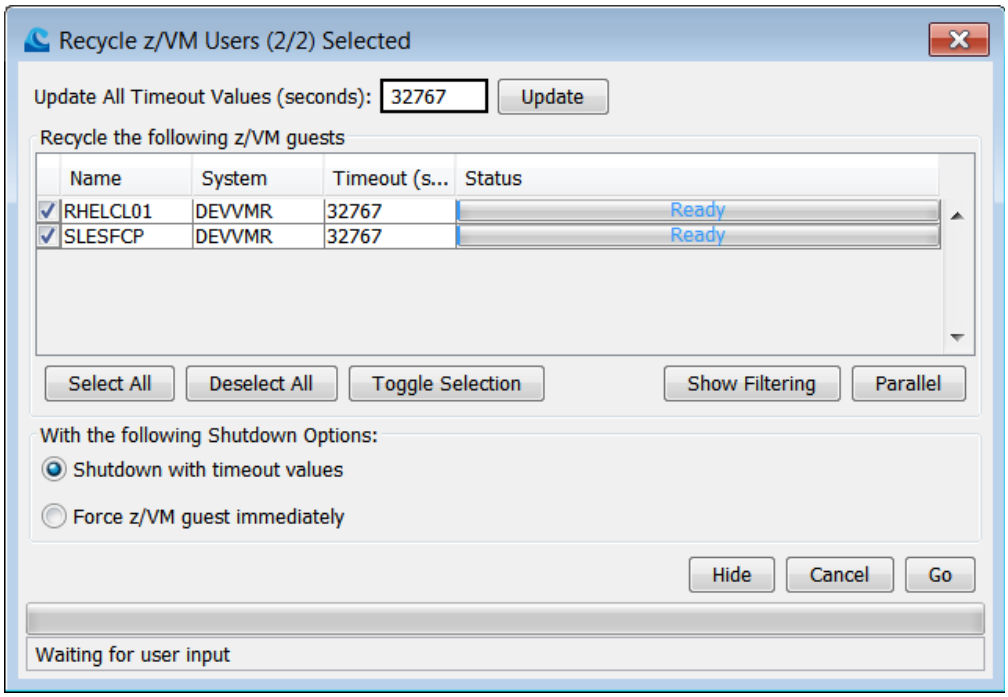

Figure 110. Recycle z/VM User

When you click **Go**, IBM Wave generates a BTS work unit for the entire recycle process. When multiple z/VM Guests are selected to recycle, multiple BTS recycle requests are included in the BTS work units. You can track the BTS work units by using the **BTS Work Units Viewer**. For more information, see "BTS Work Units Viewer" on page 19.

Note: The **Recycle** process disconnects (logs off) the z/VM Guest. Any changes to the z/VM definitions become active after the recycle is complete.

Refresh Linux Data

To refresh data for one or more Linux Guests, right-click on the Linux Guests, and then select **Refresh Linux Data**. The refresh runs a subset of the "Init For Wave" processes, including updating the hostname and querying the FCP devices.

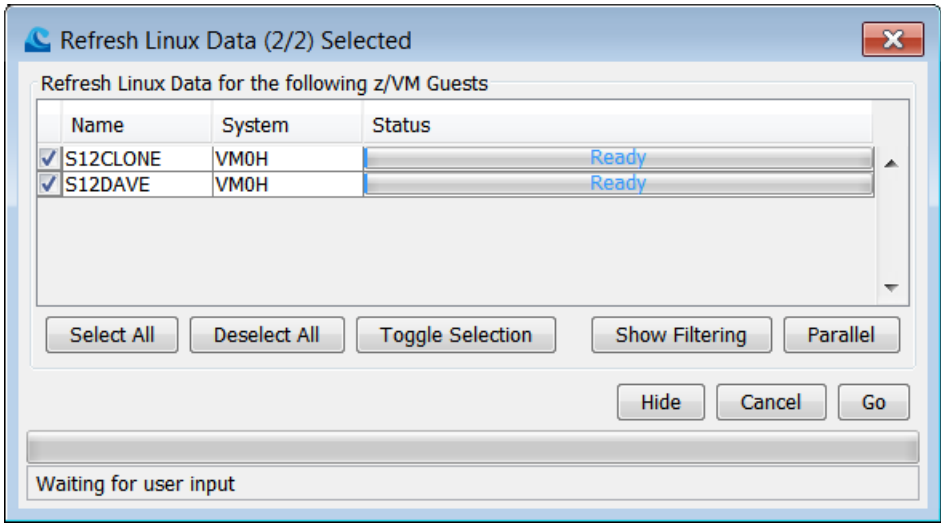

Figure 111. Refresh Linux Data

Relocate to...

Use this action for a Live Guest Relocation (LGR) against one or more z/VM Guests. Use the bottom portion of the window to indicate special options for the relocation process, such as forcing, time limitations, and others.

The LGR relocate function is available from the **Enterprise Viewer** and **Dashboard Viewer**.

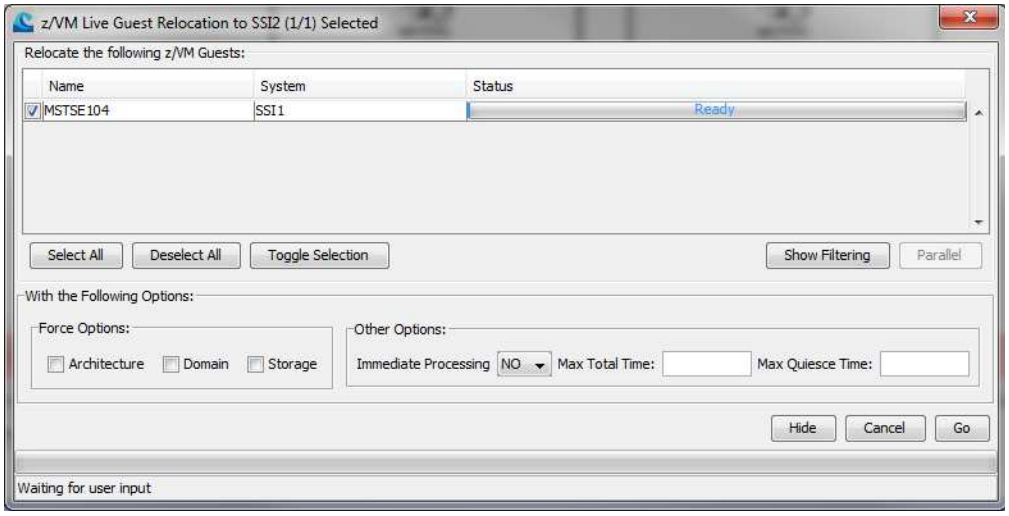

Figure 112. Live Guest Relocation (LGR)

After you select the systems to relocate, press **Go**. A BTS work unit is generated with an LGR request for each guest that is selected for relocation. The BTS work unit can be tracked from the **BTS Work Unit Viewer**.

Usage Notes:

- IBM Wave triggers the LGR process that uses the Systems Management API (SMAPI), which, in turn, uses the **VMRELOCATE** CP command. LGR processes triggered from within IBM Wave are subject to the same limitations and provisions that are imposed by the **VMRELOCATE** CP command. If an IBM Wave LGR request fails, it indicates the CP message codes that are associated with the failure in the COR output of the request.
- The **Create New z/VM Guest** action does not create guests that are eligible for LGR. The new guest has LINKs to MAINT MDISKs that reside on the SYSRES volume, which is exclusive for each member in the SSI cluster and not shared. Also, the guest does not contain the CHPIDVIRTUALIZATION Option statement.

Remove account

Use this action to remove a z/VM Account from one or more z/VM Guests.

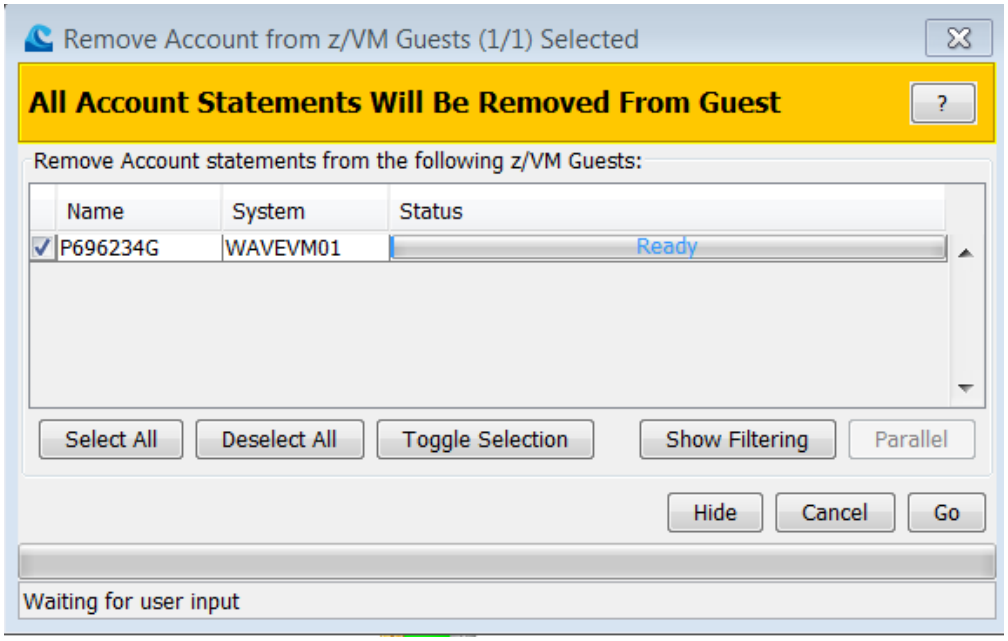

Pressing "**Go**" will submit a Work unit and request to the BTS which will perform the action.

Note: This action will remove any "ACCOUNT" statements from the directory entry of the guest. This means that apart from the primary account number, this action will also remove any secondary accounts and distribution codes associated with the z/VM Guests.

Reset installation status

Use this action to reset the Installation Status indication for the z/VM Guest. This action is only available when the z/VM Guest is currently running an interactive Linux Installation. The action will terminate any active Linux Installation BTS Requests associated with the z/VM Guest.

Resume

Use the **Resume** action to resume z/VM Guests paused using the "Pause" action. This action invokes the IBM Wave Service machines to issue the **`CP BEGIN`** command for all the selected guests.

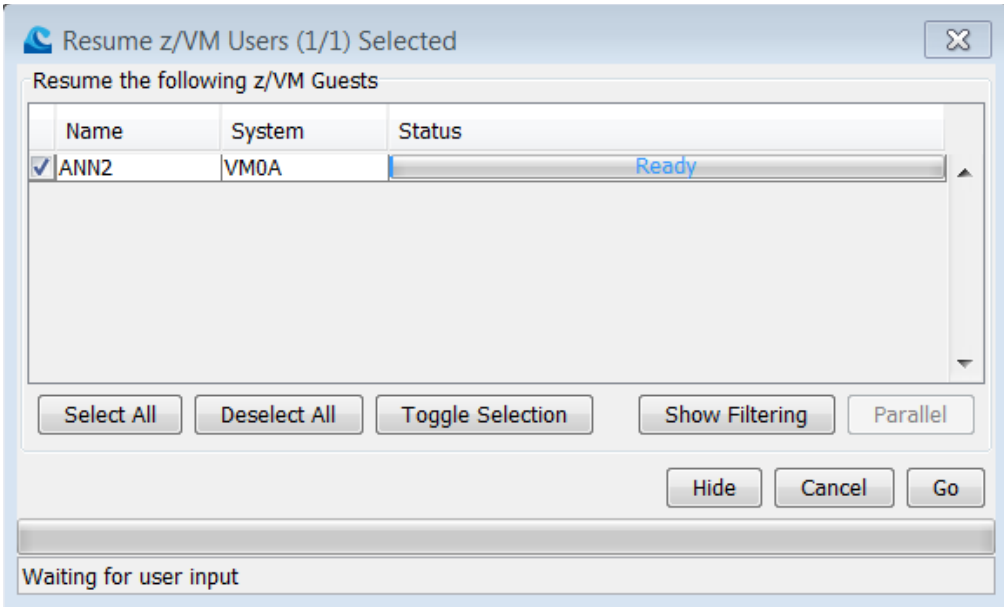

Run Automatic Guest Classification

You can run Automatic Guest Classification (AGC) against multiple z/VM guests.

To run AGC, right-click on one or more z/VM guests and then click **More Actions** > **Run AGC**.

Important: Before you **Run AGC** for the first time, be sure to review all of the "AGC" topics in .

Send message

Use this action to send a message to a z/VM Guest or users. This action connects to the Linux Virtual Server using the "root" user and issue a message to all connected Linux users.

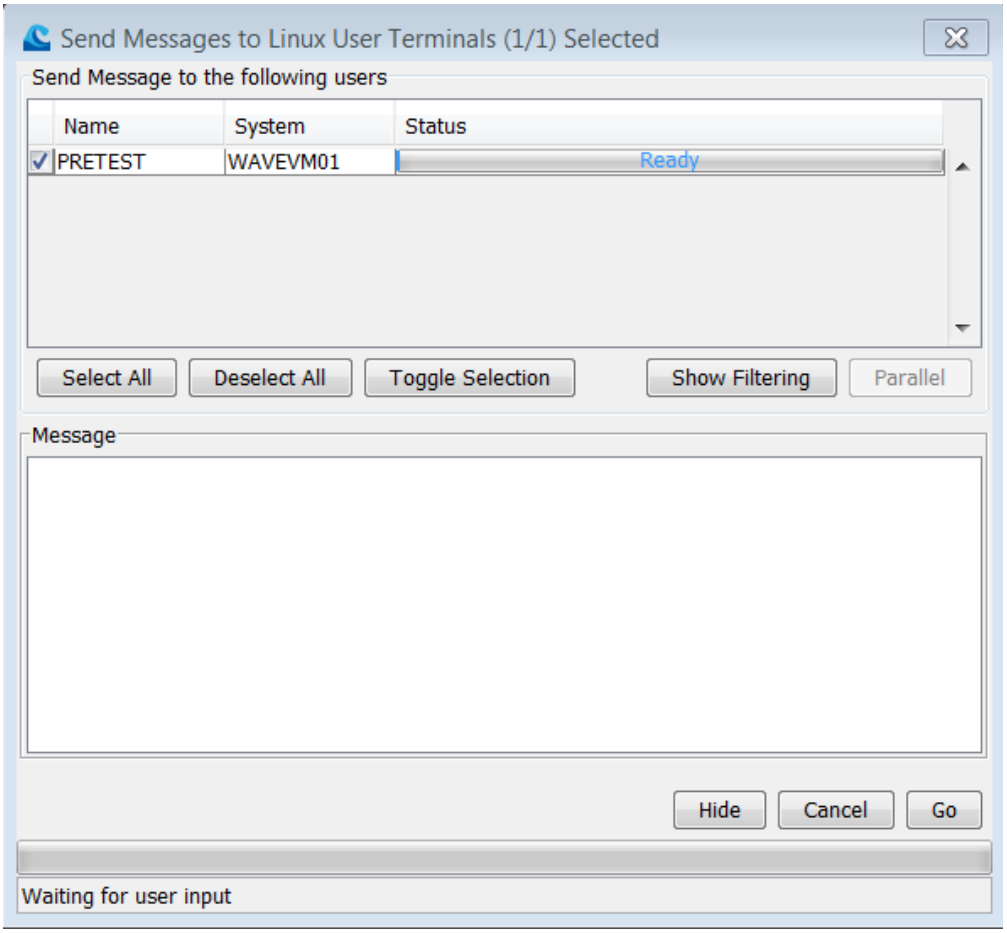

The message text box is used to specify the message. Each message sent will be prefixed and suffixed with an indication that it was sent by IBM Wave from the IBM Wave User.

Set default z/VM system

Use this action to set the default z/VM System for the selected z/VM Guests. For more information about the Default z/VM System mechanism, see Default z/VM system assignment for z/VM guests in .

Set functionality

Use this action to set the Functionality Type for the selected z/VM Guests. The Functionality Type can be later used to control the order of activation for these z/VM Guests. For more information about Functionality Types and Activation Levels, see .

Deselect as IBM Wave service machines

Use this action to deselect z/VM Guests from being marked as IBM Wave internal Service Machines. This action can be run against z/VM Guests that were marked as service machines by using the "Mark as IBM Wave Service Machine" action.

Virtual networks functions

All of the actions in this section are relevant to virtual networks and are available from the menus in the **Current System Viewer** > **Network Viewer**.

Note: All **Add**, **Delete** and **Update** actions apply to IBM Wave users who have the Network Administrator role.

Add New Guest LAN

To create a new Guest LAN for a z/VM System, right-click in the **Current System** > **Network** tab, and then select **Create New Virtual Guest** > **Add Guest LAN**.

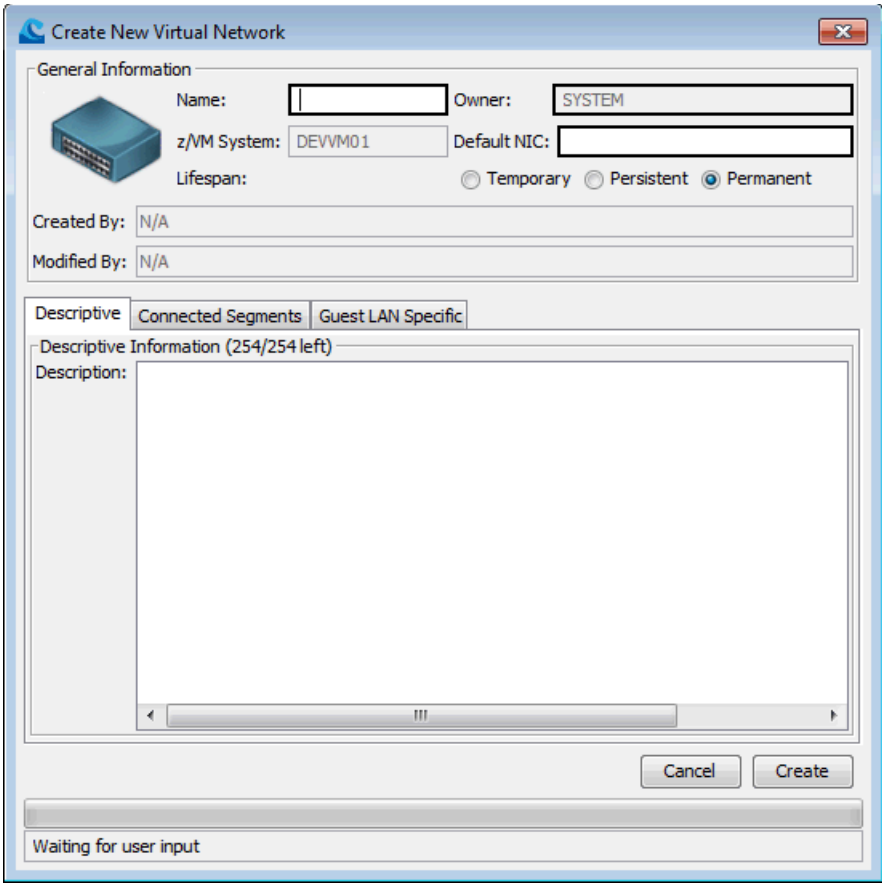

Figure 113. Create New Virtual Guest menu

Fields in this window:

- **Name** - The name of the Guest LAN, limited to 8 characters.
- **Owner** - The name of the z/VM Guest who owns the Guest LAN. For a persistent and a permanent guest LAN, this field is disabled and set to "SYSTEM".
- **Default NIC** - This virtual device address is used when connecting z/VM Guests to the Virtual Network using IBM Wave.

  Note: The address for the default NIC must be different from any other virtual device address in the user directory. Passing the same address can cause a NIC definition conflict, which can disable some of the guest features in IBM Wave.
- **Lifespan**:
  - **Temporary** - The Guest LAN is available until the z/VM Guest who owns it is deactivated.
  - **Persistent** - Guest LAN is available until the z/VM System is IPLed.
  - **Permanent** - The Guest LAN is available even after the IPL. IBM Wave adds a command to the AUTOLOG PROFILE EXEC that creates the Guest LAN.

- **Description** - An optional description.
- **Connected Segments** - This tab displays a table that contains a list of the Virtual Network Segments that are connected to the Guest LAN.

Guest LAN Specific: This tab contains options that are relevant for the Guest LAN. It permits a router machine to be defined for the Guest LAN. When the Guest LAN is created, IBM Wave performs the required z/VM actions to create the Guest LAN.

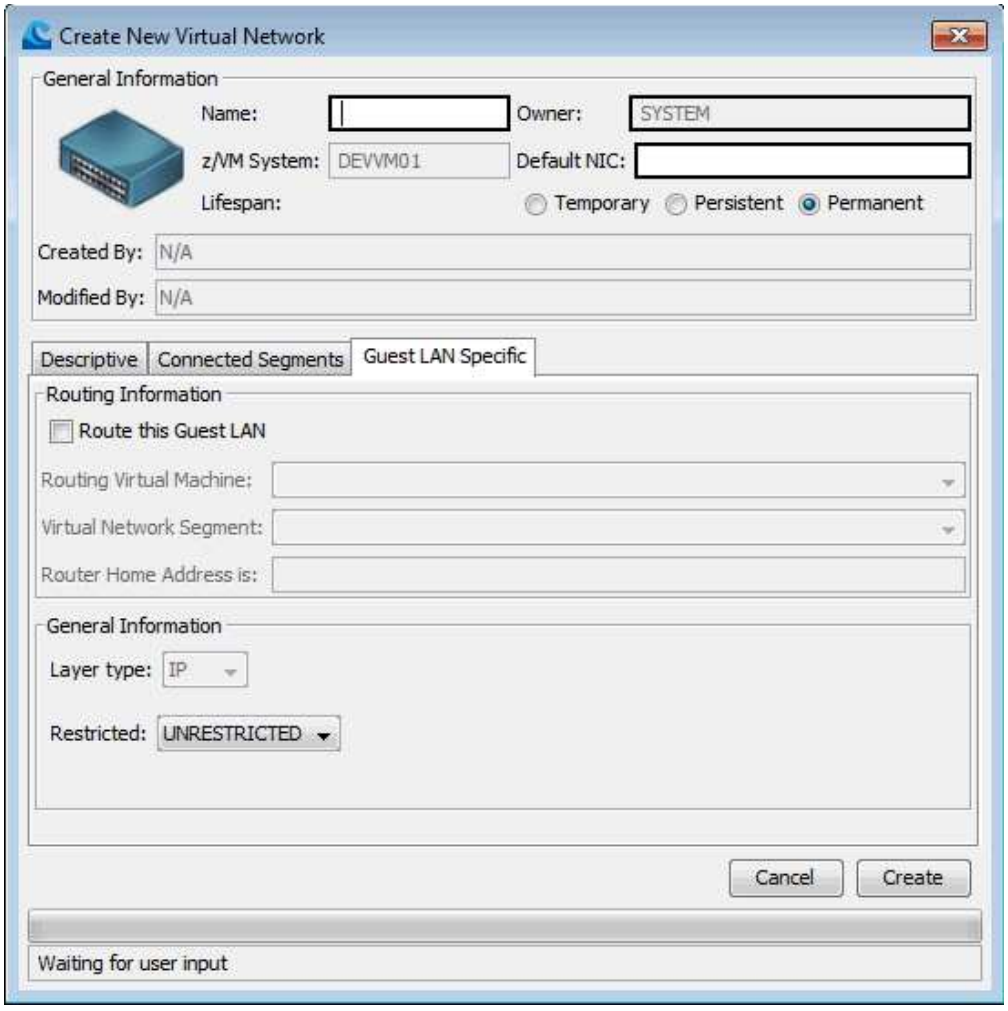

Figure 114. Create New Virtual Network pane: Guest LAN Specific tab

Fields in this tab:

- **Routing Information** - This section of the tab contains the routing information.
  Route this Guest LAN - If the **Route this Guest LAN** check box is selected, IBM Wave adds definitions to the router machine. If the Guest LAN is defined as permanent, IBM Wave updates the PROFILE file of the router machine.

Note: IBM Wave automatically creates an "Unknown IP" Virtual Network Segment (VNS) for every Guest LAN that is created. For more information about the unknown IP VNS, see the topic about "Virtual Network Segments" in .

Add New VSwitch

To define a new VSwitch in the z/VM System, go to the Network tab in the Hardware Viewer. Right-click in the white space, and select **Add New VSwitch** to display the VSwitch specific-tab in the Create New Virtual Network menu.

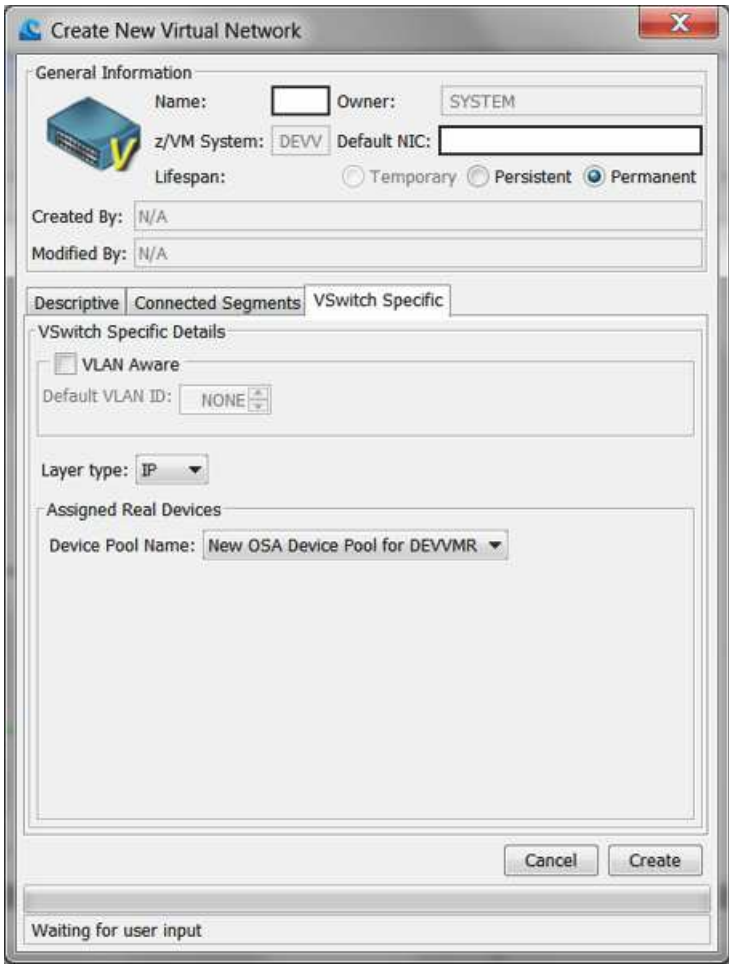

Figure 115. Add a new VSwitch

Fields in this view:

- **Name** - A name for the new VSwitch.
- **Owner** - The owner for VSwitches always defaults to "SYSTEM" because a VSwitch can be either persistent or permanent.
- **Default NIC** - This virtual device address is used to connect z/VM Guests to the Virtual Network that uses IBM Wave.

  Note: The address for the default NIC must be different from any other virtual device address in the user directory. Passing the same address can cause a NIC definition conflict, which can disable some of the guest features in IBM Wave.
- **Lifespan** - The life span of the VSwitch. Define VSwitches as either persistent or permanent. For an explanation of life spans, see "Add New Guest LAN" on page 133.
- **Description** - An optional description for the VSwitch.
- **Connected Segments** - This tab displays a table that contains all Virtual Network Segments that are connected to this VSwitch.

- **VSwitch Specific** - This tab displays information that is specific to the VSwitch creation. The IBM Wave User can define a triplet of OSA devices to which to connect the VSwitch, and the option to make the VSwitch VLAN-Aware.

Display virtual network information

Use this action to display details on a selected Virtual Network. For a description of the fields see "Add New Guest LAN" on page 133 or "Add New VSwitch" on page 135.

Update virtual network information

Use this action to update certain aspects of the selected Virtual Network. For a description of the fields see "Add New Guest LAN" on page 133 or "Add New VSwitch" on page 135.

Delete a virtual network

To delete a Virtual Network, right-click on the Virtual Network, and then select **More Actions** > **Delete**. If the Guest LAN has been defined as "Permanent", IBM Wave deletes the creation entry from the **LANPROF AUTLOG** parameter file.

Note: You cannot delete a Virtual Network that has z/VM Guests connected to it.

Connecting z/VM guests to virtual network segments

In the Network Viewer, use the "Magic Wand" tool to connect a z/VM Guest to a Virtual Network Segment. For more information about the Magic Wand, see "Special tools available in this viewer" on page 32.

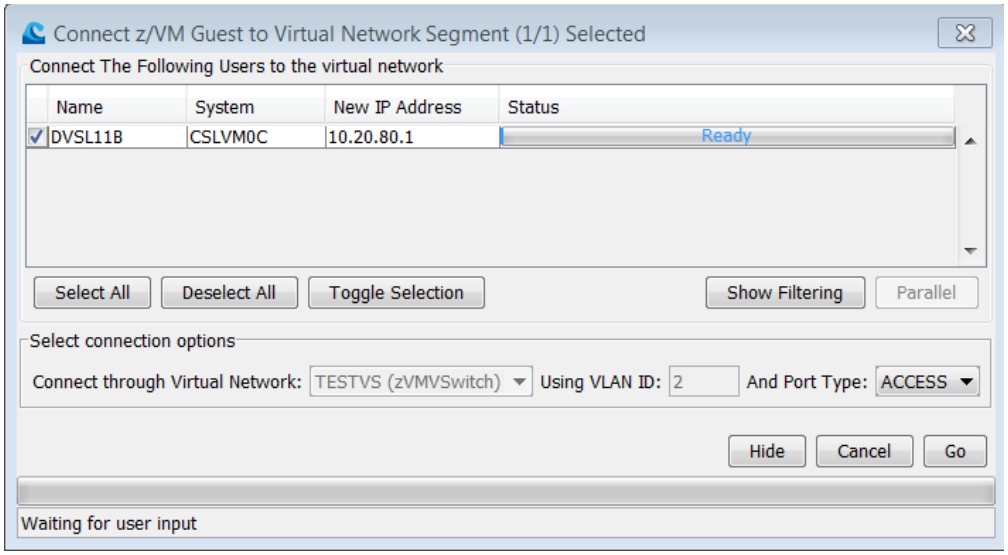

Figure 116. Connect z/VM Guest to Virtual Network Segment

IBM Wave automatically generates a new IP address that corresponds to the IP segment defined for the VNS selected. If a different IP is required, put the new IP in the **New IP Address** field. The manually provided IP address must match the network and netmask of the selected VNS.

Important: IBM Wave identifies the IP addresses of guests that it manages. However, if you set up a remote Wave server, it has no way of identifying the correct IP addresses when you clone a user. Therefore, when the Wave server is

remote and you clone a user, the automatically suggested IP address might belong to the remote Wave Server. If so, it needs to be manually changed. (A remote IBM Wave server is one that is not installed on the same z/VM system where the cloned guest is being deployed.)

If the selected VNS is connected to more than one Virtual Network, you can select the Virtual Network from the **Select connection options** pane.

IBM Wave automatically detects if the new connection requires a new Virtual NIC to be defined to the z/VM Guest. (For example, when the z/VM Guest connects to a VSwitch to which it is not already connected.)

IBM Wave automatically attempts to determine whether the virtual NIC (new or existing) is defined as "access" or "trunk". The decision is computed based on the following data:

1. The type of the selected Virtual Network (Guest LAN, VLAN Aware VSwitch, VLAN Unaware VSwitch).
2. Existing connections between the z/VM Guest and the selected Virtual Network.
3. Whether the selected VNS has a VLAN ID defined.

The following example is a network configuration:

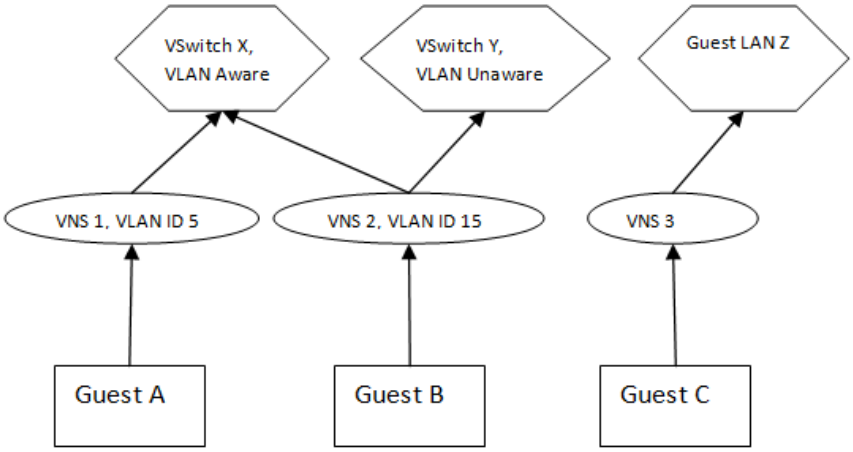

Figure 117. Example network configuration

If the user decides to connect Guest A to VNS 2, IBM Wave presents the user with two possible values in the **Connect Through Virtual Network**box: VSwitch "X" and VSwitch "Y".

- If the user selects VSwitch "X", IBM Wave automatically sets the port type of the new connection to "*Trunk*". IBM Wave uses this method because the z/VM Guest is already connected to VSwitch "X" through a VNS that is assigned a VLAN ID. Therefore, the new connection must "ride" on the same Virtual NIC address.
- If the user selects VSwitch "Y", IBM Wave automatically sets the port type to "*N/A*". IBM Wave sets the port as "*N/A*" because VSwitch "Y" is defined as

VLAN unaware. The same action occurs if the user connects Guest B to VNS 3. This action is because VNS 3 is only connected to Guest LAN "Z", which is not VLAN aware.

- However, if the user connects Guest "C" to VNS "1", IBM Wave allows the user to choose the port type ("*Access*" or "*Trunk*"). IBM Wave uses this method because VSwitch "X" is linked to VNS "1" and is defined as VLAN aware and VNS "1" is configured with a VLAN ID. If VNS "1" was not defined with a VLAN ID, IBM Wave can allow only the "*Access*" port type connection to be made.

When you click **Go**, IBM Wave connects to each (checked) z/VM Guest.

Notes:

1. For trunk connections, IBM Wave always configures the base interface (eth0, eth1, and more) as void of any IP parameters. All IP parameters are configured on the VLAN interface (vlanx, eth0.x, and more).
2. If you attempt to connect a guest to a disabled VNS, the connection is made, but the BTS cannot use it for communication to the guest (because the VNS is disabled).

Disconnecting z/VM guests from virtual networks

Use the following process to disconnect a z/VM Guest from a virtual network segment (VNS).

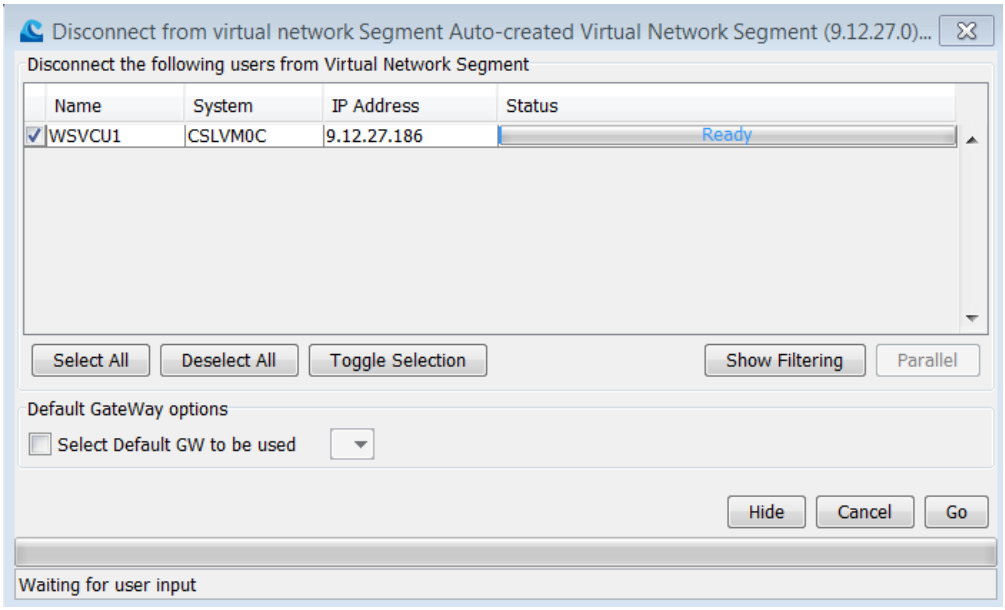

Figure 118. Disconnect a guest from a VNS

1. Open the **Network Viewer**.
2. Right-click on the connecting line between the guest and the VNS, and then click **Disconnect from the *Virtual Network Segment***.
3. If there is more than one VNS, select the virtual network segment that you want to disconnect.

   Note: You cannot disconnect a guest from the last connectable VNS. If you attempt to do so, IBM Wave displays an error message.

4. Click **Go**. IBM Wave disconnects the specified VNS from the guest.

Remove from DB

The remove from database action is relevant to only inconsistent Virtual Networks. The option lets you delete the Virtual Network from the IBM Wave database.

Read, update, or delete an IAN

Use this action to read or update an (Intelligent Active Note) IAN attached to this Virtual Network.

Recreate virtual network

This action only applies to inconsistent Virtual Networks. Because IBM Wave keeps all necessary information on Virtual Networks, it is possible to recreate a deleted Virtual Network that was deleted outside of IBM Wave. Use the **Recreate virtual network** option to do so.

Connect to router

Use this action to connect a Guest LAN to a router Machine. For a detailed explanation, see "Add New Guest LAN" on page 133.

With all visible connected users

This submenu allows the IBM Wave user to perform actions on all the visible z/VM Guests that are connected to the selected Virtual Network. The actions provided in this submenu are the same as the actions in a menu that would have been shown if all of the z/VM Guests connected to this Virtual Network were selected.

For more information about these actions, see "Overview of multiple task actions" on page 105.

Prototype functions

The prototype actions and tasks described in this information apply to a specific prototype. To access the prototype function, right-click on the prototype, and then select the action you want to perform.

Associate z/VM guest

Use this action to associate a z/VM Guest with a Prototype. This is especially useful after an auto-detection of a new z/VM System. Since IBM Wave cannot make the association by itself, the IBM Wave User has to do it.

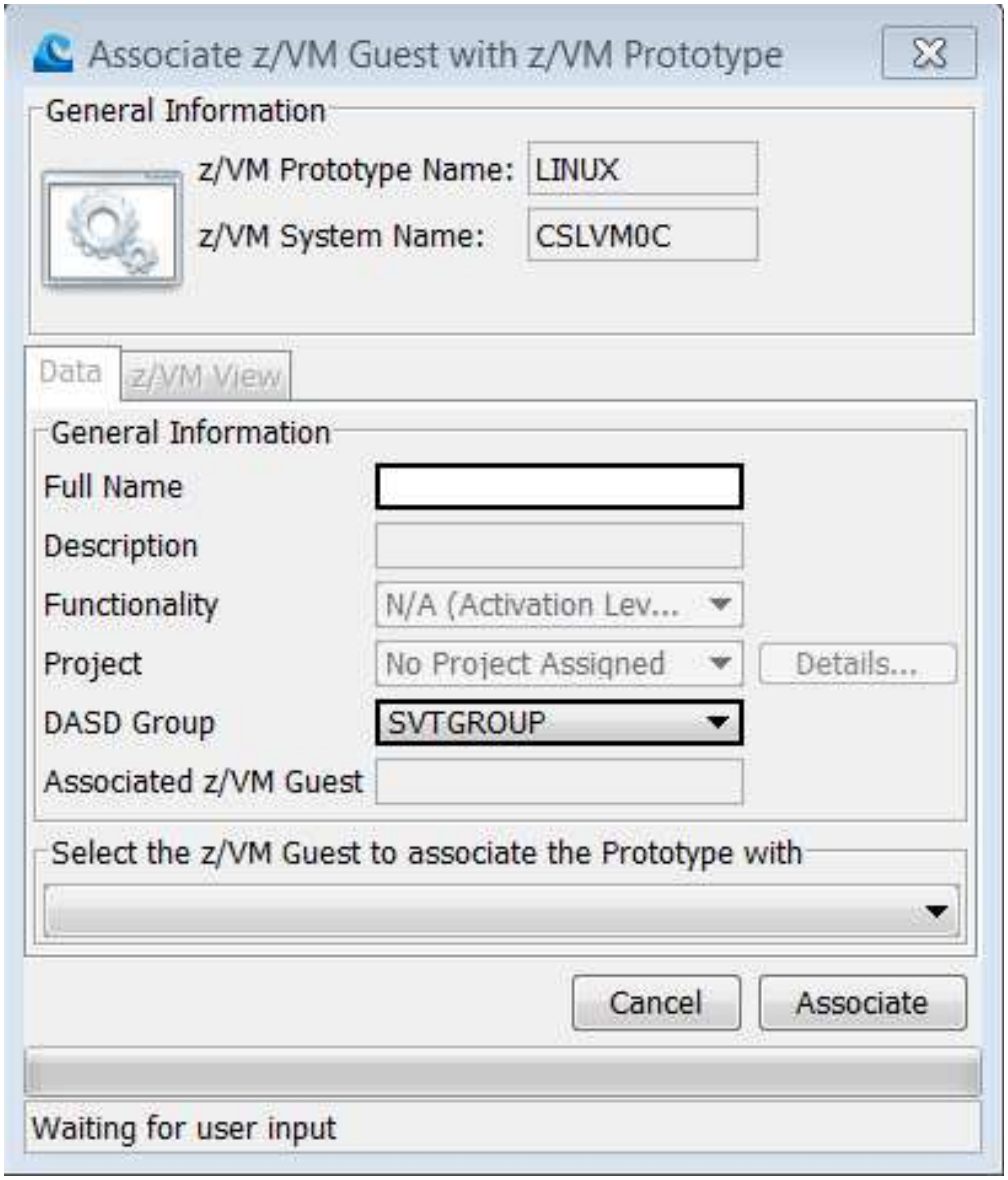

Select a z/VM Guest from the drop-down box. This action will mark the chosen z/VM Guest as a prototype User, and will disable logon to it by setting its z/VM password to "NOLOG".

Note: If a z/VM Guest exists in the z/VM System that has the same name as the prototype, it will be automatically selected when the window opens.

Once this action is complete, the chosen z/VM Guest will not appear in any view.

Note: IBM Wave will compare the selected z/VM Guest directory entry to the prototype user directory. All device-specific entries in these directory entries (MDISK, LINK, SPOOL, NICDEF, etc.) must match each other (not necessarily in the same order). Other definitions (Permissions, password, CPUs etc.) are taken from the prototype definition and these are ignored from the z/VM Guest directory entry.

Disassociate z/VM guest

Use this action to disassociate a z/VM Guest from a z/VM Prototype. The Window opened will ask for a new z/VM Password for the z/VM Guest (since once a z/VM Guest is associated with a prototype, its password is changed to "NOLOG"). Once the "Go" button is pressed, the z/VM Password will be changed, the z/VM Guest will be disassociated with the prototype and the z/VM Guest will be visible again from the z/VM Guest and Groups viewer (and the Network Viewer if this user has connections to Virtual Networks or VSwitches).

Note: Since it is possible to associate the same user with more than one z/VM Prototype, IBM Wave will only change the z/VM password for the z/VM Guest if there are no other z/VM Prototypes associated with the z/VM Guest. The same applies to the z/VM Guest being visible again in the other viewers - if there are other z/VM Prototypes associated with the z/VM Guest, the z/VM Guest will not be visible.

Clone from a Prototype

The following information explains how to clone from a prototype.

To create a clone from a prototype, select the Linux prototype that you want to clone, right-click, and select **Clone from a Prototype**.

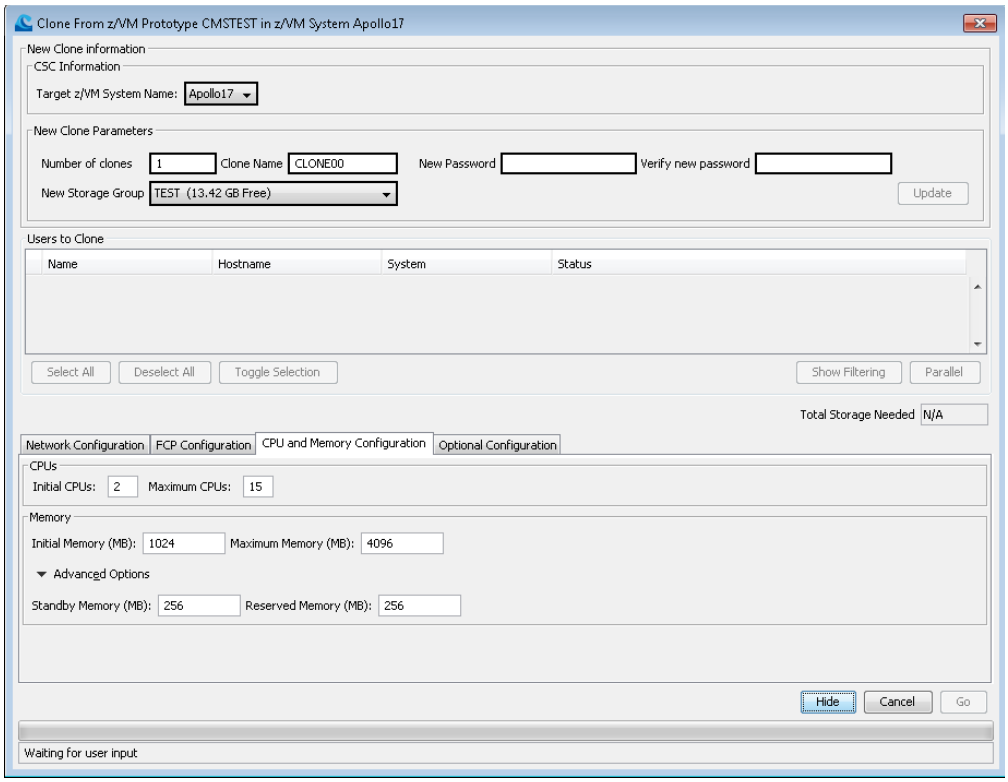

Figure 119. Clone from a Prototype

New Clone information

CSC Information

Contains the Target z/VM System from which you are creating the Cross System Clone (CSC).

New Clone Parameters

Number of clones
Enter the number of cloned virtual servers that you want to create.

Clone Name/Base Name
The field is editable and contains either:

- For one clone, the new z/VM Guest name.
- For more than one clone, the base name of the z/VM Guests.

Because the limit for the z/VM Guest's name is 8 characters, the clone's base name is limited to 8 characters.

New Password
Enter the new z/VM Guest password for all the new virtual servers created.

Verify New Password
Enter the new z/VM Guest password again to verify that it is correct.

New Storage Group
By default this field indicates the z/VM DASD Storage Group that is associated with the prototype. You can select a different z/VM DASD Storage Group.

Total Storage Needed
The **Total Storage Needed** field indicates the total amount of disk space, in GB. After you complete the information for your clones, and then click **Update**, the field is updated. The disk space is calculated by multiplying the total amount of disk space that is needed for the prototype by the number of clones requested (value in number of clones).

Users to Clone Each virtual server contains the following entries:

Name The new clone name.

Hostname
The new host name. The field is editable.

System
The z/VM System on which the new z/VM Guest is stored.

Auto-created Virtual Network Segment
The field represents the virtual network segment (VNS) to which the new z/VM Guest is connected. IBM Wave automatically suggests an available IP address, based on the VNSs to which the new clone is connected. The IP address can be changed. When you select a VNS from the "Network Configuration" table, the title of the column changes to reflect the name of the VNS. The cells are populated automatically with the IP address on the VNS' IP segment. You can connect up to three VNS to a clone. If the new virtual servers are connected to more than three segments, a warning message appears.

Important: IBM Wave identifies the IP addresses of guests that it manages. However, if you set up a remote Wave server, it has no way of identifying the correct IP addresses when you clone a user. Therefore, when the Wave server is remote and you clone a user, the automatically suggested IP address might belong to the remote Wave Server. If so, it needs to be manually changed. (A remote IBM Wave server is one that is not installed on the same z/VM system where the cloned guest is being deployed.)

When you select the VNS from the table in the "Network Configuration" tab, you can select the Virtual Network that is linked to the VNS by clicking the Virtual Network cell. If the VNS is linked to more than one Virtual Network, a list opens up for selection. You can also select the port type that the connection is assigned. Access and TRUNK are available for VLAN-aware connections, and N/A for a regular connection.

Note: You cannot connect a new clone to two different VNSs that are connected to the same Virtual Network.

Status The status of the first phase of the cloning process.

Attempt to reuse selected network configurations for clones
The check box, at the bottom of the window, can enable the reuse of valid `ifconfig` files on clone interfaces that connect to the same VNS as the master.

- When the check box is cleared, IBM Wave removes all existing `ifconfig` files when the guest is cloned, and creates new ones for each interface that is specified for the clone. This behavior is the default when cloning guests.
- When the check box is selected, IBM Wave attempts to identify the common network connections between the master and the clone. Instead of re-creating the `ifconfig` files for common interfaces, IBM Wave attempts to reuse the existing master's `ifconfig` files to preserve any network configurations set by the user. With this implementation, only the IP address, netmask, host name, and gateway configuration values are updated for each interface connection that IBM Wave recognizes as reusable. If any network connections are not recognized by Wave as reusable, they are re-created.

Network Configuration options - The Network Configuration table lists of the available VNS. The first check box is used to indicate that the new cloned virtual server is connected to the selected VNS. The "Default Gateway" check box is used to indicate which VNS is the default gateway VNS for the cloned virtual servers. If a specific VNS is connected to more than one Virtual Network, it is possible to select which Virtual Network the new cloned z/VM Guest is connected.

Note: If the target z/VM System changes, the VNS table and the list of available DASD Storage Groups changes to present the VNS and DASD Storage Group in the target z/VM System.

Click "Go" to begin the cloning processing. Each new clone is created serially. For each new clone, after the first phase is complete, IBM Wave generates work units. The work units can be viewed in the "Work Unit Viewer " minor view in the "General Status Viewer".

FCP Configuration options - After you select the number of clones for the process, the FCP Configuration area is populated with suggestions for new FCP devices. Each target clone has its own tab. For each clone, it is possible to select the new FCP devices that are assigned to it. The FCP device table shows the real FCP device address, the virtual address that is used to dedicate this device to the clone, and the source WWPN of the device. It is possible to add, remove, and change the devices by right-clicking in the table, and then selecting the appropriate action.

Devices are automatically suggested based on the following rules:

1. If the Source Guest/Prototype uses dedicated FCP devices, IBM Wave attempts to locate an FCP Device Pool whose "Default Virtual Device" matches the virtual device that is used for the dedicated device. If the Device Pool exists, and enough free FCP devices are available in the Device Pool, the table is populated with new devices from the device pool. The process is repeated for every FCP device that is dedicated to the source guest/prototype.
2. The table can also be manually completed either by using the "Global" option to add an FCP Device for all the clones in the process. Or use the "Add Device" option by right-clicking in the table.
3. Devices can be changed, removed, or added if required. For example, if a specific source WWPN is required.
4. If no FCP devices are defined for the z/VM System, the FCP Configuration tab is disabled. The status is reflected in the text that is displayed when the tab is selected.
   - When IBM Wave cannot find a suitable FCP Device Pool (in number 1), a warning message is issued. FCP devices can then be manually added to the target clones.
   - When IBM Wave finds a suitable FCP Device Pool, but not enough free devices to satisfy the request, a warning message is issued. You can decide whether to decrease the number of clones, or use a different device pool.

CPU and Memory Configuration

Initial CPUs
The initial number of CPUs assigned to the z/VM guest.

Maximum CPUs
The maximum number of CPUs that are allowed to be assigned to the z/VM Guest.

Initial Memory (MB)
The initial amount of memory, specified in MB, that is allocated to the z/VM Guest.

Maximum Memory (MB)
The maximum amount of memory, specified in MB, that can be allocated to the z/VM Guest.

Advanced Options (for Memory) - controls adding memory dynamically. To add memory dynamically, you must set Standby Memory or Reserved Memory. The options are collapsed by default, but shown in Figure 119 on page 141 for illustrative purposes.

Standby Memory (MB)
The amount of standby memory.

Reserved Memory (MB)
The amount of reserved Memory

Important: Dynamic Memory reconfiguration is only supported when the initial memory size for the guest is an exact multiple of the memory block size. To understand the calculation of memory block size, see the following topics:

- "Overview of estimating memory and CPU requirements" at www.ibm.com/support/knowledgecenter/SSB27U_6.2.0/com.ibm.zvm.v620.hcpl0/memcpuov.htm.

- "Steps for estimating memory and CPU requirements" at http://www.ibm.com/support/knowledgecenter/SSB27U_5.4.0/com.ibm.zvm.v54.hcpl0/hcsx0b3030.htm.

Optional Configuration

Select the IBM Wave Script to run after clone
Script Name
An optional script that you can select that runs during the first boot of the cloned server. To open the WAVE Script Manager, click **Browse...**, and then select a Linux shell script.

Specify REXX to run after z/VM Guest creation
REXX Name, Machine Name, and Minidisk Address
Use these fields to specify a REXX EXEC that runs as part of the clone process, per cloned z/VM Guest. The EXEC runs immediately after the new z/VM Guest is created, but before DASD storage is copied.

Dedicated Devices
Use the list to select a Device Pool from which to obtain the new dedicated devices, if the clone-source prototype or z/VM Guest uses dedicated devices. You can also use the same dedicated devices that are used by the source z/VM Guest or prototype.

Descriptive fields

Project
This field allows the IBM Wave user to assign a project to the newly cloned virtual servers. By default, this field contains the project that is associated with the prototype.

Functionality
A descriptive field for the functions of the newly created cloned virtual servers.

Description
A short description for the newly created virtual servers.

Optional Linux parameters

Domain
The domain of the new virtual servers.

Regenerate SSH Keys
When selected, IBM Wave removes all files that begin with "`/etc/ssh/ssh_host_`", which causes the SSH daemon to regenerate the files during startup.

When selected on an Ubuntu Server, the default HostKey files are regenerated automatically by IBM Wave. Any HostKey statements in `/etc/ssh/sshd_config` are removed to enforce the defaults.

Important: Using the same dedicated devices as the source z/VM Guest or prototype can lead to problems when both the source z/VM Guest and the cloned z/VM Guest are activated concurrently. Use this option only in cases where a mirror image of a z/VM Guest must be created.

Click **Go** to begin the cloning processing. Each new clone is created serially. For each new clone, after the first phase is complete, IBM Wave generates work units. View the work units in **General Status Viewer** > **BTS Work Units**.

Clone logic

IBM Wave clones from the source to the target (or targets) in two major phases:

1. Phase 1 - During Phase 1, IBM Wave performs various actions that are synchronous in nature. Progress is indicated on the clone pane:
   a. **Create the new Virtual Server** - For each new Clone, IBM Wave uses the SMAPI server to create a new Virtual Server that matches the Prototype on the target z/VM System.
   b. **Register new network information** - For each new Clone, IBM Wave registers the new IP address, hostname, startup script and others for the new cloned server.
   c. **Initiate the background copying process** - IBM Wave submits a request to the BTS, which starts the background disk copying.
2. Phase 2 - Phase 2 happens in the BTS. In this phase, the GUI client submits a BTS Workunit that contains DASD copy requests. Such BTS Workunits are created per cloned z/VM Guest. For example, if the process is cloning three new z/VM Guests, three BTS work units are generated. The BTS Workunit contains the following BTS requests:
   a. If a REXX is selected for execution as part of the clone process, a REXX Execution BTS Request.
   b. A DASD copy BTS Request for each MINIDISK or DASD Dedicate device in the source z/VM Guest (or prototype), based on the type of statement:
      1) “Clone MDisk from <source guest> to <target guest> vdev <vdev>” - For regular MDisk statements.
      2) “Clone MDisk from <source system> <source guest><Target system><target guest> vdev <vdev>” - For Cross System Clone requests.
      3) “Clone MDisk from <source guest><source vdev> to <Target guest><Target vdev>” - For minidisk statements that specify the “END” keyword in the statement.
      4) “Clone Dedicated DASD from RDEV <source real device> to RDEV <target real device>” - For dedicated DASD devices.

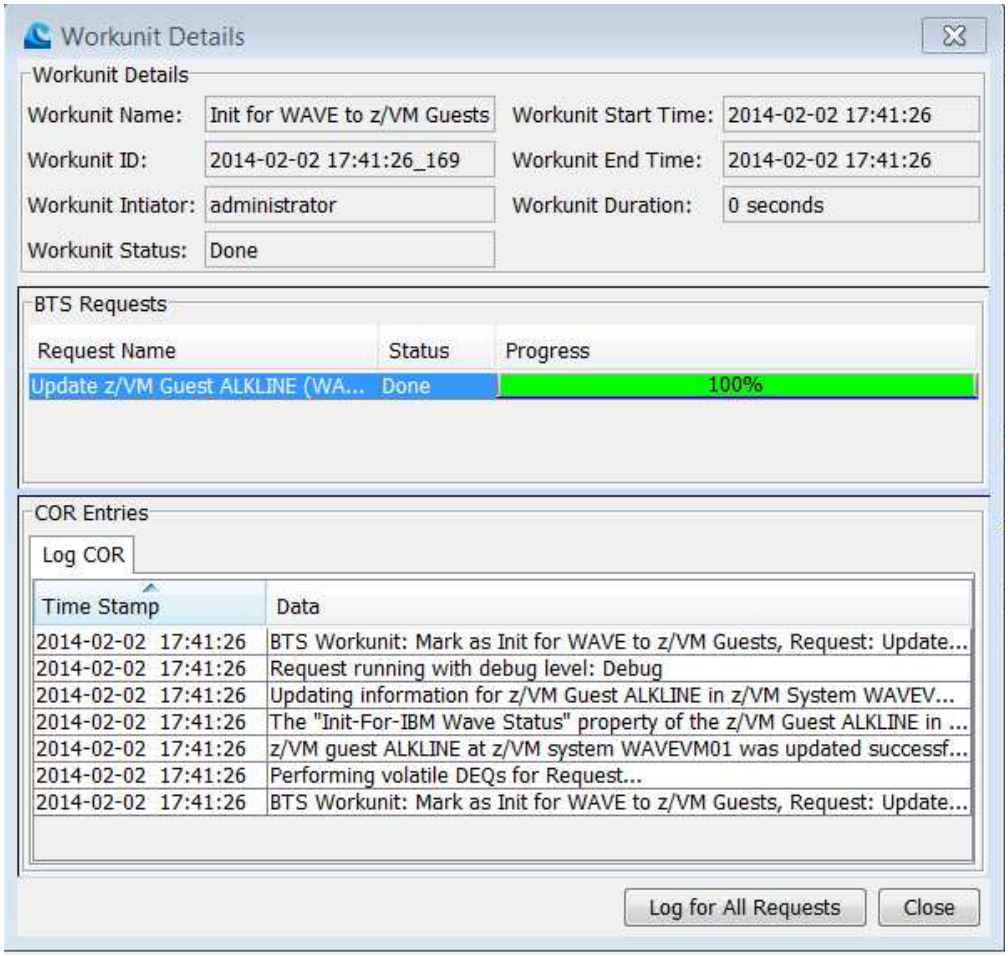

Figure 120. BTS workunit details for a clone request

The progress of the second phase can be tracked from the BTS Workunit Tab in the General Status Viewer.

When the process completes, the new z/VM Guest appears as disabled. When the Virtual Server is enabled for the first time, during the boot process, the following actions are taken:

- IBM Wave changes the IP address, hostname, and other network aspects of the new cloned server.
- An IBM Wave Post Clone exit runs locally on the new z/VM Guest.
- If an initial boot script was specified, it is run.

Delete

Use this action to delete a prototype. This action is relevant both for associated prototypes and for disassociated prototypes. When a disassociated prototype is deleted, only the prototype directory manager entry is removed. When an associated prototype is deleted, the following panel appears:

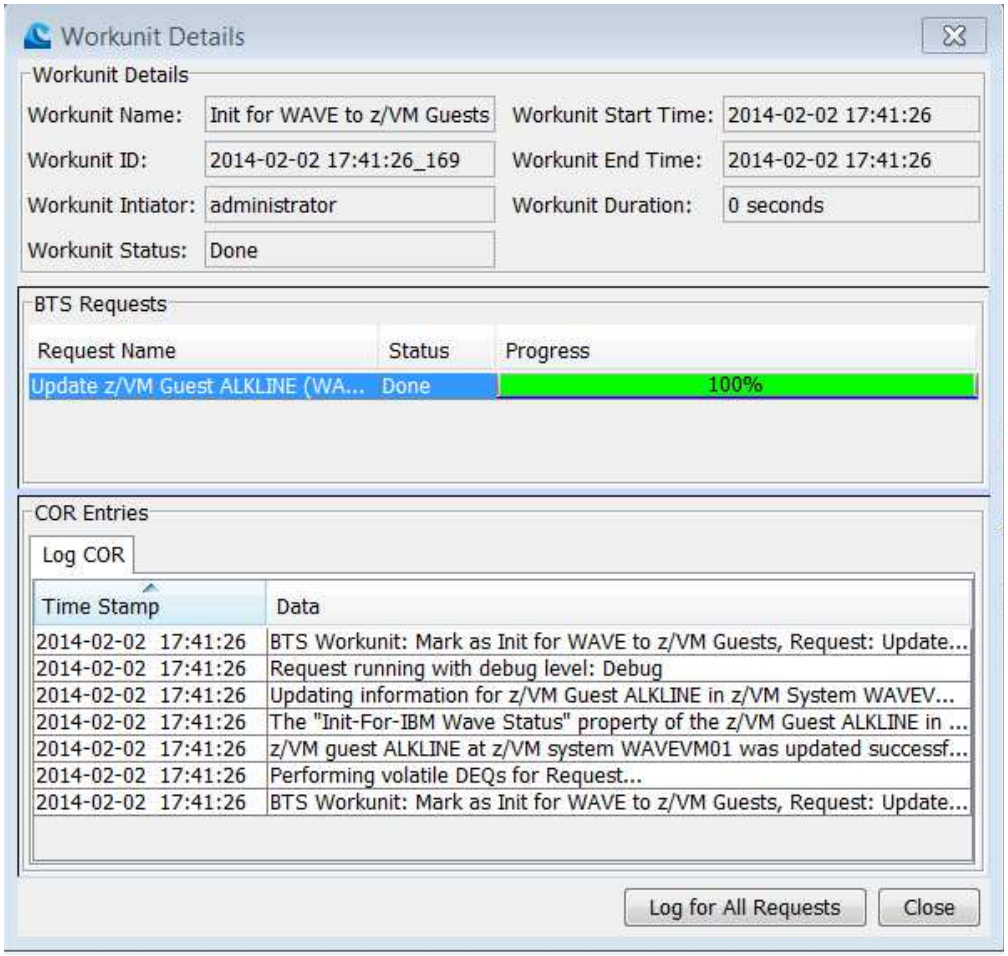

Since the prototype is associated with a z/VM Guest, IBM Wave needs to know what to do with that z/VM Guest. The IBM Wave User can choose to delete the associated z/VM Guest altogether, or to keep the associated z/VM Guest. If the latter is chosen, IBM Wave will change the z/VM Guest's z/VM password to the value specified in the "New VM Password for User" field. IBM Wave will remove the prototype mark from this z/VM Guest and it will be displayed in the z/VM Guests and Groups Viewer and the Network Viewer (if relevant).

Display information

Use this action to display information about the prototype.

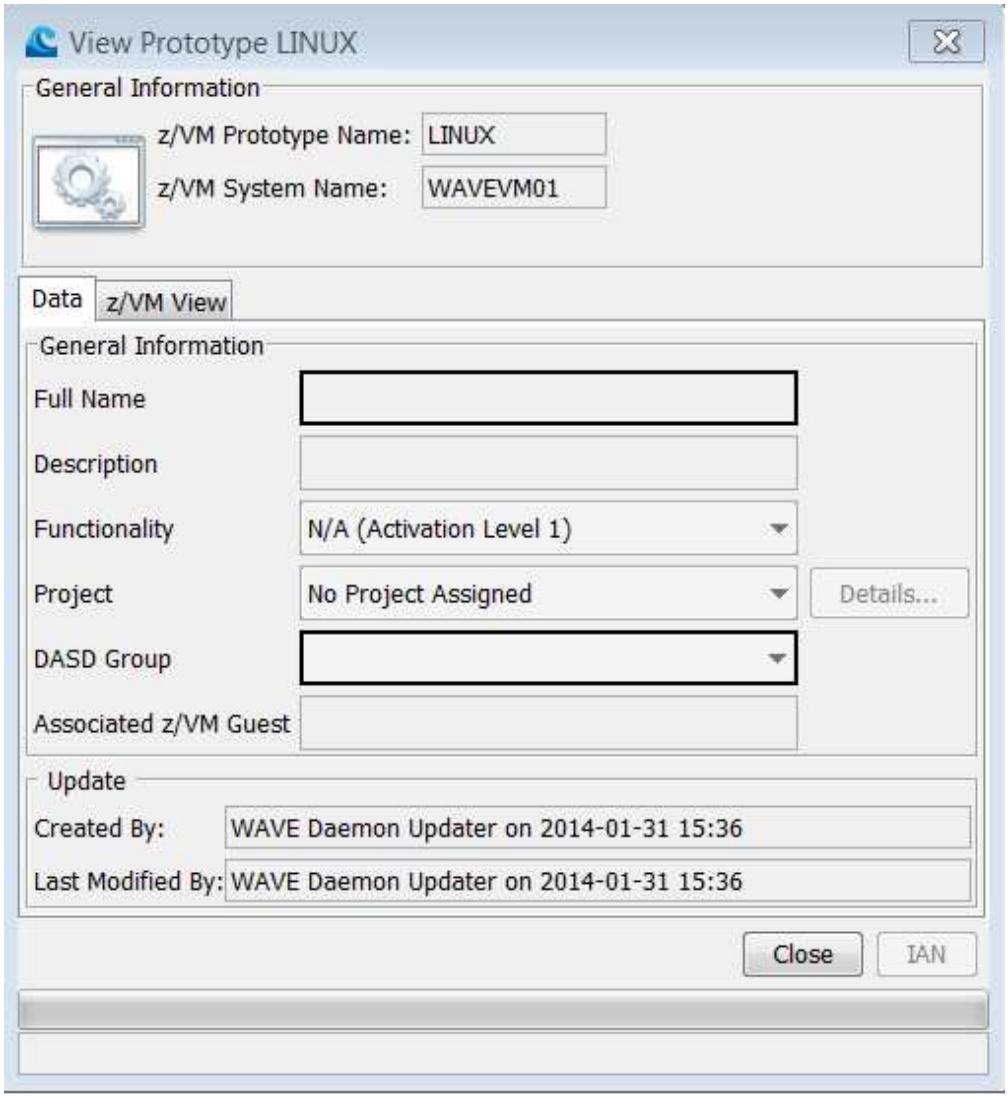

The fields in this window are:

- **Full Name** - The full name of the prototype.
- **Description** - An optional description.
- **Functionality** - An optional functionality.
- **Project** - The project to which the prototype is assigned. Press the details button to view details on the project.
- **DASD Group** - The DASD Storage Group from which to take disk space when cloning from this prototype.
- **Associated VM User** - The z/VM User associated with this prototype.

The z/VM tab displays the directory entry for this prototype.

Duplicate prototype

Use this action to duplicate a prototype entry. This action will only duplicate the directory manager prototype entry. The action is especially useful, if one prototype has to serve several projects, or if the same prototype has to take disk space from several different DASD Storage Groups.

- Prototype - The new name for the prototype.

- Full Name - The new full name for the prototype.
- Description - The new prototype's description.
- Functionality - The new prototype's functionality.
- Project - The new prototype's project.
- DASD Group - A new DASD Storage Group for the prototype.

The VM tab displays the actual directory entry for this prototype.

Duplicate prototype user definition

Use this action to perform a physical clone of the prototype's User.

This action has the same effect as duplicating the z/VM definitions of a z/VM Guest. For more information, see "Duplicate z/VM User Definition" on page 85.

Read/update/delete IAN

Use this Action to read or update the IAN attached to this z/VM Prototype.

Lock/unlock

Use this action to lock or unlock the z/VM Prototype.

Update information

Use this function to update some parameters of the prototype. For a complete description of the fields, see "Display VNS information" on page 168.

Storage functions

All the following actions are relevant to Storage administration. There are three types of storage functions:

1. DASD Storage Group functions
2. DASD Volume functions
3. DASD Emulated device (EDEV) functions

All of the actions are accessible from the popup menu of the selected object, except for "Add New DASD Group" which is accessible from the DASD Group viewer popup menu.

Note: The term "DASD Storage Group" refers to DIRAMINT DASD Groups and VM:Secure / VM:Direct Subpools.

DASD Storage Group functions

This section covers the following topics:

Add new DASD Group

To add a new DASD storage group to the z/VM System, go to the **Storage Viewer** > **Group** tab. Right-click in the white space, and then click **Add new DASD**

Storage Group. The **Create Storage Group** window opens.

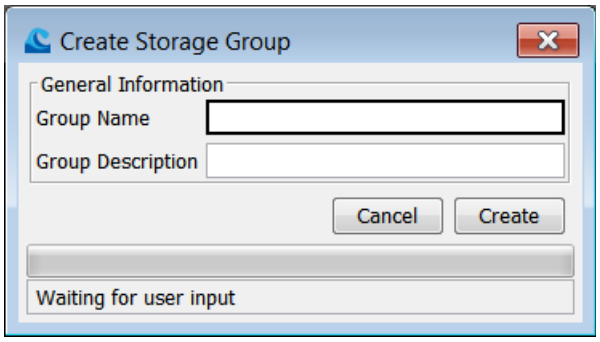

Figure 121. Add new DASD Group

- **Group Name** - The name for the new DASD Storage Group. The DASD Storage Group name must be unique for each z/VM System.
- **Group Description** - An optional description for the DASD Group.

When you click **Create**, the new DASD Storage Group is defined in the z/VM System. Because z/VM does not permit empty DASD Groups, the new DASD Group is assigned a dummy region. The dummy region does not show in the viewer. For consistency, the data displayed are the actual objects on the z/VM System.

Delete

Use this action to delete a DASD Group. This deletes the DASD Group in the z/VM System and reassigns the dummy region.

Lock/unlock

Use this action to lock or unlock the DASD Group.

Read or update the IAN

Use this action to read or update the IAN attached to the DASD Group.

Update information

Use this action to update information on the DASD Group. For a description of the fields see "Add new DASD Group" on page 150.

DASD volume functions

IBM Wave has the following types of DASD Volume actions.

- Multiple selection actions - Actions that can be performed on a number of DASD volumes.
- Single selection actions - Actions that can be performed on one DASD volume at a time.

Multiple selection actions

This section includes the following topics:

"Mark as page or spool volume" on page 155
"Remove page or spool volume indicator" on page 155
"Add to z/VM system as page or spool volumes" on page 155
"Vary online" on page 155
"Vary offline" on page 155

Attach or detach DASD volumes: Use this action to attach or detach DASD Volumes form the System. This action can only be applied on DASD Volumes with the "Free" status.

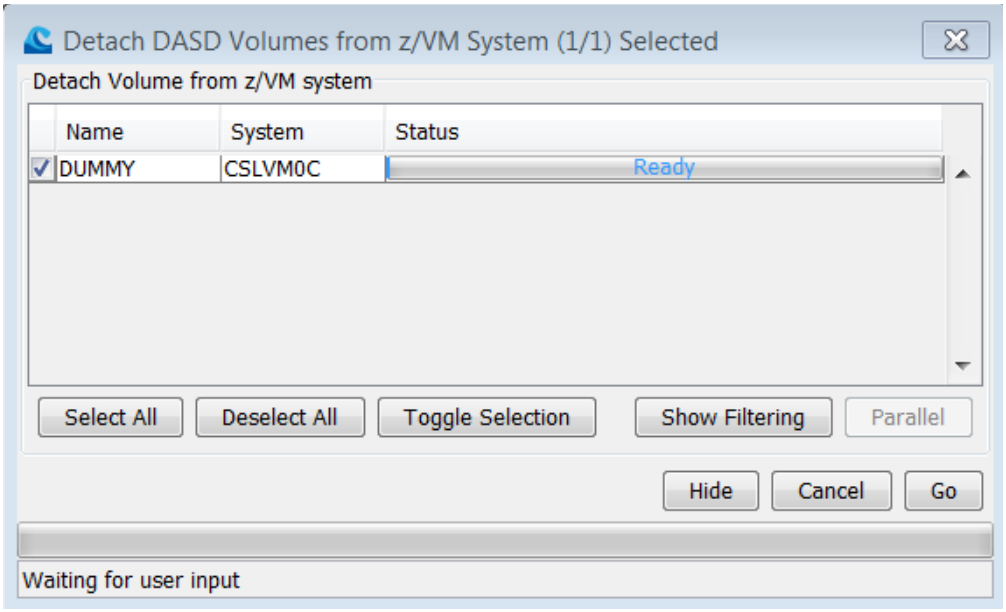

Define region: Use this action to define a DIRMAINT region on the selected DASD Volume. This action changes the status of the DASD Volume to "Defined". Although DIRMAINT supports multiple regions on a single DASD Volume, this functionality is not supported by IBM Wave. IBM Wave always defines a single region on the DASD Volume that is the size of the DASD Volume starting from the first cylinder (cylinder 1). IBM Wave creates this region with the same name as the DASD Volume serial number.

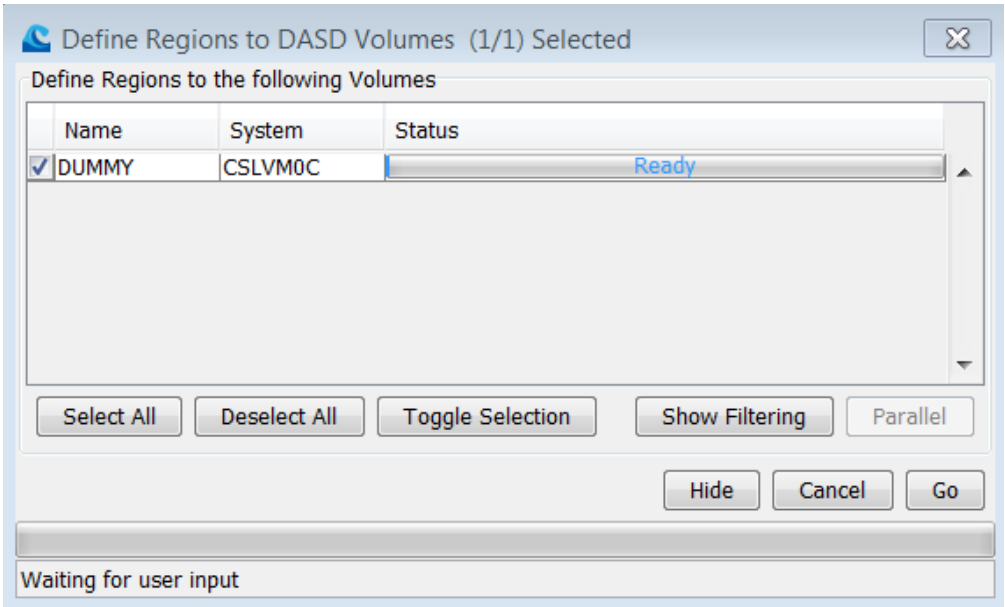

Note: If the z/VM System's directory manager is defined as VM:Secure or VM:Direct this option is unavailable.

Undefine region from volume: Use this action to unassign the DIRMAINT region assigned on the selected DASD Volume. This action changes the status of the DASD Volume to "Enabled".

Note: If there are other regions defined on the DASD Volume (regions that were defined outside of IBM Wave, since IBM Wave does not support multiple regions on DASD Volumes), only the region that has the same name as the DASD Volume serial number is deleted. Also, if the z/VM System's directory manager is defined as VM:Secure or VM:Direct this option is unavailable.

Assign DASD Volumes to DASD Group: Use the **Assign DASD Volumes to DASD Group** action to assign a DASD volume to a DASD storage group. Right-click on the defined volser, and then click **Assign DASD Volumes to DASD Group** and select the DASD storage group from the list of groups.

Unassign volume from DASD group

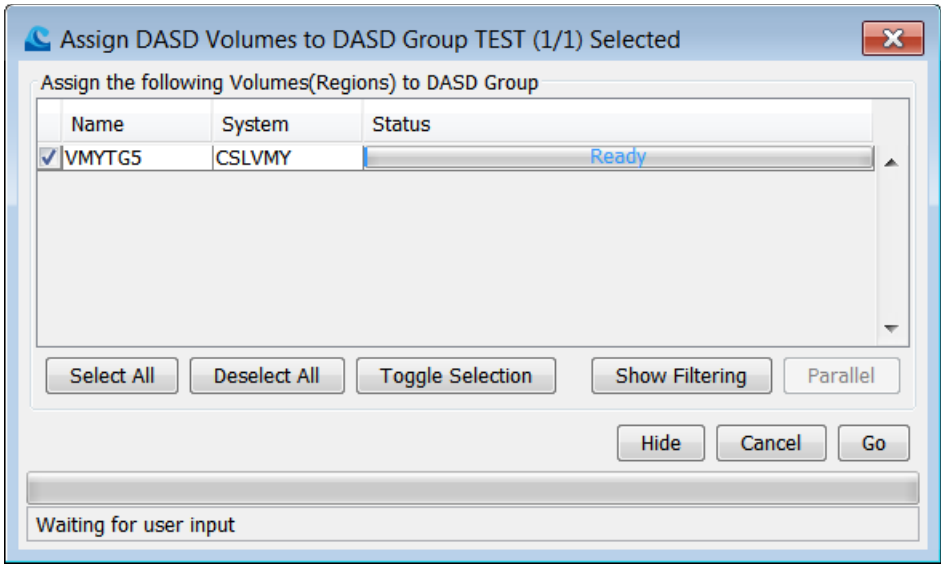

Figure 122. Assign DASD Volumes to DASD Group

Press **Go** to assign all of the selected DASD Volumes to the DASD Storage Group.

Unassign Volumes from DASD Group: To remove a DASD Volume from a DASD Storage Group, right-click on the assigned volume, and then click **Unassign Volumes from DASD Group**. Because IBM Wave supports only a one to one ratio between a specific Volume and a DASD Storage Group, DASD Storage Group selection is irrelevant. The DASD Volume gets removed from the DASD Storage Group in which it was assigned.

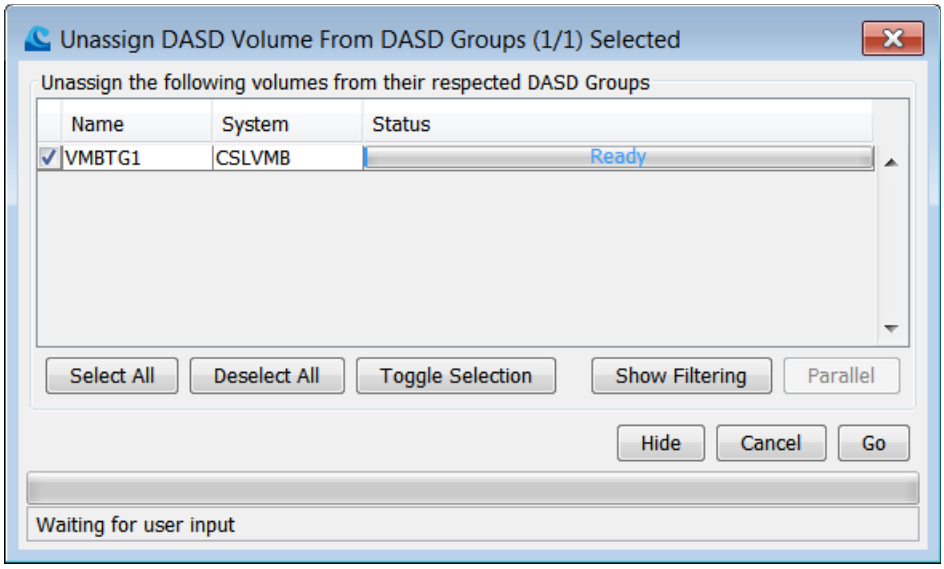

Figure 123. Unassign Volumes from DASD Group

Click **Go** to remove the selected volumes.

Lock or unlock: Use this function to lock or unlock the DASD Volume. Locking is used to indicate that the locked volume's status should not be changed. You cannot perform any action on a locked DASD Volume, except unlock. Assigned DASD Volumes cannot be locked.

After you press the **Go** button, IBM Wave marks all the (checked) DASD Volumes as locked.

Note: The lock is purely meta-data, and applies to IBM Wave functionality. There is no physical locking of the DASD Volume, so theoretically, the z/VM System programmer can still manipulate this DASD Volume from outside of IBM Wave.

Mark as page or spool volume: Use these actions to mark the DASD Volumes as pre-formatted Page, Spool, or both DASD volumes. The volumes can then be attached to the z/VM System using the **Add Page** or **Add Spool** z/VM System action.

Remove page or spool volume indicator: Use these actions to remove the "Page" or "Spool" indicator marks from one or more DASD Volumes. After the action is complete, the volumes are not allowed to be added to the z/VM System as Page or Spool Volumes.

Add to z/VM system as page or spool volumes: Use this action to add the selected DASD Volumes to the z/VM System as Page or Spool volumes. For more information, see "Add Page and Add Spool" on page 65.

Vary online: Use the "Vary Online" action to vary a DASD Volume online in a z/VM System. This action interacts with the XVDSKON user exit in order to perform the action on the z/VM System. For more information, see the XVDSKON user exit in .

Vary offline: Use the "Vary Offline" action to vary a DASD Volume offline in a z/VM System. This action interacts with the XVDSKOFF user exit in order to perform the action on the z/VM System. For more information about the XVDSKOFF user exit, see "XVDSKOFF - Before Varying DASD Volume offline" in .

Single selection actions

This topic covers:

- "Display Volume information"
- "Read or update an IAN" on page 156

Display Volume information: To view information about a DASD Volume, right-click on the Volser, and then select **Display Information**.

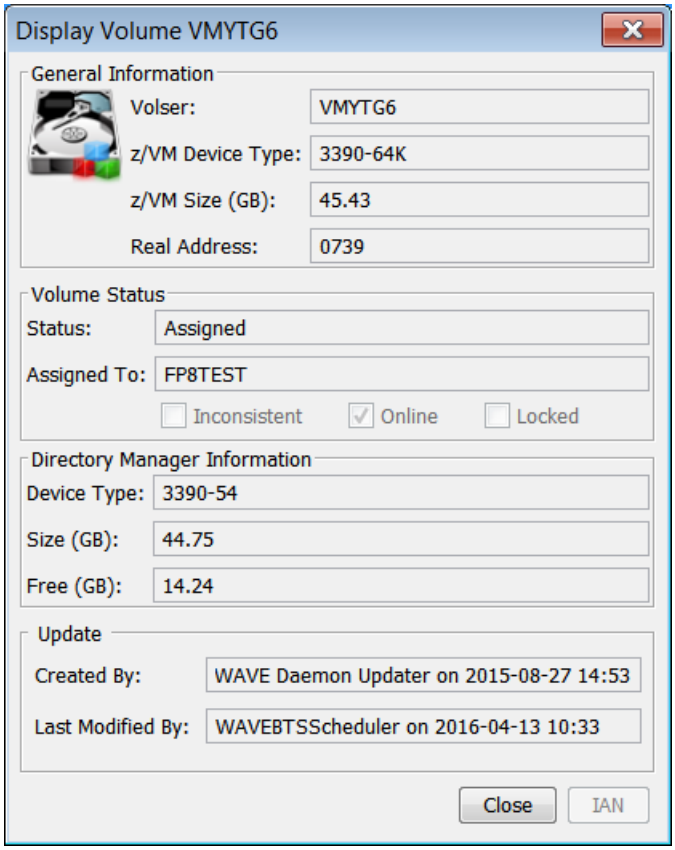

Figure 124. Display Volume Information

The fields in this window are:

- **Volser** - The VOLSER of the selected DASD volume.
- **Device Type** - The device type (3390 model).
- **Size** - The size, in GB, of the DASD volume.
- **Real address** - The real device address of the DASD volume.
- **Status** - The status of the DASD volume.
- **Assigned to** - If the DASD volume's status is "Assigned", the DASD storage group to which the volume belongs is displayed.
- **Inconsistent** - Indicates when the DASD volume is inconsistent.
- **Online** - Indicates when the DASD volume is online.
- **Locked** - Indicates when the DASD volume is locked.

Read or update an IAN: Use this action to read or update an Intelligent Active Note (IAN) that is attached to the DASD volume.

Emulated device functions

An emulated device (EDEV) is a representation of Fibre Channel Protocol (FCP) storage within z/VM, which emulates its appearance to the system as fixed-block architecture (FBA) volumes.

IBM Wave provides support for EDEV management. You can use the user interface to create and modify an EDEV and differentiate between real device volumes and EDEV volumes.

Prerequisites:

- z/VM 6.2 or higher (If you are running EDEV on systems before z/VM V6R3, verify that you have the correct service levels (PTFs) installed.)
- SCSI storage that is attached to a z/VM environment.

The following icon represents an EDEV volume in IBM Wave.

The following actions relate to emulated device (EDEV) functions:

- "Specify an EDEV Address Range"
- "Create an EDEV" on page 159
- "Delete an EDEV" on page 161
- "Display an EDEV" on page 162
- "Update information for an EDEV" on page 164
- "Discover an EDEV" on page 165

Specify an EDEV Address Range

You can specify a range of preferred addresses for an emulated device (EDEV) on each z/VM system. IBM Wave selects the first available address when it creates the new EDEV.

Note: The address value is not final. You can specify a different address manually.

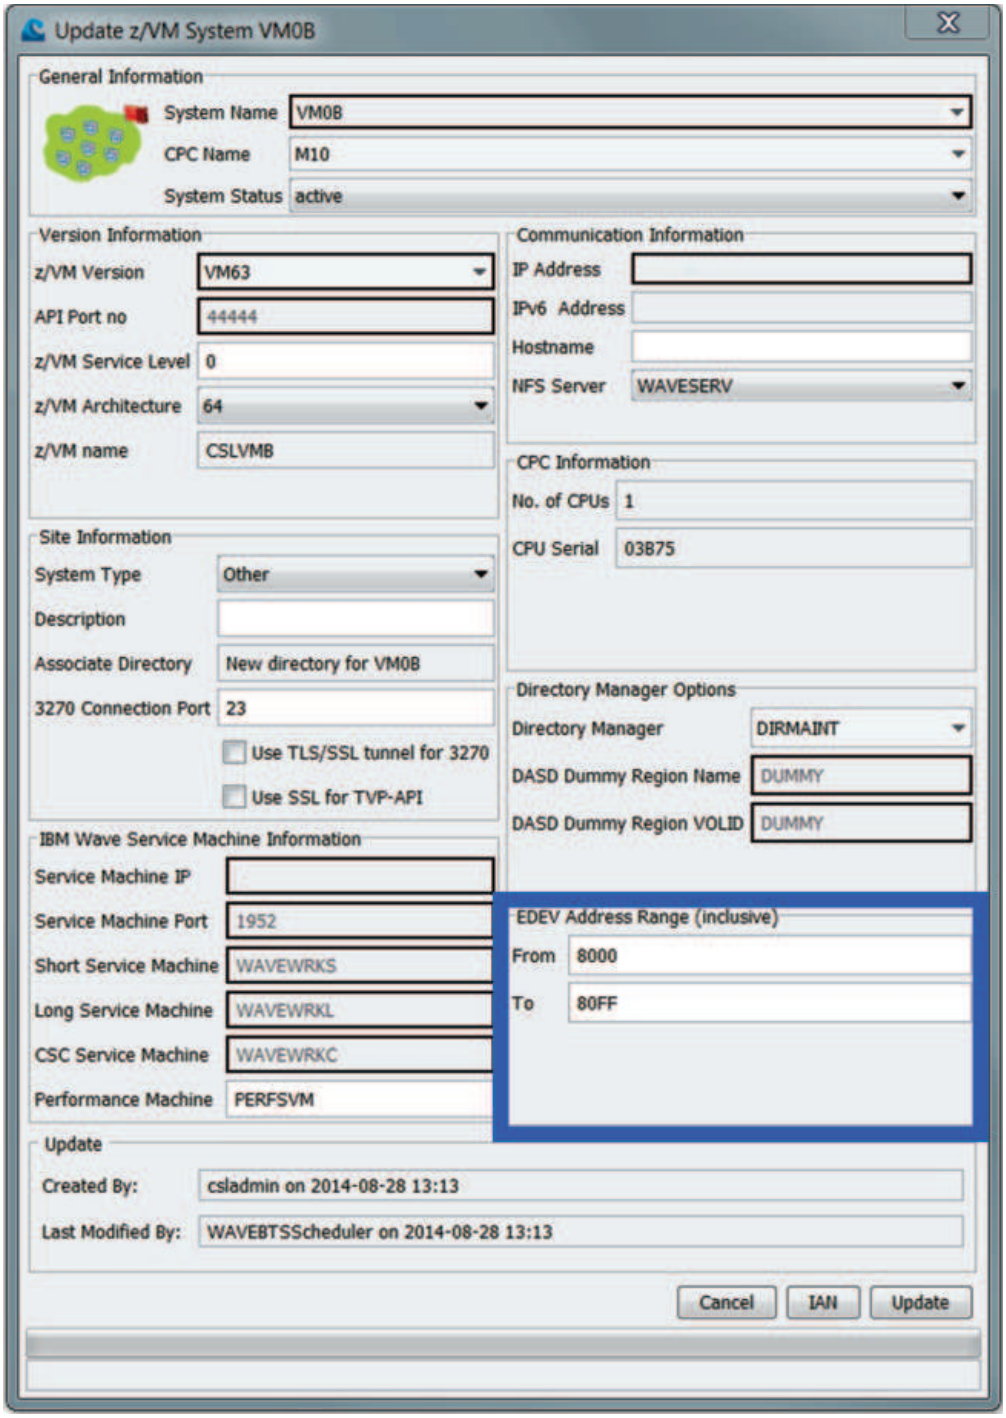

Figure 125. EDEV Address Range

EDEV Address Range fields:

- **EDEV Address Range (inclusive)** - The range of preferred addresses for the emulated device (EDEV) definitions that can be automatically used when you define an EDEV for SCSI storage on VM systems.
- **From** - The starting address range. The value must be four digits hexadecimal only.
- **To** - The end of the address range. The value must be four digits hexadecimal only.

Note: For a complete list of menu item descriptions, see "Add New System" on page 60.

Create an EDEV

To create an emulated device (EDEV), go to the **Current System Viewer** and select **Storage > Volumes** tab. Right-click in the white space to display the **Add SCSI Storage** option. After you choose the **Add SCSI Storage** option, the **Add SCSI Storage (EDEV)** menu appears.

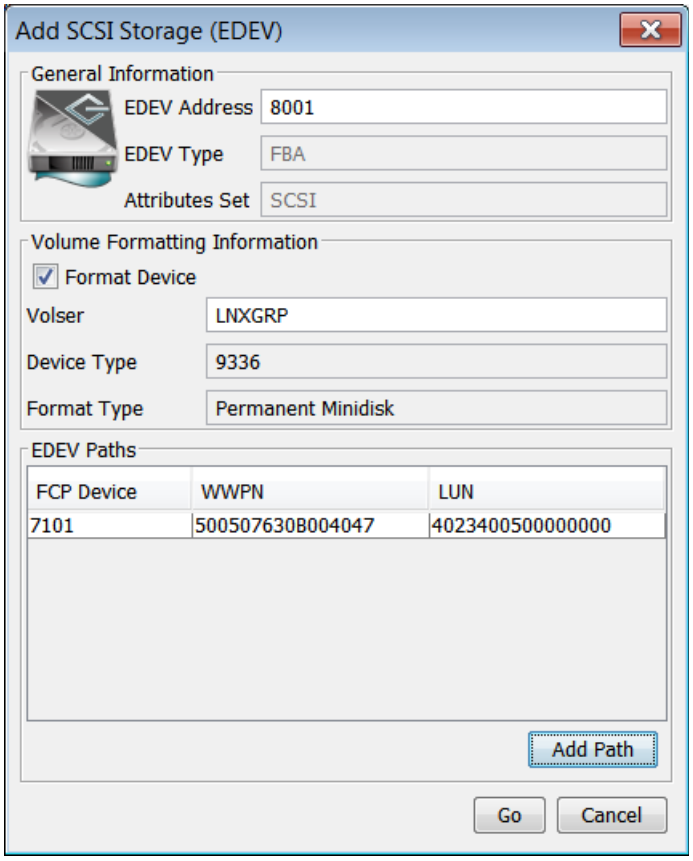

Figure 126. Add SCSI Storage (EDEV)

Fields in this pane:

- **EDEV Address** - The address at which to define an EDEV.
- **EDEV Type** - Currently only FBA is supported.
- **Attribute Set** - Currently only SCSI is supported.
- **Format Device Checkbox** - When selected, the new EDEV is formatted upon creation. (The amount of time this action takes to complete depends on the size of the device to format. It can take several minutes to complete.)
- **Volser** - six-digit alphanumeric serial of the volume to be formatted.
- **Device Type** - The device type of the DASD volume. Currently, only 9336 is supported.
- **Format Type** - Currently only Permanent Minidisk is supported.
- **EDEV Paths** - A table that contains predefined FCP device, worldwide port name (WWPN), and logical unit number (LUN) information for each EDEV path that is specified.

- **Add Path** - A feature to help you automatically or manually add FCP paths for a new EDEV.

When you click **Add Path**, the **Add FCP Path for New EDEV** menu opens:

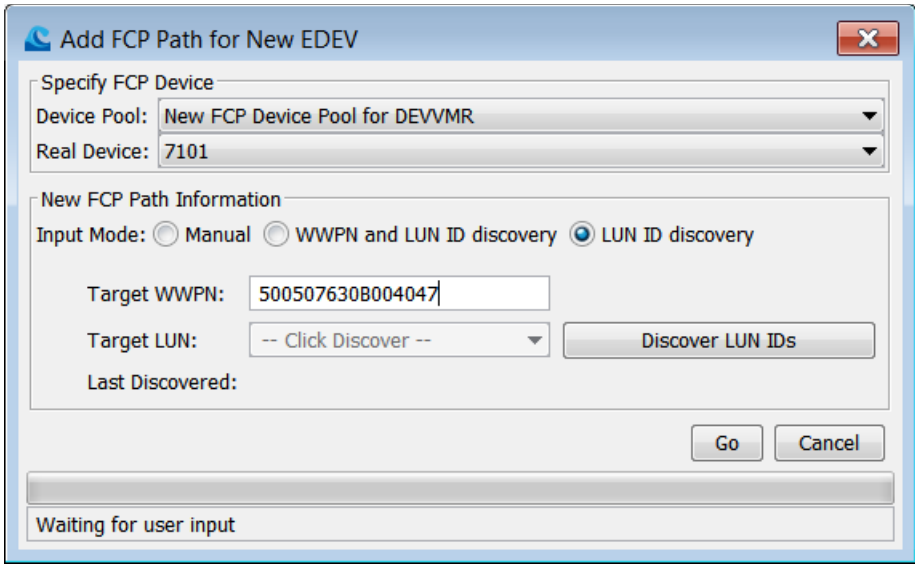

Figure 127. Add FCP Path for New EDEV

Fields in the "Specify FCP Device" pane:

- **Device Pool** - FCP Device Pool from which to select a real device.
- **Real Device** - The FCP device that is used to obtain the target WWPN and Logical Unit Number (LUN) ID.

For the **Target WWPN** and **Target LUN ID** fields in the "New FCP Path Information" pane, you can select the "Input Mode" as:

- **Manual**: When you know the target WWPN and LUN IDs, select the **Manual** option and type in the WWPN and LUN IDs.
- **WWPN and LUN ID discovery**: When you want IBM Wave to discover the WWPN and LUN ID, use the **WWPN and LUN ID discovery** option. Click **Discover WWPN and LUN IDs**.
- **LUN ID discovery**: When you know the WWPN and want to discover the LUN ID, use the **LUN ID discovery** option. Type in the WWPN, and then click **Discover LUN IDs**.

Note: Depending on your environment, the discovery of WWPN and LUN identifiers can take several minutes. However, the discovery of only the LUN identifiers, when you already know the WWPN, can be more efficient.

After you are satisfied with the SCSI selection, click **Go**. IBM Wave validates the path.

- When the validation is successful, the Unique ID and size of the disk that was found displays to confirm that it is correct. Confirmation happens on the first path. Subsequent paths bypass the confirmation step. When one or more paths are specified, they appear in the original menu and show all paths simultaneously before IBM Wave attempts to create the EDEV.
- When the validation fails, an error message appears with a dialog with LUN or WWPN error details, or both.

Remove a path: Before you click **Go** to create an EDEV on the **Add SCSI Storage (EDEV)** pane, you can also remove a specified path. Right-click on a row in the **EDEV Paths** table, and select **Remove Path**.

What happens during validation: IBM Wave attempts to discover the available SCSI paths during data input. When you select the FCP device, a request for discovery is sent to the background task scheduler (BTS). The request uses the standard SCSI discovery protocols to discover all of the available WWPNs and LUNs that are accessible through the specific FCP device.

- When the discovery request ends successfully, the WWPN field is populated with every WWPN that is accessible through the specific FCP device. When you select a WWPN, the LUN field is populated with every LUN that is accessible through the specific FCP device and WWPN. After you complete the final LUN selection and click **Go**, the Unique ID and size of the disk that was found displays to confirm that it is correct.
- If the discovery fails, for any reason, the WWPN and LUN fields are replaced with fields to manually input the data. After you input the WWPN and LUN data and click **Go**, a validation request is sent to the BTS. The validation request verifies that the specified WWPN and LUN are accessible through the specified FCP device. If so, IBM Wave displays the Unique ID and size of the disk to confirm that it is correct.

Delete an EDEV

You can delete an EDEV from a z/VM system.

From the **Storage > Volume** tab, right-click an EDEV volume and select the menu option to delete the device. The **Delete EDEV Volume** *name* menu opens:

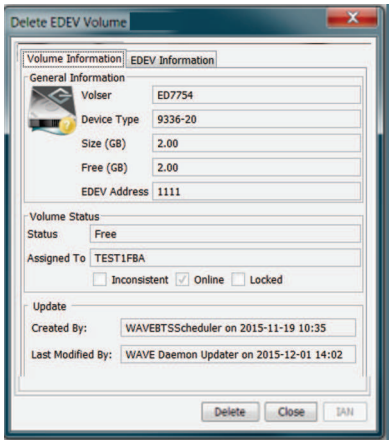

Figure 128. Delete an EDEV

Fields in this pane:

- **Volser** - The six-digit alphanumeric serial of the volume.
- **Device Type** - The device type of the DASD volume.
- **Size (GB)** - Total size of the FCP volume.
- **Free (GB)** - Free space on the FCP volume.
- **EDEV Address** - The address for an EDEV definition that can be automatically used when you define an EDEV for SCSI storage on VM systems.
- **Status** - The unassigned or assigned status of the volume.
- **Assigned To** - The name of the z/VM user.

- **Inconsistent/Online/Locked** - The status of the volume.

This menu is identical to the "Display an EDEV" with the addition of a **Delete** button. With this option, you can review the volume's information and confirm that it is the correct volume to delete before you delete it.

Display an EDEV

You can see detailed information about an emulated device (EDEV), such as the number of blocks that are defined or used.

An EDEV is viewed in the same way that a DASD volume is viewed. You can either:

- Double-click the volume wanted
- Right-click the volume, and then select **Display Information**.

Using either method, opens the **Display Information EDEV Volume** menu.

The **Display Information EDEV Volume** menu has two tabs:

1. **Volume Information**: Identical to the information that can be displayed for traditional DASD volumes in IBM Wave. However, the title bar and volume icon varies from what is used in the display menu for Extended Count Key Data (ECKD) volumes.
2. **EDEV Information**: Contains specific information for an EDEV including all defined paths:

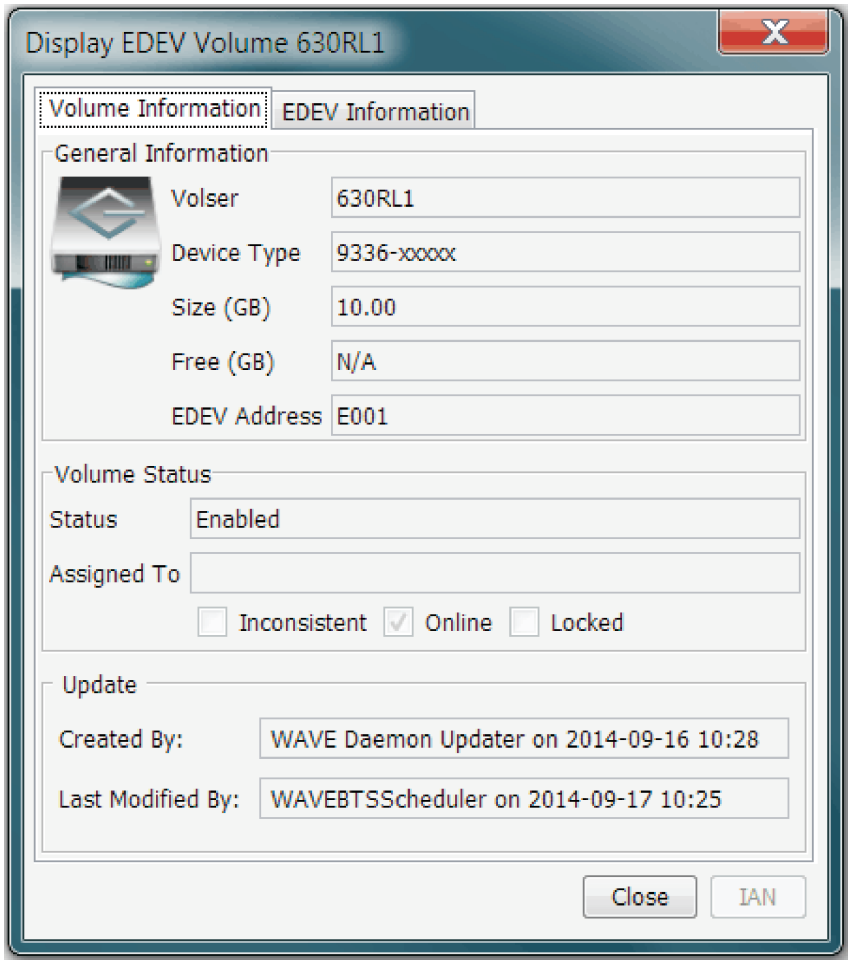

Fields in the Volume Information pane:

- **Volser** - 6-digit alphanumeric serial of the volume
- **Device Type** - The device type of the DASD volume.
- **Size (GB)** - Total size of the FCP volume
- **Free (GB)** - Free space on the FCP volume.
- **EDEV Address** - The address for an EDEV definition that can be automatically used when you are defining an EDEV for SCSI storage on VM systems.
- **Status** - Volume is unassigned or enabled.
- **Assigned To** - Name of z/VM user.
- **Inconsistent/Online/Locked** - Status of the volume.

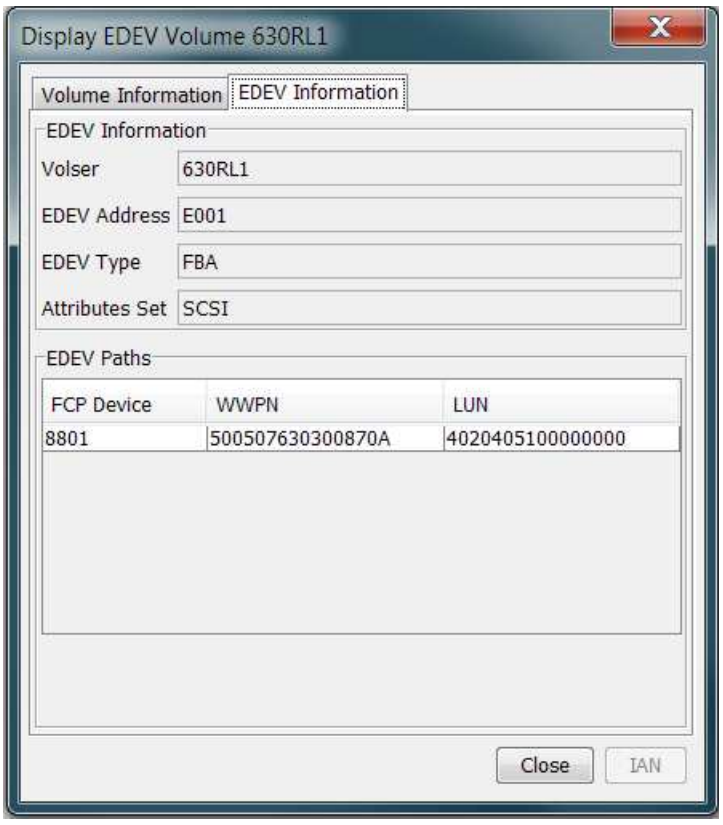

Fields in the **EDEV Information** pane:

- **Volser** - 6-digit alphanumeric serial of the volume.
- **EDEV Address** - The address for an EDEV definition that can be automatically used when you are defining an EDEV for SCSI storage on VM systems.
- **EDEV Type** - Currently only FBA is supported.
- **Attributes Set** - Currently only SCSI is supported.
- **EDEV Paths** - FCP device, WWPN, and LUN information for each path that is specified in the EDEV.
- **FCP Device** - 4-digit hexadecimal number of the FCP device.
- **WWPN** - worldwide port name.
- **LUN** - Logical Unit Number.

Update information for an EDEV

You can update information for existing EDEVs that are managed by IBM Wave.

From the "**Storage**" tab of IBM Wave, right click an EDEV volume and select the menu item to update information. The **Update Information for EDEV Volume** opens and displays the **EDEV Information** tab. The tab displays the relevant EDEV details and the paths associated with the EDEV.

Fields in the **EDEV Information** tab:

- **Volser** - six digit alphanumeric serial of the volume
- **EDEV Address** - The address at which the EDEV is defined.
- **EDEV Type** - Currently only FBA is supported.
- **Attributes Set** - Currently only SCSI is supported.

- **EDEV Paths** - FCP Device, WWPN, sand LUN information for each EDEV path specified.

  FCP Device
  Four digit hexadecimal number of the FCP device.

  WWPN
  World Wide Port Name

  LUN
  Logical Unit Number.

You can add or update paths in the same manner you used to "Create an EDEV" on page 159.

Discover an EDEV

IBM Wave can discover online and offline emulated device (EDEV) definitions that are created or modified within or outside of IBM Wave as part of the **Update Storage Aspect** action.

Direct-attached SCSI Devices

Small Computer System Interface (SCSI) is an ANSI standard electronic interface that allows systems to communicate with peripheral hardware, such as storage devices with faster speed and more flexibly than other interfaces. After you work with your Linux Administrator and z/VM Storage Administrator, IBM Wave can use the serial SCSI command protocol to enable you to add storage for new z/VM guests or duplicate z/VM definitions.

Prerequisites: IBM Wave does not provision new storage. You must work with your Linux Administrator and z/VM Storage Administrator to control devices that are connected through the mainframe-specific channel subsystem for IBM z Systems (direct-attached SCSI devices, SAN network, relevant storage controllers, and others). Your Linux Administrator must receive the WWPN and LUN identifiers and configure IBM Wave to use them. Afterward, you can define storage when you create a new z/VM Guest or define a duplicate z/VM User. For information, see the topic about "Storage Management" in .

The following topics are about adding SCSI storage and MDisks to z/VM guests and Linux users in IBM Wave:

- For more information about Linux multipath support, see "Multipath support for FCP devices" on page 97.
- For the steps for adding minidisks and SCSI devices, see "Adding SCSI Storage."
- To create a new z/VM Guest and define SCSI storage, see "Create New z/VM Guest" on page 72.
- To clone an existing Linux user and define SCSI storage, see "Duplicate z/VM User Definition" on page 85.

Adding SCSI Storage

Use the following procedure to add SCSI storage.

Before you begin

You must work with your Linux Administrator and z/VM Storage Administrator to define storage. After FCP device pools are defined, you can add SCSI storage when you create a new a z/VM Guest or duplicate an existing z/VM user. For more information, see "Display device pools" on page 53 and "Multipath support for FCP devices" on page 97.

Procedure

1. The **Storage Configuration** pane in the **Create New z/VM Guest**, **Duplicate z/VM User Definition** and **Update z/VM User** contains an option to **Add SCSI Storage**. For either action, select **Add SCSI Storage**.

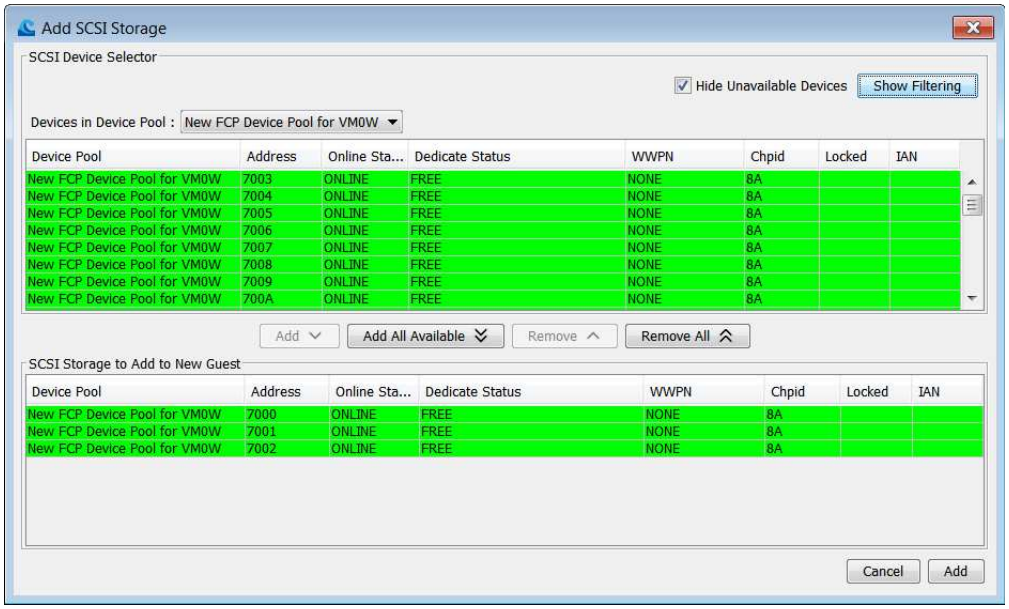

Figure 129. Add SCSI storage

2. To **Add SCSI Storage**, select a device pool to populate the "Devices in Device Pool" table with a list of available devices to add to the guest. By default, the "Hide Unavailable Devices" check box is selected, so only the available devices are shown.
   - Select one or more devices from the "Devices in Device Pool" table and click **Add** to add them to the "SCSI Storage to Add to New Guest" table. Click **Add All Available** to add all of the available devices without preselecting them.
   - To remove items from the "SCSI Storage to Add to New Guest" table, select them and then click **Remove**. Click **Remove All** to remove all of the devices without preselecting them.

   When you add storage, it is possible that another user might also be defining and adding storage simultaneously. To safeguard against a lock, if a device becomes unavailable, IBM Wave notifies you of the conflict and removes any unavailable devices from the table. Before you complete the storage definition, you can return to the **Add SCSI Storage** window to add more available devices.

Results

You added SCSI storage for a new z/VM Guest or for a duplicated z/VM user definition.

"Create New z/VM Guest" on page 72

"Duplicate z/VM User Definition" on page 85

Virtual network segment functions

This section describes actions for Virtual Network Segments.

Create new virtual network segment

This action is accessible from the Virtual Network Segment manager and is used to create a new Virtual Network Segment. For information about the various fields, see “Display VNS information.”

Display VNS information

To display information about a Virtual Network Segment (VNS), from the **Network Viewer**, right-click on the VNS and then click **Display Information**.

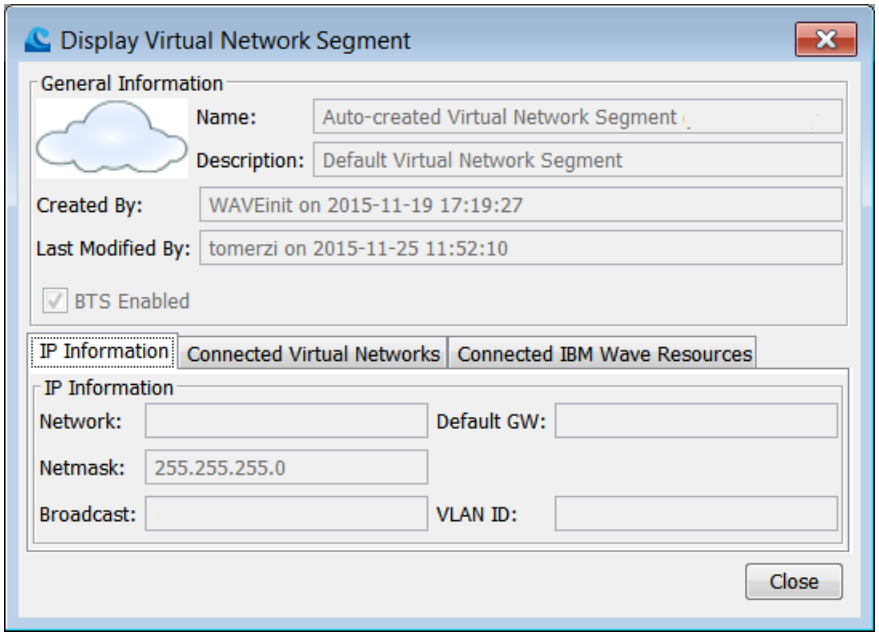

Figure 130. Display Virtual Network Segment

For descriptions of the fields in **Display Virtual Network Segment**, see “Update VNS information.”

Update VNS information

To update information for a Virtual Network Segment (VNS), from the **Network Viewer**, right-click on the VNS, and then click **Update Information**.

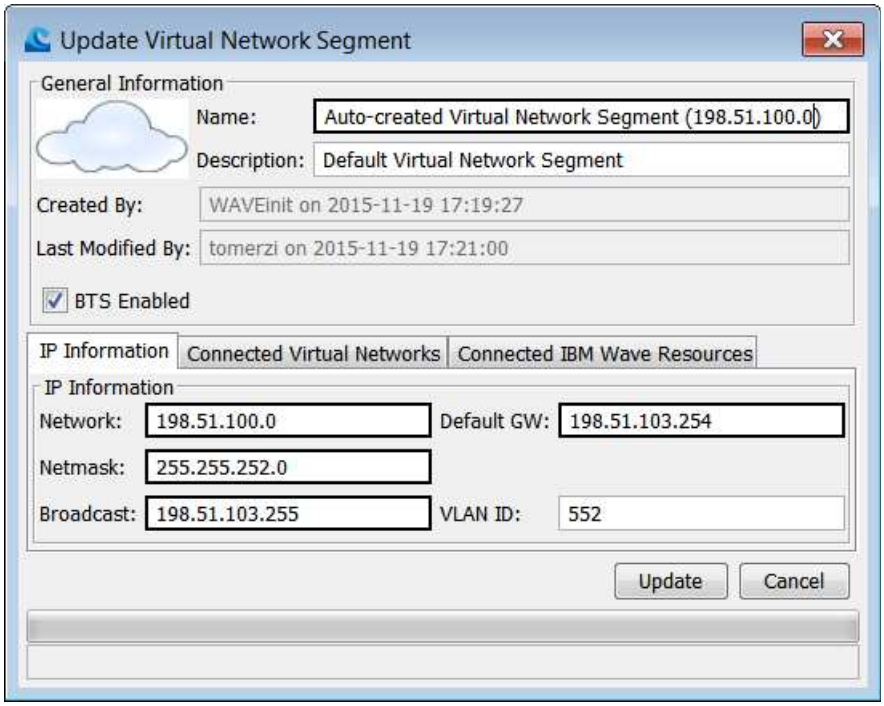

Figure 131. Update Virtual Network Segment information

The following fields are in the **Update Virtual Network Segment**:

- **Name** - The name of the VNS.
- **Description** - The description of the VNS.
- **Created By/Last Modified By** - The user or process that created and last modified the VNS with the time stamp.
- **BTS Enabled** - Indicates whether the BTS can attempt to use connections that are routed through this VNS to connect to guests. When the check box is selected, the BTS attempts to use IP addresses from this VNS to perform managed guest actions such as querying performance data, or managing storage. When the check box is not selected, the BTS ignores any IP addresses that come from this VNS, and cannot use them to connect to guests.

  Note: When the check box is not selected, because the BTS cannot use the VNS to communicate with the guest, it is possible that some guests might appear as if they are not connectable.
- **IP Information**
  - **Network** - The network of the VNS. The network can be either a full IP segment, or a subnet that is combined with the contents of the Netmask field.
  - **Netmask** - The netmask that is defined for this VNS.
  - **Broadcast** - The broadcast address that is used by this VNS.
  - **Default GW** - The default gateway to be used to connect z/VM Guests to the VNS.
  - **VLAN ID** - The optional indicator for whether the VNS is connected to a VLAN.
- **Connected Virtual Networks** - This tab contains a table with a list of the Virtual Networks that are connected to the VNS. You can click **Connect to Virtual Network** to add another network.

- **Connected IBM Wave Resources** - This tab contains a table of the IBM Wave Managed entities that are associated with the VNS. You can click **Connect an External Entity** to connect the VNS to a router.

Remove

Use this action to remove a Virtual Network Segment (VNS). A VNS that has one or more z/VM Guests connected to it cannot be removed.

External entities functions

To access the **IBM Wave External Entities Manager**, from the **IBM Wave Main Menu**, click **Administrative** > **Site Management** > **Manage External Entities**.

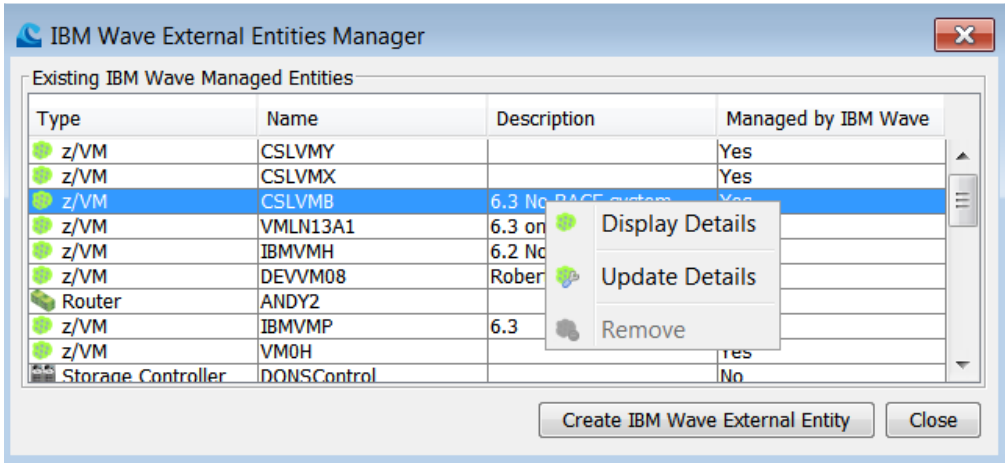

Figure 132. External Entities Manager

The fields in this window are:

- **Type** - The type of the Managed Entity such as a z/VM system, router, or storage controller.
- **Name** - The 4 - 8 character name of the Managed Entity.
- **Description** - The optional description for the Managed Entity.
- **Managed by Wave** - The table displays whether the entity is managed by IBM Wave.

Create a new external entity

Before a z/VM system or other entity can be defined and managed by IBM Wave, it must be added to the **External Entities Manager**.

Before you add a new z/VM system, router, or controller, you must define it as an **External Entity** that IBM Wave can manage.

From the **IBM Wave Main Menu**, click **Administrative** > **Site Management** > **Manage External Entities**, and then click **Create IBM Wave External Entity**.

Select the type of managed entity from the menu, and fill in the name and optional description of the entity.

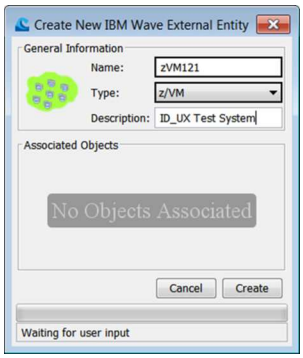

Figure 133. Create a New External Entity

The fields in the **Create IBM Wave External Entity** pane are:

- **Name** - The name of the Managed Entity. The name must be 4 - 8 characters, or it is not accepted.
- **Type** - The type of the Managed Entity.
- **Description** - The optional description for the Managed Entity.
- **Associated Objects** - The table displays a list of the objects associated with the managed entity. The associated objects can be either Device Pools or Virtual Network Segments, but the objects do not appear until you establish the connection.

Display, update, or remove an external entity

To display, update, or remove an external entity, right click on the entity as shown in Figure 132 on page 171. Select one of the following options:

- **Display Details** to view the entity and any associated objects.
- **Update Details** to update the entity.
- **Remove** to remove the entity.

For more information, see .

Related information:

"Add New System" on page 60

Device pool functions

The following actions relate to Device Pools:

- "Create New Device Pool"
- "Display device pool information"
- "Update device pool information" on page 174
- "Read, update, or remove an IAN" on page 174
- "Remove a device pool" on page 174

Create New Device Pool

Access the **Create New Device Pool** action from the **Device Pool Manager**. See "Display device pool information" for a description of the fields.

Display device pool information

Click **Display** to display information about an existing Device Pool.

Figure 134. Display device pool information

The fields in the **Display** are:

General Information

- **Device Pool Name** - The name of the Device Pool.
- **Type** - The type of the Device Pool.
- **Created By** - The IBM Wave user who created the Device Pool and the timestamp.
- **Modified By** - The IBM Wave user who last modified the Device Pool and the timestamp.

Descriptive Information

- **Default Virtual Device** - An optional device address (4 hexadecimal digits), which specifies the default virtual device to use when dedicating real devices from the device pool to z/VM guests.
- **Description** - An optional description for the Device Pool.

The **Devices** tab displays all the devices owned by the Device Pool.

- **System Name** - The **System Name** table displays all z/VM Systems associated with the Device Pool. When a z/VM System is selected, the real devices table is populated with the real devices from the selected z/VM System owned by the Device Pool.
- **Real Devices** - The **Real Devices** table lists the real devices system-addresses that belong to the selected z/VM System and owned by the Device Pool. Each line in the table provides information about one real device's address, online status, and dedicated status.

The **Connected IBM Wave Resources** tab displays a table that lists all the IBM Wave Managed Entities that are connected to the Device Pool.

Update device pool information

Use this action to update an existing Device Pool. For a description of the various fields, see "Display device pool information" on page 173. When moving to the Devices tab, it is possible to associate a z/VM System with the Device Pool. When moving to the Connected IBM Wave resources tab, it is possible to connected an IBM Wave Managed Entity to this Device Pool.

Read, update, or remove an IAN

Use this action to read, update, or remove an IAN attached to a Device Pool.

Remove a device pool

Click **Remove** to remove a Device Pool.

Note: Device Pools that own real devices cannot be removed.

Chapter 3. IBM Wave reporting subsystem

IBM Wave contains a powerful reporting subsystem. You can use the IBM Wave Reporting subsystem to generate reports about the z/VM and IBM z Systems environment. You can also customize, print, and export reports to comma-separated value (.csv) files. In addition, you can generate specific reports for the custom attributes that are defined in your environment.

- For an overview of the IBM Wave reporting subsystem, see "Manage report templates" on page 176.
- To work with data created in an IBM Wave V1R1 report, see "Report behavior: V1R1 and V1R2 reports" on page 182.
- To create a new report, see "Create a new report" on page 177.
- To customize a report with the data you require in the report, see "Customize a report" on page 183.
- To load an existing report, see "Load or update a report" on page 185.
- To save a report, see "Save a report" on page 187.
- To understand scopes, permissions, and locking reports, see "Report scopes and permissions" on page 189.
- To lock or unlock a report, see "Lock and unlock a report" on page 190.
- To export data to a spreadsheet program, see "Exporting a report" on page 191.
- To print a report, see "Print a report" on page 192.
- To delete a report, see "Delete a report" on page 193.
- For a list of all the IBM Wave report descriptions, see "IBM Wave report descriptions" on page 195.

© Copyright IBM Corp. 2007, 2018

Manage report templates

To access the IBM Wave report subsystem, from the main menu in IBM Wave, select **Reports > Manage Report Templates**.

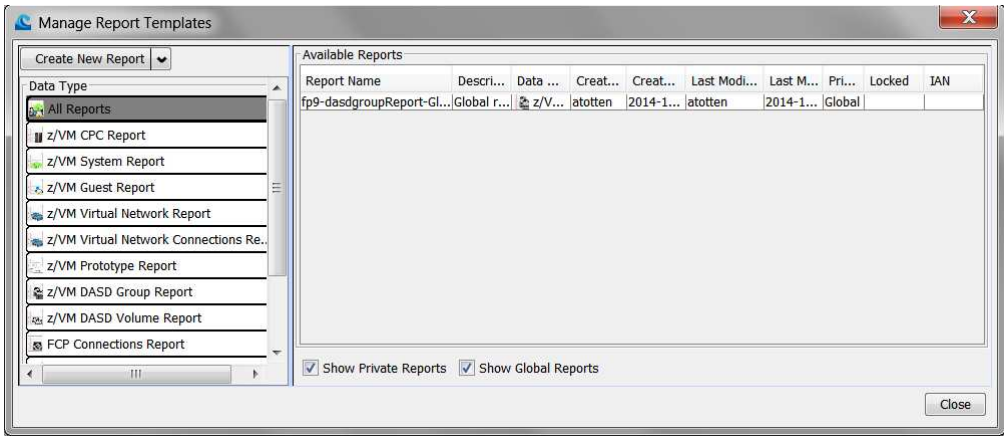

Figure 135. IBM Wave Manage Report Templates

The **IBM Wave Manage Report Templates** window opens. The list of report types are displayed in the **Data Type** pane on the left side of the window. Saved global and private reports are listed in the **Available Reports** pane on the right side of the window. Each line in the **Available Reports** pane represents a single report. For details about roles and permissions, see "Report scopes and permissions" on page 189.

To access the report template:

- Click **Create new report**, and then select the report type from the menu.
- Right-click on the report type, under **Data Type**.

When you move the mouse over a line that contains a report in the **Available Reports** pane, a summary appears displaying the following information:

- The name of the report as provided by the IBM Wave user who created the report.
- A description of the report (when an IBM Wave user who created the report enters a description).
- A reminder to double-click the report to open it.
- If an IAN is attached, the name of the creator, the date and time the IAN was created, and the message that is in the IAN."
- To understand how to work with reports, see "Create a new report" on page 177.
- For a complete list of reports, see "IBM Wave report descriptions" on page 195.

Create a new report

Use the steps in the following procedure to create an IBM Wave report.

About this task

Use the following steps to create a new IBM Wave report:

Procedure

1. To access the report manager, select **Reports > Manage Report Templates** from the main menu in IBM Wave.
2. Press **Create New Report** and then select the report from the list of available report types. The **IBM Wave Report Generator <*Report Name*>** opens.
3. In the **Report Parameters** section, adjust the filter parameters to contain the data you want appears in the report. For the report parameters, you can set all "AND" or all "OR" operands. You cannot use mixed operands, "AND" and "OR" values.
4. In the **Report Fields** section, press "**Add**" to move a field in a hidden column to the visible columns. Press "**Remove**" to move a field from the visible columns to the hidden columns.
   - The following options are available:
     - Select one or more contiguous fields from the hidden columns and press "**Add**" to make the selected columns visible. Select one or more contiguous fields from the visible columns and press "**Remove**" to hide the selected columns.
     - Press "**Add All**" to make all of the hidden columns visible. Press "**Remove All**" to remove all of the visible columns.
     - To reorder one or more visible columns in the report, select the column to move and then use the blue up or the down arrow, which is located on the right of the visible columns. To move a field up, press the up arrow. To move a field down, press the down arrow.

   When you are satisfied with the parameters in the report, you are ready to generate the report.
5. Press "**Generate Report**" to run the report. All reports are generated with a default scope of private.

Results

The report results are displayed in **Report Results**.

Example

Figure 136 on page 178 is an example of the IBM Wave Report Generator - Custom Attribute Report after the report was generated. You can sort the data in **Report Information** pane to display the data in ascending or descending order, and move the columns in any order you prefer.

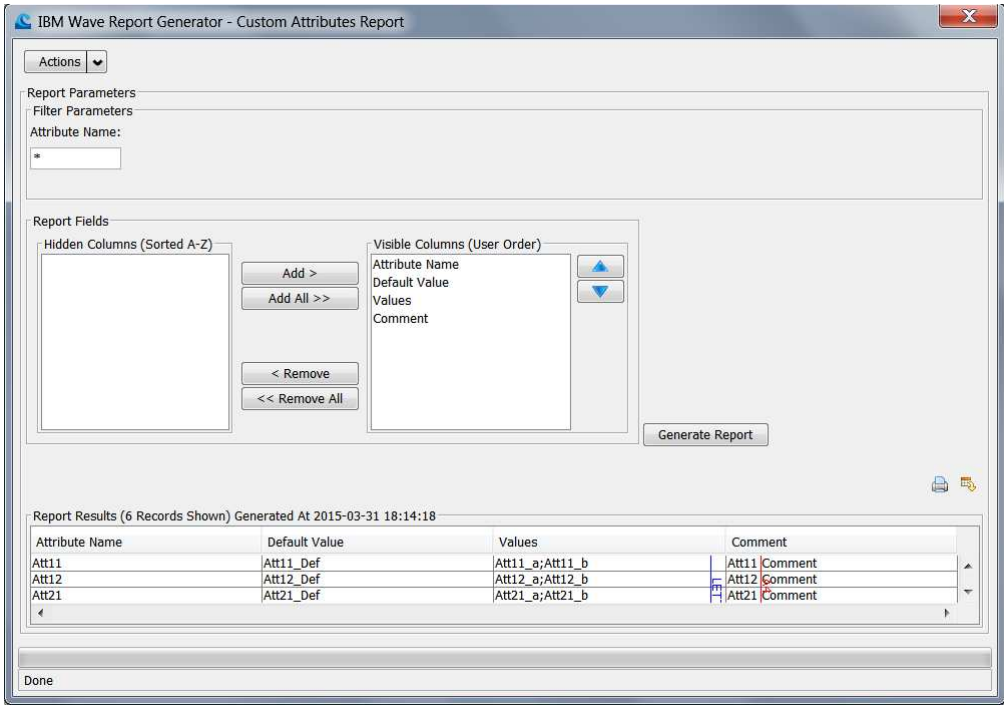

Figure 136. IBM Wave Custom Attribute Report

What to do next

When you are satisfied with the generated data, you can:

- Customize. For more details, see "Customize a report" on page 183.
- Save as. For information, see "Save a report" on page 187.
- Export. For information, see "Exporting a report" on page 191.
- Print. For information, see "Print a report" on page 192.

IBM Wave Report Generator

Using the **IBM Wave Report Generator** you can create concise, detailed reports. Reports can be filtered, and then printed or exported for further analysis.

The **IBM Wave Report Generator** window contains the following panes:

- **Report Parameters**: The first pane in IBM Wave Report Generator contains the parameters for the report. The **Report Parameters** provides you with a method of entering complex filters for the report. Each report contains a unique set of filters that match the objects in the report. Only all “AND” or only all “OR” operands can be set. Mixed “AND” and “OR” operands are not supported. The pane also contains any additional filters that apply to the report. For example, in the *z/VM System Report*, the *Active* and the *Suspended* filters to display the state of the system.
- **Report Fields**: The second pane contains the report fields. The following options are available:
  - Select one or more contiguous fields from the hidden columns, and then click "**Add**" to make the selected columns visible.
  - Select one or more contiguous fields from the visible columns, and then click "**Remove**" to hide the selected columns.
  - Click "**Add All**" to make all of the hidden columns visible.
  - Click "**Remove All**" to remove all of the visible columns.
  - To reorder one or more fields in the report, use the blue up or down arrow. Click the up arrow to move up the field. Press the down arrow to move a field down.
  - Click "**Generate**" to run the report. "**Generate**" is on the right of **Report Fields**.
- **Report Results**: The last pane contains the report results. The fields in the **Report Results** are not visible until you click "**Generate**" to create the report.

  After the report runs, the number of records that are in the report and a time stamp when the report was generated appears above **Report Results**. The pane also contains a **Print** icon and a **Save** icon, which are at the upper right of the **Report Results** pane.

Figure 137 on page 180 shows an example of the *IBM Wave Report Generator: IBM z/VM System Report* before “**Generate**" is clicked.

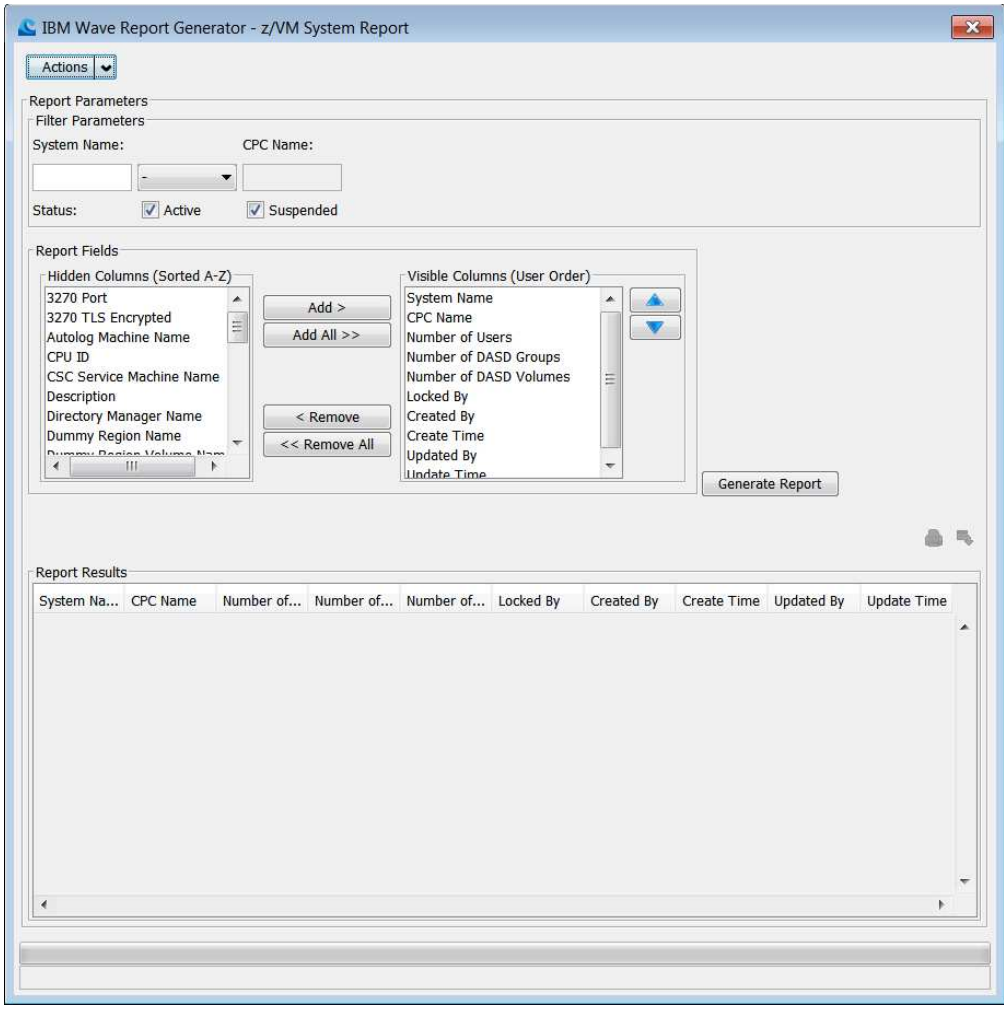

Figure 137. The IBM Wave Report Generator: z/VM System Report

Figure 138 on page 181 shows an example of a generated report: *IBM Wave Report Generator: IBM z/VM System Report*.

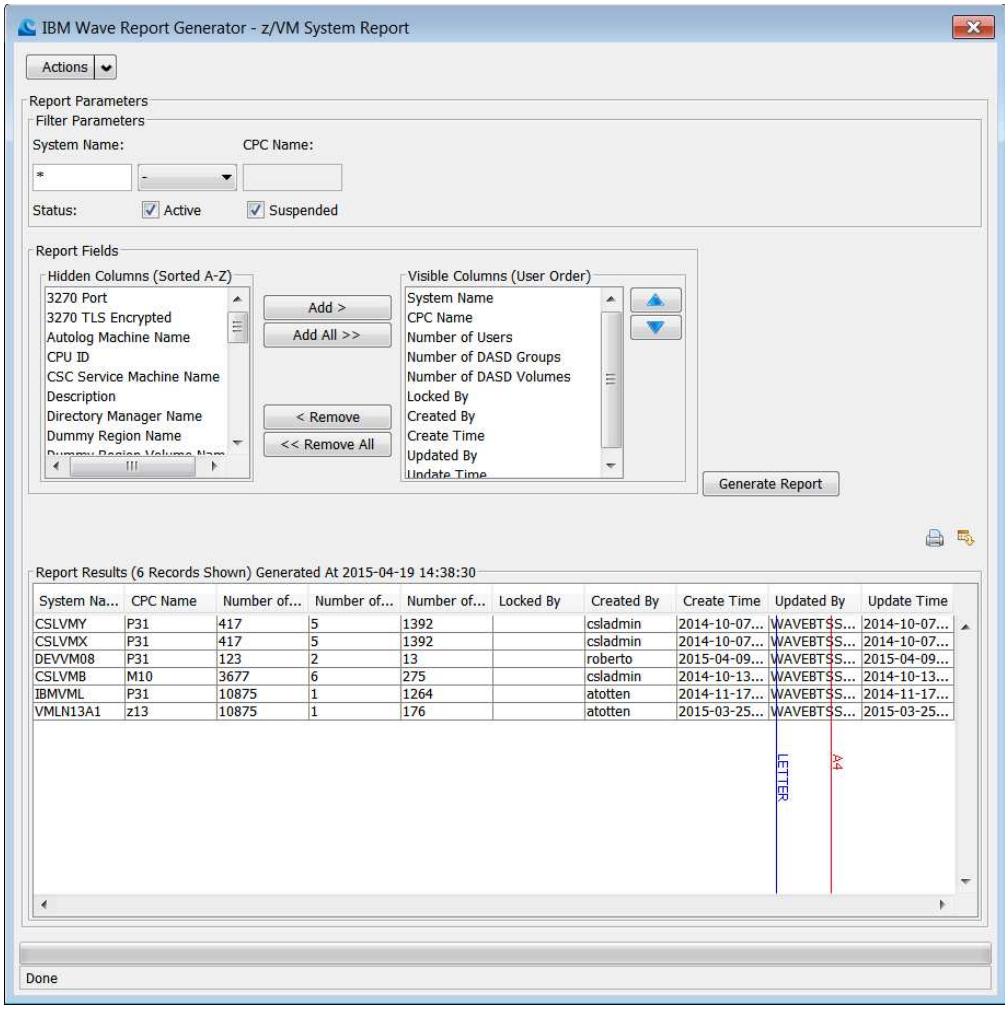

Figure 138. Example of a generated report

- For information about scopes and permissions, see "Report scopes and permissions" on page 189.
- For instructions about generating a new report, see "Create a new report" on page 177.
- For information about how to customize the report after the table is populated, see "Customize a report" on page 183.

Report behavior: V1R1 and V1R2 reports

The following information explains how to work with V1R1 report data by using IBM Wave V1R2.

In IBM Wave V1R1, you were able to create a report that uses filters with mixed "AND" and "OR" operands. IBM Wave V1R2 can use either all "AND" or all "OR" operands. The filters do not support mixed "AND" and "OR" operands. The IBM Wave V1R1 reports with mixed filters are still usable, but you are unable to update or change the report template. You can generate report data, print, or export a report, but you cannot save the V1R1 report.

When you view V1R1 reports, IBM Wave retains all saved attributes, which includes the size of the report window. When you open an IBM Wave V1R1 report with IBM Wave V1R2, if any section of the report is cut off or concealed, resize the window to the appropriate size to see all objects.

Report behavior: removing a custom attribute from IBM Wave V1R2

All custom attributes must be assigned by using the **Custom Attribute Manager**. If you remove a custom attribute from the Custom Attribute Manager, it will affect any Custom Attribute reports that include the removed field.

If a custom attribute is removed and either of the following scenarios apply:

- While in the middle of updating a report template that contains the removed custom attribute.
- If you remove a custom attribute from the **Custom Attribute Manager**, it affects any custom attribute reports that include the removed field.

IBM Wave displays a warning message, at the top of the loaded report template, that indicates the custom attribute was removed. The warning message indicates the names of the removed attributes to inform you that a data point was removed. After the report is saved, without the removed attribute, the warning message goes away. Alternatively, you can close the warning message.

Customize a report

You can customize the IBM Wave reports with the fields for objects that you want to display. Each report contains filters that match the objects that are available for each report. You can add and remove filters to make the report contain only the objects you want in the report.

Fields in the **IBM Wave Report Generator**:

Report Parameters

Each report has a unique set of filter parameters that you can tailor in the **Report Parameters** pane. Mixed filters, the "AND" and "OR" operands, are not supported. Only all "AND" or only all "OR" operands can be set. All filters that have an asterisk show up in the report.

Report Fields

Each report has fields that can be added and removed from the report.

- To add or remove one or more fields from the *Hidden Columns (Sorted A-Z)* or the *Visible Columns (Use Order)* . Select one or more contiguous fields, and then press "**Add**" or "**Remove**" to add or remove.
- To add or remove the entire list of fields from the *Hidden Columns (Sorted A-Z)* or the *Visible Columns (Use Order)* . Press "**Add All**" or "**Remove All** " to add or remove all of the fields.
- To change the use order, press the blue up or down arrow to move the value to the location you want. You can also press "**Add All**" or "**Remove All**" to move all the report fields.

Tip: You can also press the control (Ctrl) key while selecting one or more values, and then press "**Add All**" or "**Remove All**" to move fields to the desired column.

Report Results

Each value that is entered in the **Report Parameters** pane, and the "Visible Columns (Use Order)" in the **Report Fields** pane appears in the **Report Results**. After the report is generated, you can continue to modify the column size or sort the report based on your own preferences.

- Change the order of the columns by dragging the header of the column to its new location.
- Change the width of the columns by sliding the separator.

After a report is generated, you can print it, or export it to manipulate the data in a spreadsheet or other program.

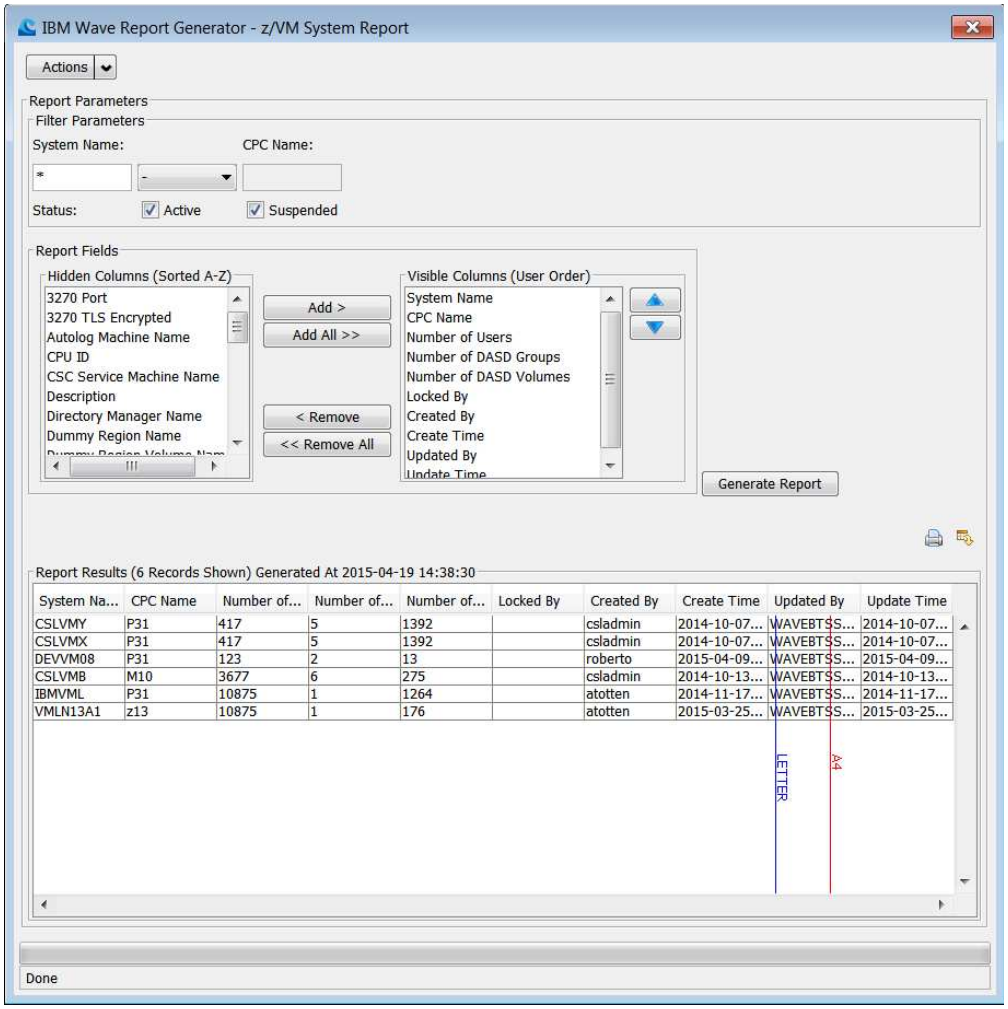

Figure 139. The z/VM System Report

- For information about export and print, see "Exporting a report" on page 191 and "Print a report" on page 192.
- For a complete list of all the reports that IBM Wave can generate, see "IBM Wave report descriptions" on page 195.

Load or update a report

Use the steps in this procedure to load or update a saved report.

About this task

Use this task to load a saved report. The fields in the report are set according to the saved information unless you change the values before you generate the report. For information about reports that were generated by using IBM Wave V1R1 and earlier, see "Report behavior: V1R1 and V1R2 reports" on page 182.

Procedure

1. Select **Reports > Manage Report Templates**. The **Manage Report Templates** window opens.
2. In the **Available Reports** pane, on the right side of the menu, take one of the following actions:
   a. Double-click the line that represents the report you want to load.
   b. Right-click and select **Update Report** to load the report.

   The **IBM Wave Report Generator** opens.
3. To generate the report, press "**Generate Report**".

Results

IBM Wave runs the report, and then loads the data in the **Report Results**.

Example

Figure 140 on page 186 shows an example of the generated report. Notice how the **Report Results** pane displays the number of records in the generated report immediately followed by the time stamp of when the report was generated.

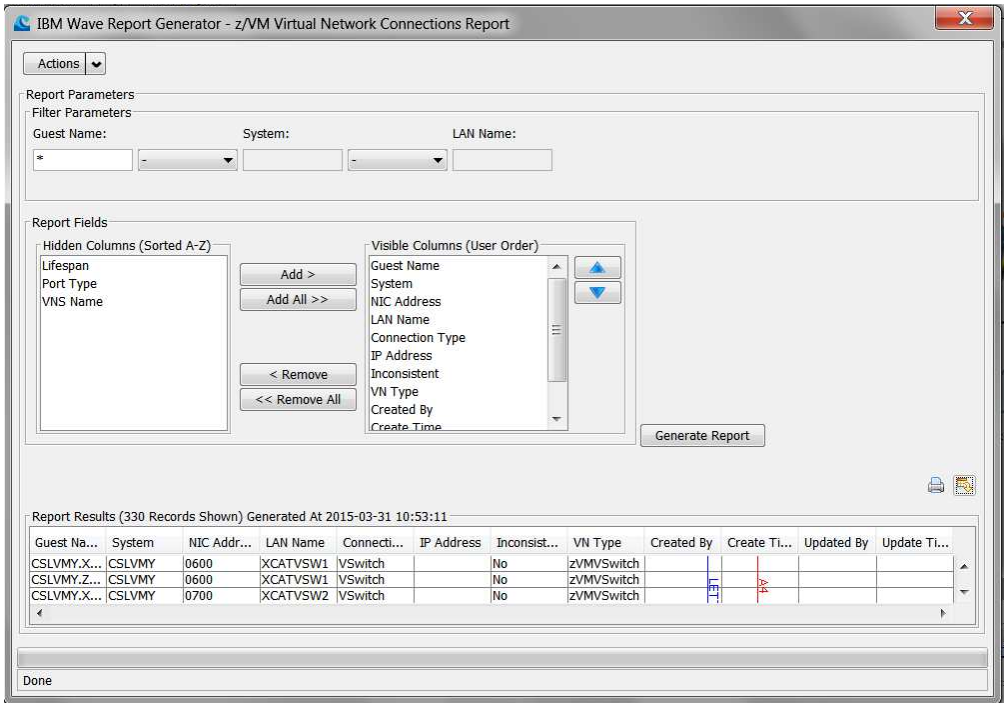

Figure 140. Report example: Virtual Network Connections Report

Save a report

With the IBM Wave Report Generator, you can export and save a generated report for future use. This is especially useful for reports that incorporate complex filters and customization.

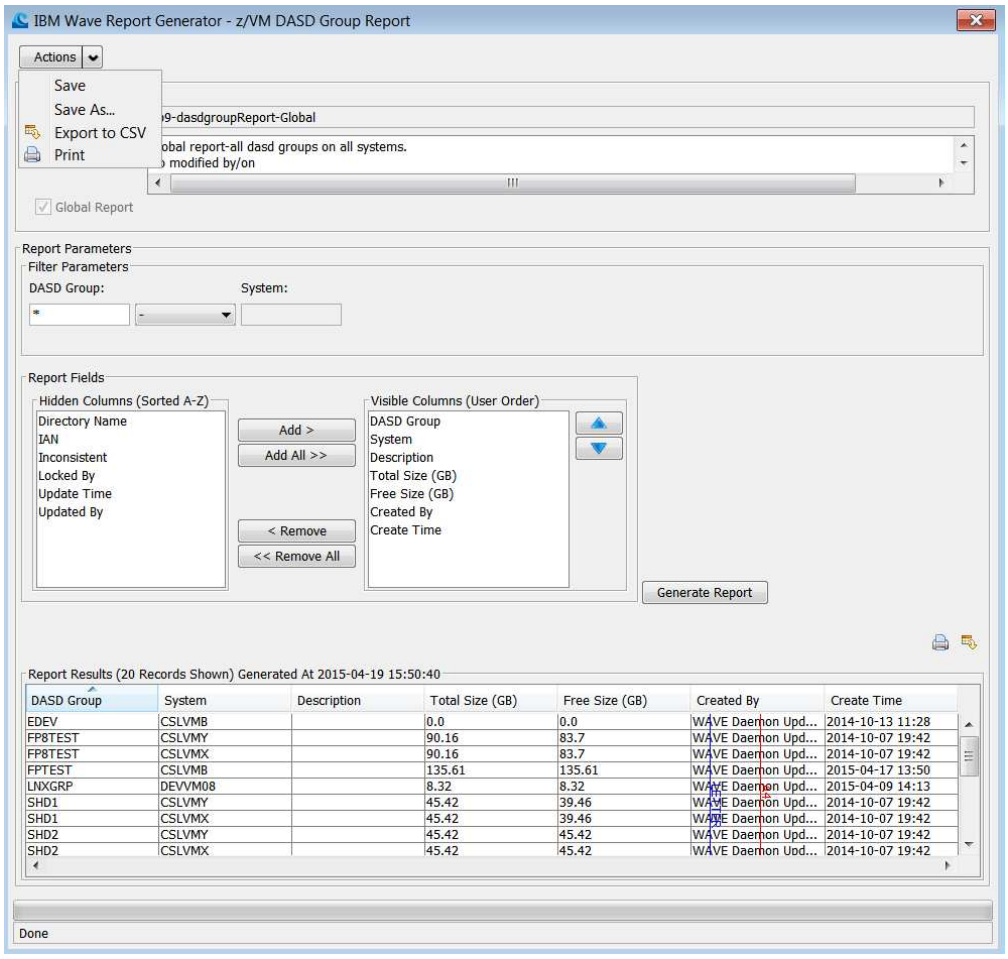

Figure 141. Save an IBM Wave report

After the report is generated, saving can be done in two ways:

1. "**Save As...**": Saves a newly created report, and specifies a name and description for the new report.
2. "**Save**": Only available when you update an already saved report. You can load a saved report, make changes, and then save the changes.

Some important concepts about saving reports:

- Each report has a creator, which corresponds to the initial IBM Wave user who created the report.
- Only IBM Wave users with the site level administration permissions can create and save global reports. The same is true for converting global reports to private reports. Only the site level administrator can convert a global report to a private report. To change the report scope to a global report, select the *Global Report* check box.
- When converting a report from global to private, the report becomes a private report of the IBM Wave user that first created the report (called the report

creator). This is an important concept because the IBM Wave user who converts the report might not be the IBM Wave user who created the report.

Report scopes and permissions

There are two report scopes:

Global reports
Global reports can be viewed, accessed, and generated by an IBM Wave administrator with site level administrator (SLA) permissions. Only an administrator with SLA permissions can save changes to the report template.

Private reports
Private reports can be viewed, accessed, and generated by the user who originally created the report. Only the report creator can view and save changes to the report template.

The data in a generated report contains objects that are in the IBM Wave user's scope. A user cannot generate data in a report for objects that are not in their scope. For example, a user with networking administrator permissions cannot generate the z/VM DASD Group report or the z/VM DASD Volume report. In contrast, an IBM Wave administrator, who has SLA permissions, can view all reports.

Multiple users who have SLA permission can open and use report templates within their scopes and permissions concurrently. Any updates that are made by one user on a saved report template are automatically refreshed and displayed to all other users that are viewing the same report template. To disable this feature, and ensure that a global report is not modifiable by any other users, you can lock the report. To lock and unlock a report, see “Lock and unlock a report” on page 190.

If you used IBM Wave V1R1 to create reports, see “Report behavior: V1R1 and V1R2 reports” on page 182.

Lock and unlock a report

How to lock and unlock an IBM Wave report.

Before you begin

In the IBM Wave **Manage Report Templates** window, select the All Reports value in the **Data Type** pane to see a list of reports. The reports appear in the **Available Reports** pane. In Figure 142, notice how the example report displays a padlock icon in the “Locked" column to indicate the report is locked. The padlock indicates the report is locked.

When a global report is locked, the report creator is the only user who can modify the report.

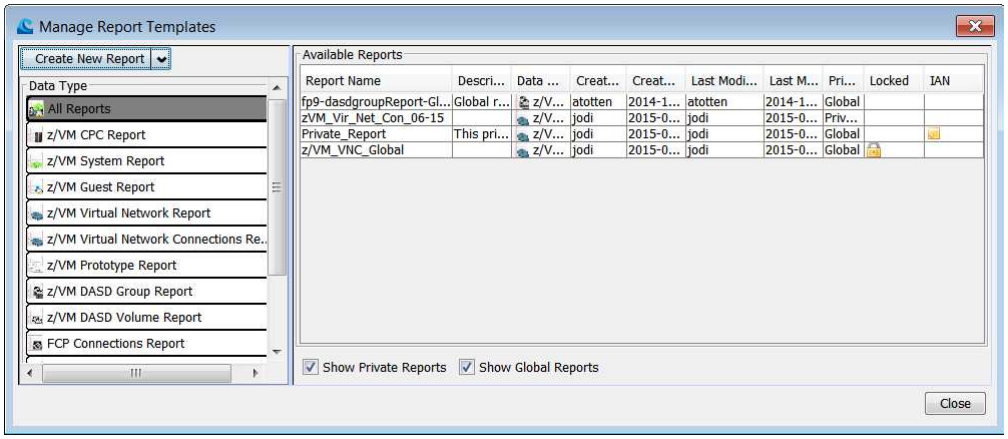

Figure 142. Lock and unlock report templates pane with locked and unlocked reports

About this task

An IBM Wave user with SLA permission can lock or unlock a global report. An administrator might want to lock a global report to keep it from being modified by other users.

Use the following steps to lock or unlock a report:

Procedure

1. Select the report name from the **Available Reports** pane, which is on the right side of the **Manage Report Templates** window.
2. Right-click on the report.
   - Select “**Lock**” to lock the report.
   - Select “**Unlock**” to unlock the report.

   If a report has an intelligent active note (IAN) attached to it, IBM Wave issues a *IAN items-Lock Report* message. To lock the report, you must check the "Run action on this item" box.

Results

When you are done, the report is locked or unlocked.

Exporting a report

Use this task to export any of the IBM Wave reports to a .csv file.

About this task

Use this task to export an IBM Wave report in tabular format. In the **IBM Wave Report Generator**, you can export a generated report by using the export icon or the menu options **Actions > Export to CSV**. Tabular data is typically exported directly to spreadsheets and databases. For a sample of exported data, see Appendix B, "Sample CSV file for importing guest attributes," on page 247.

After an IBM Wave report is generated, the report can be exported in the comma-separated value (.csv)[2] format. Afterward, you can open the report data in a spreadsheet or other program that presents data in a tabular format, and work with the data.

Use the following steps to export an IBM Wave report with the `.csv` file extension, and save it to your local file system.

Procedure

1. After a report is generated, select **Actions > Export to CSV** from the top of the IBM Wave Report Generator window. The **Windows Save** menu opens.
2. Enter a file name that has the `.csv` file extension, and save the file to your local file system.

Example

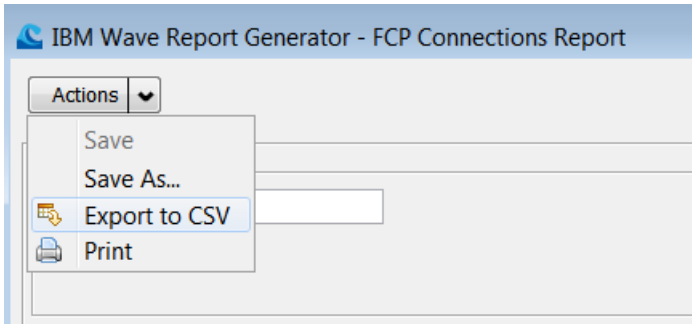

Figure 143. Exporting IBM Wave report data

What to do next

You can now open the .csv file from your local file system and manipulate the tabular data with your own software.

Notes: The following notes apply to working with `.csv` data:

- The report is exported from IBM Wave in standard `.csv` format. IBM Wave follows all `.csv` standards when it exports data to a `.csv` file.
- For details about working with the data in reports, see the instructions from the independent software vendor who created the software.

2. Character-separated value file extension.

Print a report

You can print a new or saved report from the Manage Report Templates window.

Before you begin

In the **IBM Wave Report Generator**, you can print a generated report by using the printer icon or the menu options **Actions > Print**. The print preview window gives you the ability to customize the report before it is printed. You can customize page definitions along with the column order and width. Up to ten lines of text are displayed for each row of the printed report.

About this task

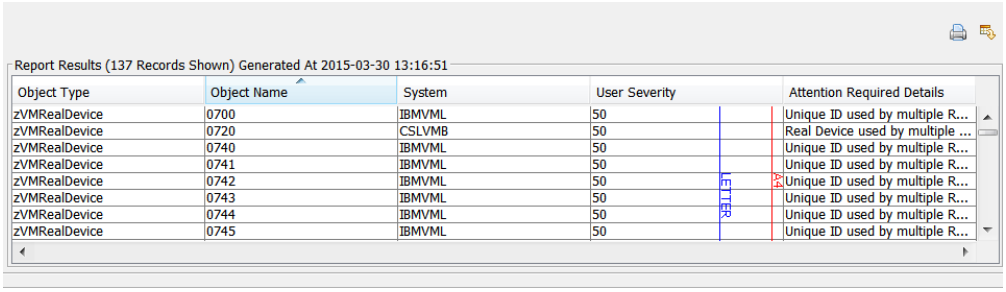

Figure 144. IBM Wave LETTER and A4 print options

IBM Wave has two preset printer values in the **Report Results** pane of the **Manage Report Templates** window. The values can help you arrange the report columns to make the data suitable for printing.

- The blue line with the label "LETTER" represents a page size that is 8.5 inch (216 mm) × 11 inch (279 mm).
- The red line with the label "A4" represents a page size that is 8.27 inch (210 mm) × 11.69 inch (297 mm).

On the print preview menu, in the lower right corner, IBM Wave indicates the number of pages in the report.

Procedure

1. Select **Reports > Manage Report Templates** by from the menu bar at the top of the IBM Wave window. The IBM Wave **Manage Report Templates** window opens.
   - To generate a new report, see “Create a new report” on page 177.
   - To load a saved report, see “Load or update a report” on page 185.
2. Preview the report before you print it. The ***Report Name*** **Print Preview** window opens a preview as shown in Figure 145 on page 193. The preview displays how all of the pages in the printed report appear when the report is printed. The other options in the ***Report Name*** **Print Preview** window include:
   - **Report > Page Setup** to adjust the paper, orientation, and margins.
   - **Report > Go to Page** to enter the page number, or use the slider to review any page in the report.
   - **View > Zoom** to adjust the size of the preview area.
   - **View > User Defined Zoom** to control the size of the preview area by using the slider.

Note: Currently, the print behavior for a saved report is as follows: When report data goes beyond the A4 format border, any fields that are outside of the border are printed on a new row that is below the existing row.

3. When you are satisfied with the way the report looks, select one of following options to print or close the report without printing.
   - In the **IBM Wave Report Generator** window in the upper-left corner of the window, select **Actions > Print**.
   - Press the printer icon, which is above the **Report Results** pane on the right.
   - To exit without printing, select **Report > Close**.

Example

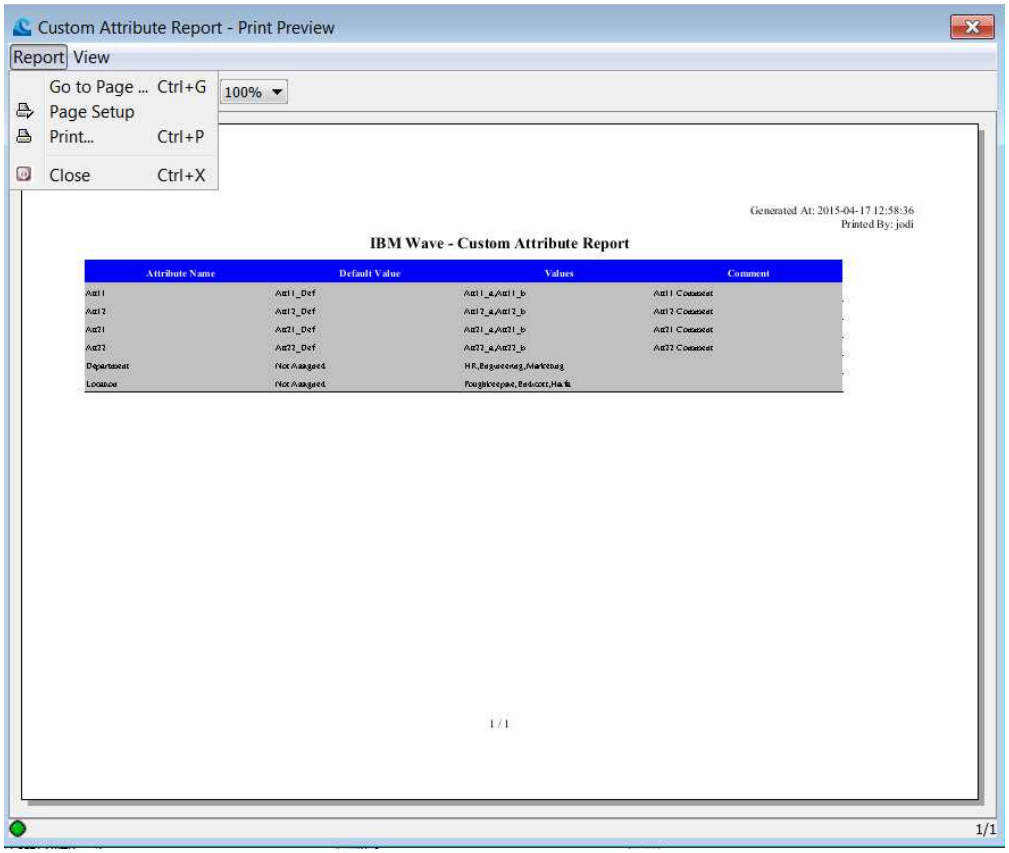

Figure 145. IBM Wave print preview window

Delete a report

An IBM Wave user with the correct permissions can delete a report that is listed in the **Available Reports** pane of the **Manage Report Templates** window.

A user who has the site level administrator (SLA) role can delete a global report, but only the creator of a private report can delete a private report.

To see the scope of a report, see the values in the Privacy column. You can filter the view by selecting the **Show Private Reports** check box or **Show Global Reports** check box. (By default both check boxes are selected.)

To delete a saved report, right-click on the line that represents the report and select "**Delete Report**" from the list of menu options. After you select "**Delete Report**",

the workunit is submitted to the BTS, and then the report is permanently removed from IBM Wave.

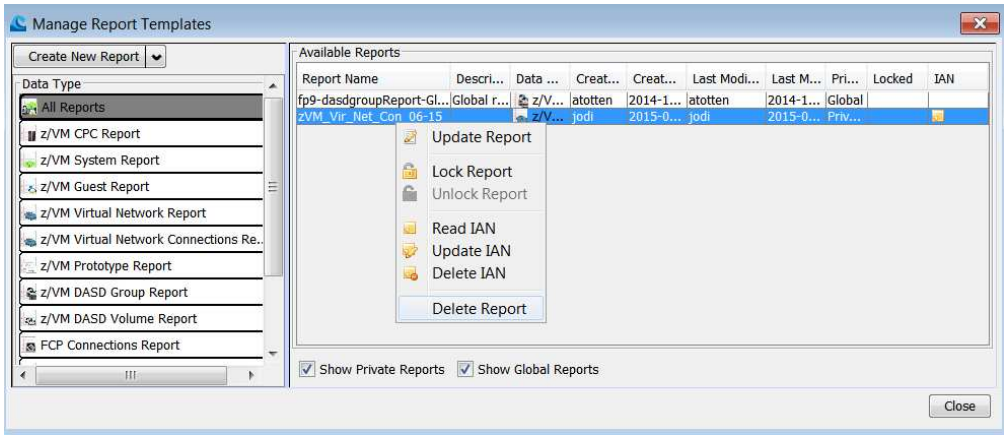

Figure 146. Manage Report Template - Delete a report

When you delete a report that has an intelligent active note (IAN) attached to it, IBM Wave displays a yellow warning message with the following text:

"The selected object has an IAN item attached".

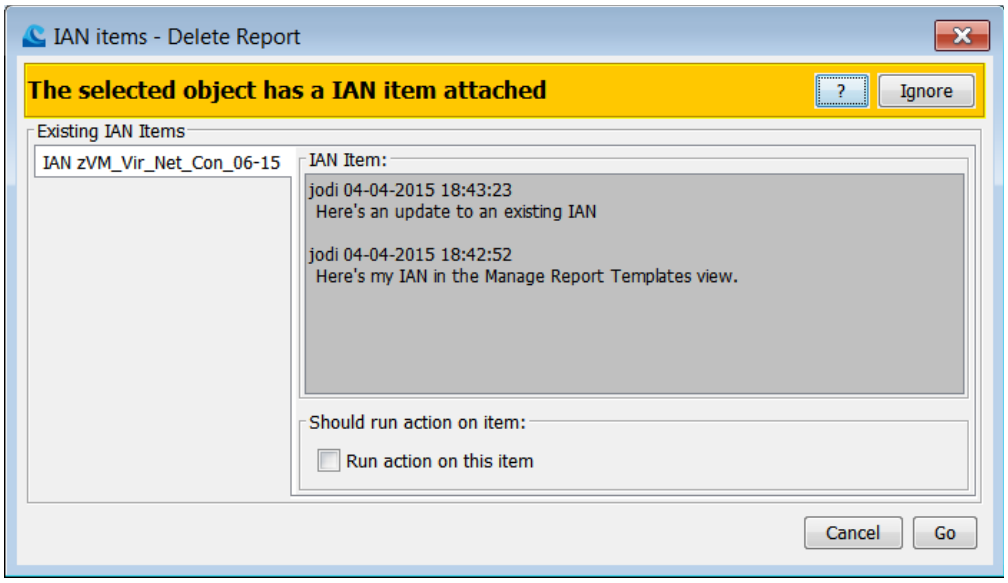

Figure 147. IBM Wave warning message: IAN items - Delete Report

Select the question mark ("**?**") for more information.

Select "**Ignore**" to close the message.

The name of the report is listed in the Existing IAN Items pane. The complete contents of the IAN is listed in the IAN Item pane, which is on the right side of the window.

IBM Wave report descriptions

IBM Wave contains the following reports:

- All Reports: Contain a list of all the saved global and private reports.
- z Systems CPC report: “CPC Report” on page 196
- z/VM System report “z/VM System Report” on page 198
- z/VM Guest report “z/VM Guest Report” on page 201
- z/VM Virtual Network report “z/VM Virtual Network Report” on page 205
- z/VM Virtual Network Connections report “z/VM Virtual Network Connections Report” on page 207
- z/VM Prototype report “z/VM Prototype Report” on page 209
- z/VM DASD Group report “z/VM DASD Group Report” on page 211
- z/VM DASD volume report “z/VM DASD Volume Report” on page 213
- FCP Connections report “FCP Connections Report” on page 215
- Flagged z/VM Object report “Flagged z/VM Object Report” on page 217
- IBM Wave Custom Attribute report: “IBM Wave Custom Attribute Report” on page 219
- IBM Wave User report “IBM Wave User Report” on page 221
- IBM Wave User Permission report “IBM Wave User Permission report” on page 223

CPC Report

CPC report for each of the systems managed by IBM Wave.

The CPC Report is used to generate a report about each IBM Z central processor complex (CPC) that is being managed by IBM Wave.

The **Report Parameters** pane contains the following **Filter Parameters** that you can customize for your environment and add to your report:

CPC Name
The name of the CPC that you want to use as a filter.

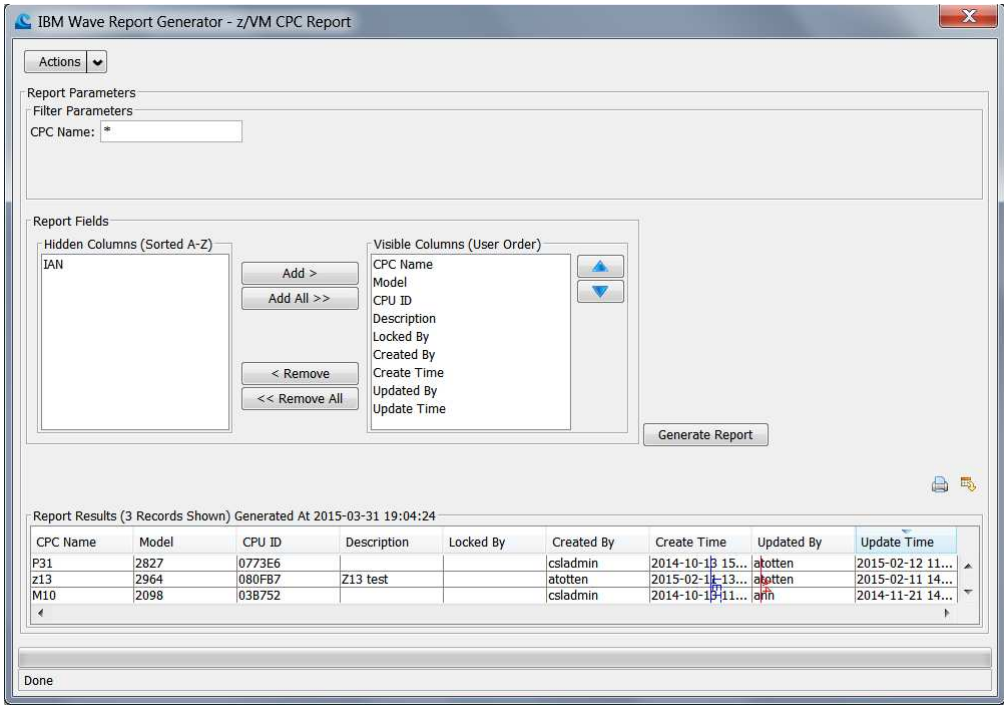

Figure 148. IBM Wave Report Generator - CPC report

Table 13. CPC report fields

The Report Fields in the *Hidden Columns (Sorted A - Z)***:**	**The Report Fields in the** *Visible Columns (Use Order)***:**
IAN The name of the user who created the IAN.	**CPC Name** The CPC Name. **Model** The CPC IBM model number. **CPU ID** The CPU ID. **Description** A description of the CPC. **Locked By** Describes who holds the lock. **Created By** The name of the user who created the CPC in IBM Wave. **Create Time** The time stamp that indicates when the CPC was first created in IBM Wave. **Updated By** The name of the user who last updated the CPC. **Updated Time** The time stamp that indicates when the CPC was updated.

z/VM System Report

The z/VM System Report is used to generate a report about all the z/VM systems (LPAR) defined in IBM Wave.

The **Report Parameters** pane contains the following **Filter Parameters**:

System name
The name of the z/VM system that you want to use as a filter.

CPC name
The name of the CPC that you want to use as a filter.

Status By default, both "Active" and "Suspended" z/VM systems are displayed. To filter on systems that are active, select the "Active" check box. To filter on suspended systems, select the "Suspended" check box.

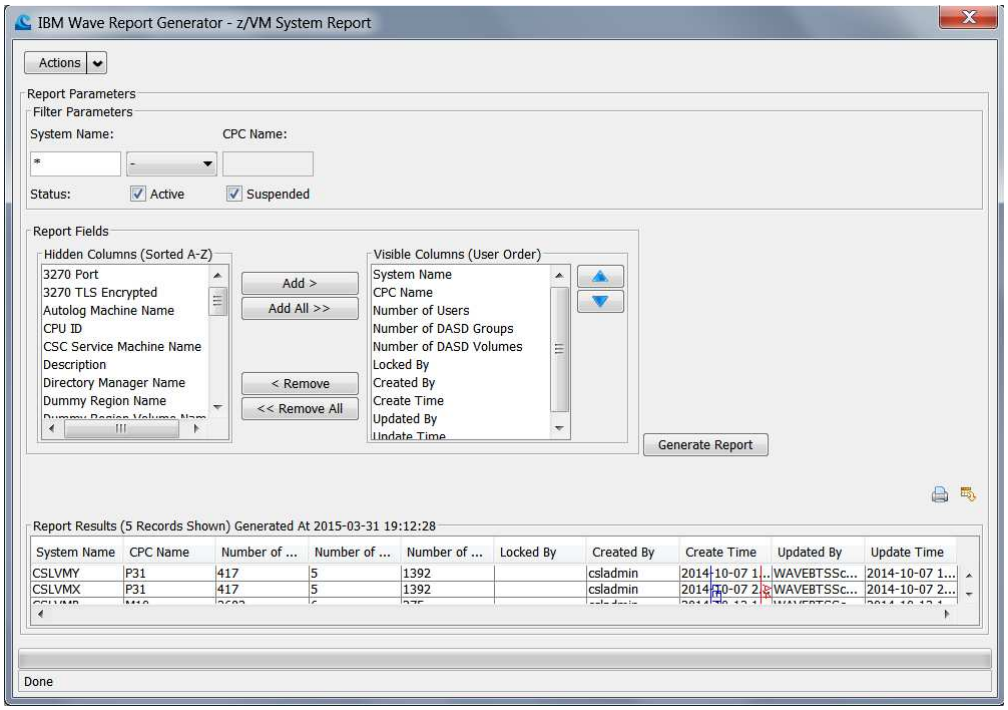

Figure 149. IBM Wave Report Generator: z/VM System Report

Table 14. Report fields z/VM System Report

Report fields in the *Hidden Columns (Sorted A - Z)*:	**Report fields in the** *Visible Columns (Use Order)*
The following fields apply: **3270 Port** The port for the 3270 emulator. **3270 TLS Encrypted** A yes or no field that indicates the state. **Autolog Machine Name** The name of the AUTOLOG machine. **CPU ID** The ID number for the central processing unit (CPU). **CPS Service Machine Name** The name of the CPS service machine. **Description** A description of the CPC. **Directory Manager Name** The name of the directory manager. **Dummy Region** The name of the dummy region. **Dummy Region Volume Name** The name of the dummy region volume. **EDEV Range** The EDEV address range. **Hostname** The host name. **IAN** The name of the user who created the IAN. **Long Service Machine Name** The name of the long service machine. **NFS Server** The name of the NFS server. **Number of CPUs** The number of processors.	The following fields apply: **System Name** The name of the z/VM system. **CPC Name** The name of the CPC under which this z/VM system was defined to IBM Wave. **Number of Users** The number of z/VM users. **Number of DASD Groups** The total number of DASD groups that are defined in the z/VM system. **Number of DASD Volumes** The total number of DASD volumes available in the z/VM system. **Locked By** Describes who holds the lock. **Created By** The name of the user who created the z/VM system on IBM Wave. **Create Time** The time stamp that indicates when the z/VM system was first created. **Updated By** The name of the user who last updated the z/VM system. **Updated Time** The time stamp that indicates when the z/VM system was updated.

Table 14. Report fields z/VM System Report (continued)

Report fields in the *Hidden Columns (Sorted A - Z)*:	**Report fields in the** *Visible Columns (Use Order)*
Number of Prototypes The total number of prototypes that are defined in IBM Wave. **Number of virtual networks** The total number of virtual networks that are defined for the z/VM system. **Performance Machine** The name of the performance machine. **SMAPI Port** The port number for SMAPI. **SMAPI TLS Encrypted** A yes or no field that indicates the state. **Service Machine IP address** The IP address that is defined for the service machine. **Short Service Machine Name** The name of the short service machine. **Short Service Machine Port** The port number for the short service machine. **Signal Shutdown Timeout** The default shutdown timeout used to deactivate Linux guests. **Status** The status for the z/VM System. Works with the check box defined in the **Filter Parameters**. **Type** The type of z/VM System. **z/VM IP** The IP address for the z/VM system. **z/VM IP6** The IP6 address for the z/VM system. **z/VM Name** The name of the z/VM system. **z/VM Service Level** The service level that is applied to the z/VM system. **z/VM Version** The version of the z/VM operating system.	

z/VM Guest Report

The z/VM Guest report generates a report listing all of the z/VM Guests from all of the z/VM Systems defined in IBM Wave.

The **Report Parameters** pane contains the following **Filter Parameters** that you can use as filters for the report data:

Guest name
The name of the z/VM guest.

Project
The project name.

System
The z/VM System name.

Distribution
The type of distribution.

Group The name of the group.

Status By default, both "*Active*" and "*Inactive*" z/VM guests are displayed. To filter on systems that are active, select the "*Active*" check box. To filter on inactive systems, select the "*Inactive*" check box.

Type Select "*Linux Only*" to view only the Linux guest distributions in the report data.

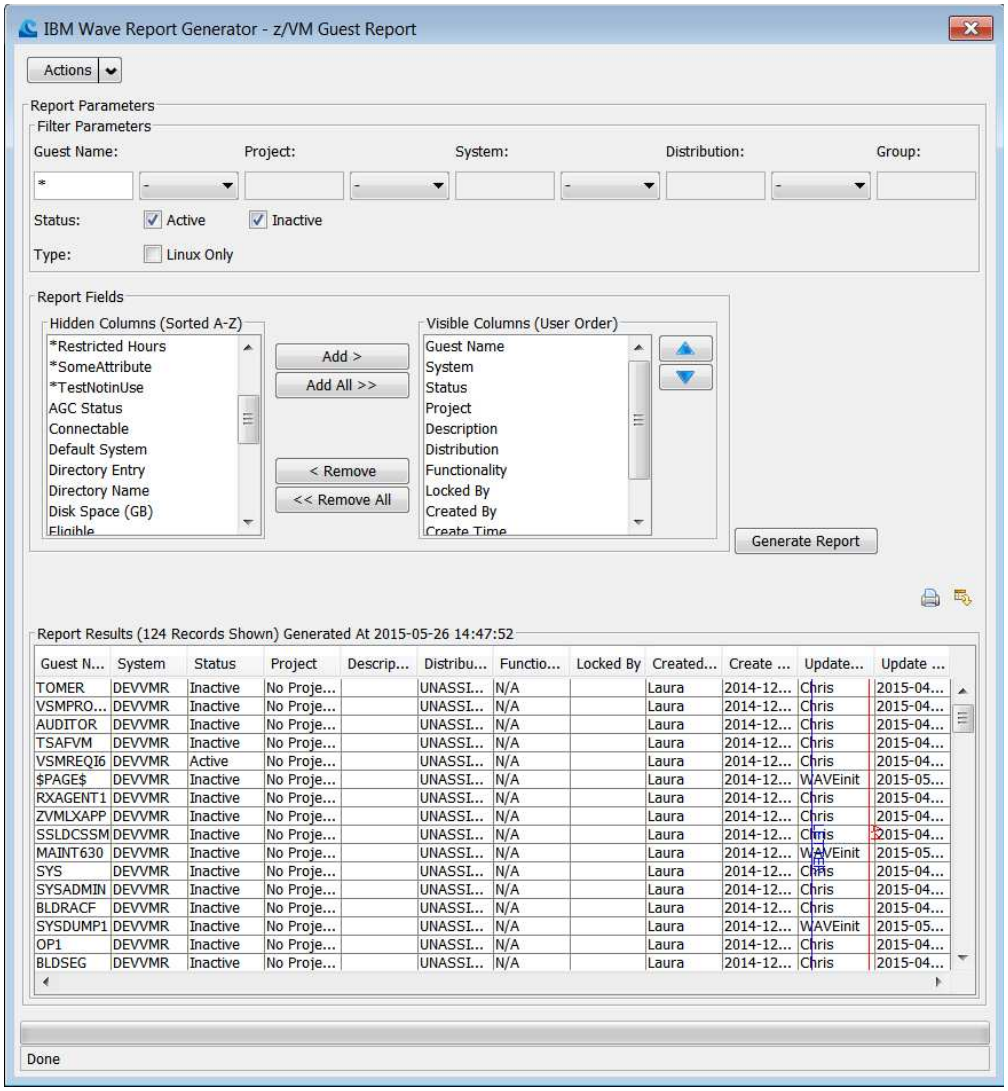

Figure 150. IBM Wave Report Generator: z/VM Guest report

Table 15. z/VM Guest report.

The Report Fields in the *Hidden Columns (Sorted A - Z)***:**	**The Report Fields in the** *Visible Columns (Use Order)***:**
Custom attributes By default, custom attributes that are assigned to the objects in the report appear in the hidden column and are preceded by an asterisk. For example, the "*RestrictedHours*" attribute in Figure 150 on page 202. **AGC Status** The status for Automatic Guest Classification (AGC). **Connectable** A yes or no field that indicates the state. **Default System** The name of the default system. **Directory Entry** The directory entry. **Directory Name** The directory name. **Disk Space (GB)** The disk space in gigabytes. **Eligible** A yes or no field. **Group** The name of the group. **IAN** The name of the user who created the IAN. **Inconsistent** Indicates whether the object in consistent or inconsistent. **Linux Hostname** The Linux host name. **Memory Max** The maximum amount of memory. **Memory Min** The minimum amount of memory. **Number of CPUs** The number of processors. **Original Image** The name of the original image.	**Guest Name** The name of the z/VM guest. **System** The z/VM system to which the z/VM guest belongs. **Status** The status of the z/VM guest. **Project** The name of the project to which the z/VM guest belongs. **Description** The description of the z/VM guest. **Distribution** The operating system distribution for the z/VM guest. **Functionality** The functionality of the z/VM guest. **Locked By** Describes who holds the lock. **Created By** The name of the user who created the z/VM guest. **Create Time** The time stamp for when the guest was first created. **Updated By** The name of the user who last updated the z/VM guest. **Updated Time** The time stamp that indicates when the z/VM guest was updated.

Table 15. z/VM Guest report (continued).

The Report Fields in the *Hidden Columns (Sorted A - Z)***:**	**The Report Fields in the** *Visible Columns (Use Order)***:**
Prototype The prototype that is defined. **Wave Init Status** The status of the initialization process for the IBM Wave server. **z/VM Account** The account number of the z/VM system. **z/VM Authorization Class** The authorization class for the z/VM system.	

z/VM Virtual Network Report

The z/VM Virtual Network Report generates a list of all the virtual networks from all z/VM systems defined to IBM Wave.

The **Report Parameters** pane contains the following **Filter Parameters** that you can customize for your environment and add to your report:

Device Name
The name of the device.

System
The name of the z/VM system.

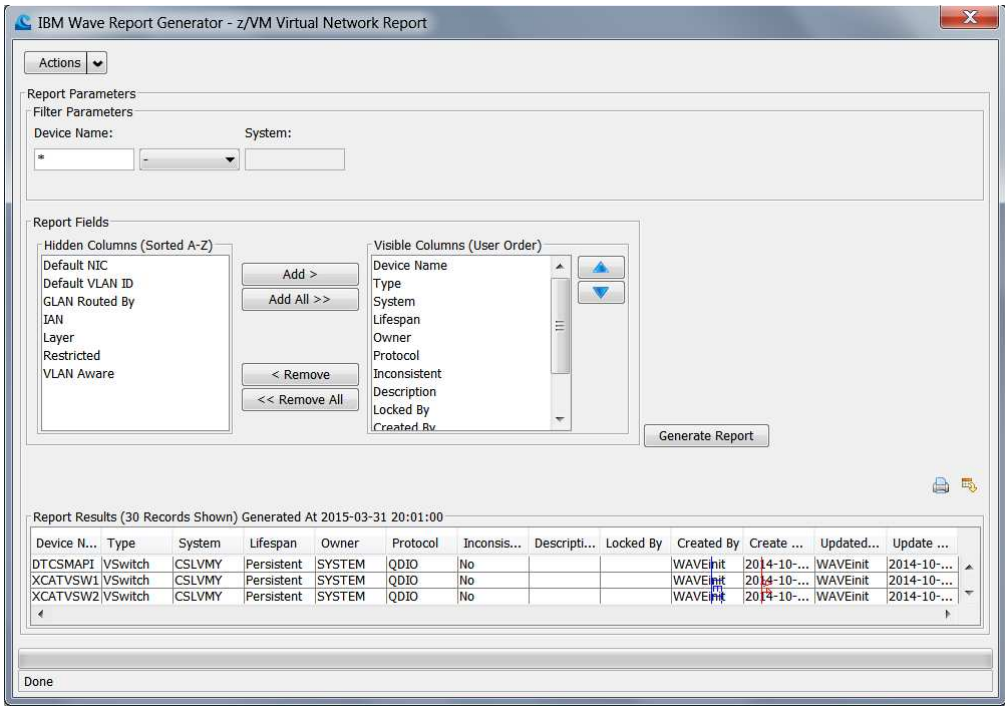

Figure 151. IBM Wave Report Generator: z/VM Virtual Network Report

Table 16. z/VM Virtual Network report

The Report Fields in the *Hidden Columns (Sorted A - Z)*:	**The Report Fields in the** *Visible Columns (Use Order)*:
Default NIC The default NIC. **Default VLAN ID** The default VLAN ID. **GLAN Routed By** The routing information for the Guest LAN **IAN** The name of the user who created the IAN. **Layer** The layer type for the virtual network. **Restricted** Indicates if the device is restricted. **VLAN Aware** The VLAN aware status.	**Device Name** The device name. **Type** The type of the virtual network. **System** The system name. **Life Span** The life span of the virtual network. **Owner** The owner of the virtual network . **Protocol** The protocol in use for the virtual network. **Inconsistent** Indicates whether the virtual network is in an inconsistent state. **Description** A description of the virtual network. **Locked By** Indicates who locked the virtual network. **Created By** The name of the user who created the virtual network. **Create Time** The time stamp for when the virtual network was created. **Updated By** The name of the user who last updated the virtual network. **Updated Time** The time stamp for when the virtual network was updated.

z/VM Virtual Network Connections Report

The z/VM Virtual Network Connections Report generates a report of the network connections that are defined in the z/VM LPAR.

The **Report Parameters** pane contains the following **Filter Parameters**:

Guest name
The name of the guest.

System
The name of the system.

LAN Name
The name of the LAN.

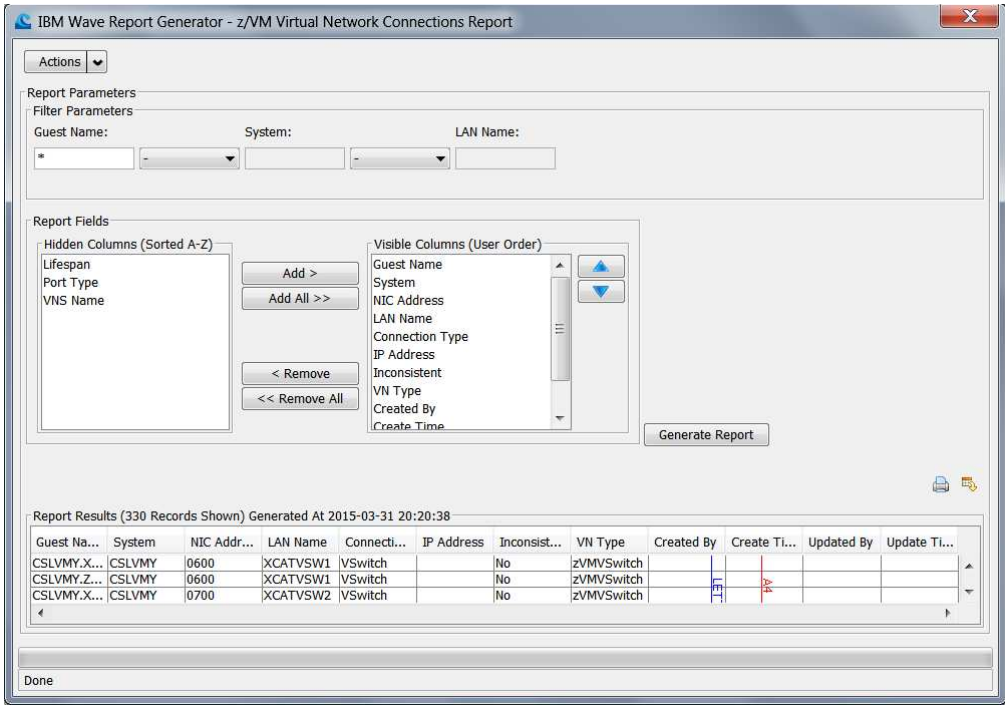

Figure 152. IBM Wave Report Generator: z/VM Virtual Network Connections report

z/VM Virtual Network Connections Report

Table 17. z/VM Virtual Network Connections Report

The Report Fields in the *Hidden Columns (Sorted A - Z)*:	**The Report Fields in the** *Visible Columns (Use Order)*:
Lifespan The life span of the VSwitch. **Port Type** The port type. **VNS Name** The name of the VNS.	**Guest Name** The name of the z/VM guest. **System** The name of the z/VM system. **NIC Address** The virtual NIC device for the connection. **LAN Name** The name of the GLAN/VSWITCH. **Connection Type** The type of connection (temporary, permanent). **IP address** The IP address the IBM Wave detected for this interface. **Inconsistent** Indication if this connection is inconsistent (either the LAN or the z/VM guest were deleted or are in some sort of error). **VN Type** The virtual network type. **Created By** The name of the user who created the virtual connection. **Create Time** The time stamp that indicates when the virtual connection was first created. **Updated By** The name of the user who last updated the virtual connection. **Updated Time** The time stamp when the connection was updated.

z/VM Prototype Report

The z/VM Prototype report generates a list of all the prototypes that are defined in all the z/VM Systems that are defined in IBM Wave.

The **Report Parameters** pane contains the following **Filter Parameters**:

Prototype name
The name of the prototype.

System
The name of the z/VM system that you want to use as a filter.

Project
The name of the project.

DASD Group
The name of the DASD group.

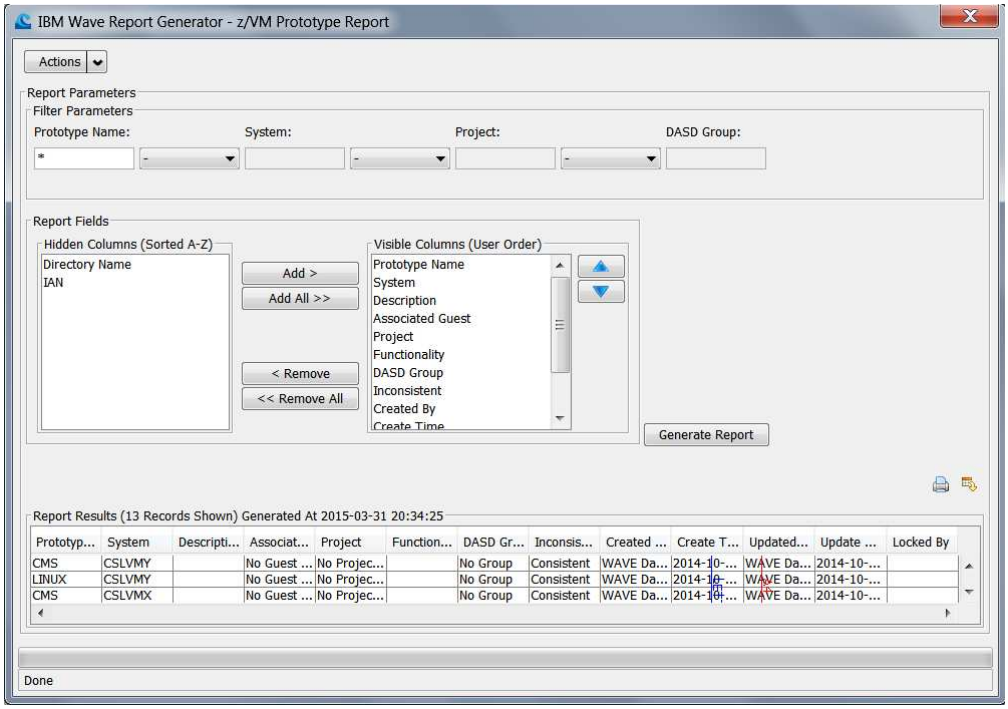

Figure 153. IBM Wave Report Generator: z/VM Prototype Report

z/VM Prototype Report

Table 18. z/VM Prototype report

The Report Fields in the *Hidden Columns (Sorted A - Z)*:	**The Report Fields in the (***Visible Columns (Use Order)*:
Directory name The name of the directory. **IAN** The name of the user who created the IAN.	**Prototype Name** The prototype name. **System** The z/VM system to which the prototype belongs. **Description** The description of the prototype. **Associated Guest** The z/VM guest that is associated with the prototype. **Project** The project to which the prototype belongs. **Functionality** The functionality that is assigned to the prototype. **DASD Group** The DASD group that is assigned to the prototype. **Inconsistent** Indication whether the prototype is in the inconsistent state. **Created By** The name of the user who created the prototype. **Create Time** The time stamp that indicates when the prototype was first created. **Updated By** The name of the user who last updated the prototype. **Updated Time** The time stamp that indicates when the prototype was updated. **Locked By** The name of the user who holds the lock.

z/VM DASD Group Report

The z/VM DASD Group Report generates a list of all DASD groups that are defined in the z/VM Systems managed by IBM Wave.

The **Report Parameters** pane contains the following **Filter Parameters**:

DASD Group
The name of the DASD Group.

System name
The name of the z/VM system.

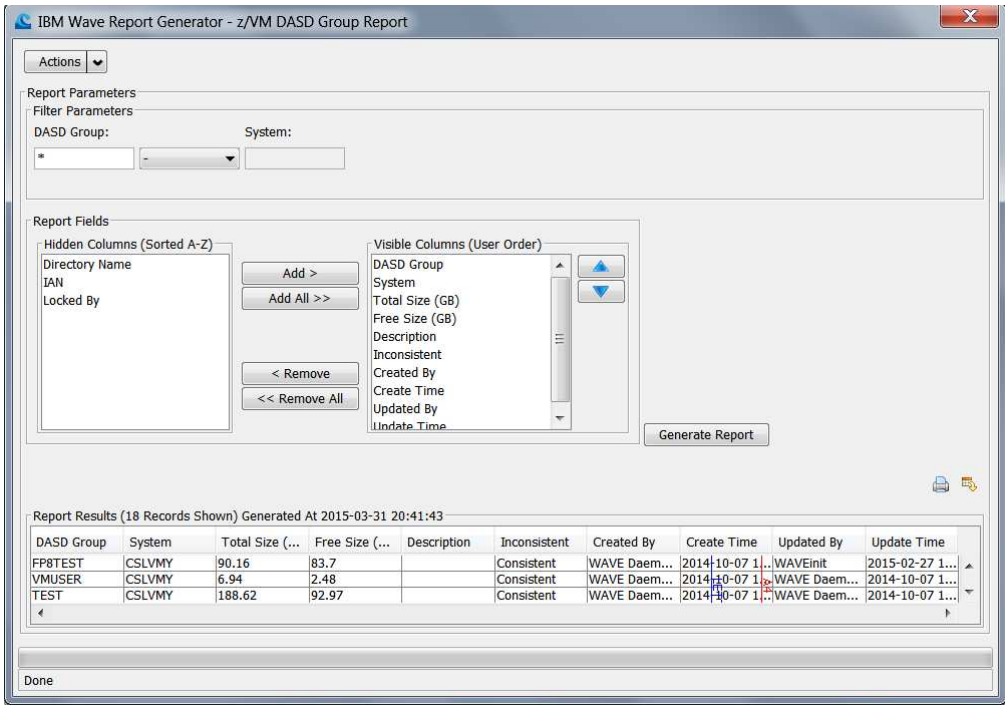

Figure 154. IBM Wave Report Generator: z/VM DASD Group report

z/VM DASD Group Report

Table 19. z/VM DASD Group Report

The Report Fields in the *Hidden Columns (Sorted A - Z)***:**	**The Report Fields in the (***Visible Columns (Use Order)***:**
Directory name The name of the directory. **IAN** The name of the user who created the IAN. **Locked By** Describes who holds the lock.	**DASD Group** The name of the DASD group. **System** The z/VM system to which the DASD group belongs. **Total Size** The total size, in GB of the DASD group. **Free Size** The total free space, in GB, in the DASD group. **Description** The description of the DASD group. **Inconsistent** Indication whether the DASD is in the inconsistent state. **Created By** The name of the user who created the DASD group. **Create Time** The time stamp that indicates when the DASD group was first created. **Updated By** The name of the user who last updated the DASD group . **Updated Time** The time stamp that indicates when the DASD group was updated.

z/VM DASD Volume Report

Describes the z/VM DASD Volume Report.

The z/VM DASD Volume Report generates a list of all the DASD volumes available for the z/VM systems that are managed by IBM Wave.

The **Report Parameters** pane contains the following **Filter Parameters**:

DASD Volume
The name of the DASD Volume that you want to use as a filter.

System
The name of the z/VM System that you want to use as a filter.

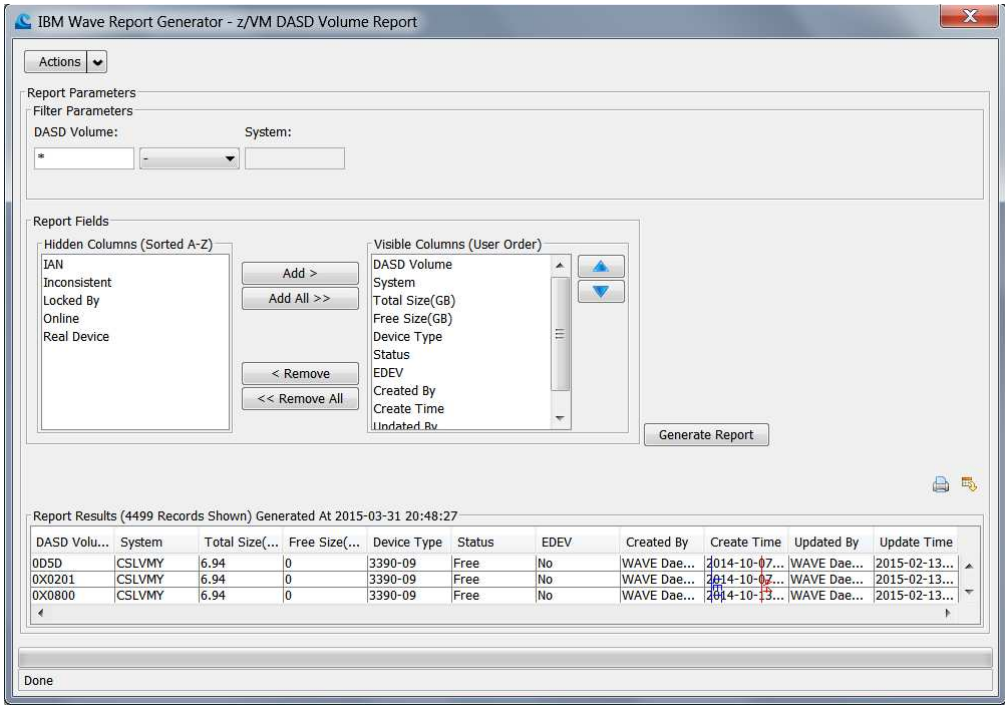

Figure 155. IBM Wave Report Generator: z/VM DASD Volume report

z/VM DASD Volume Report

Table 20. z/VM DASD Volume report

The Report fields in the *Hidden Columns (Sorted A - Z)*:	**The Report fields in the** *Visible Columns (Use Order)*:
IAN The name of the user who created the intelligent active note (IAN). **Inconsistent** A Yes or No field that indicates whether the volume is inconsistent. **Locked By** The name of the user who holds the lock. **Online** A Yes or No field that indicates whether the volume is online. **Real Device** The ASID of the real device.	**DASD Volume** The VOLSER of the DASD. **System** The z/VM system to which the DASD volume belongs. **Total Size** The size, in GB, of the DASD. **Free Size** The total free space, in GB, in the DASD group. **Device Type** The device type of the DASD volume. **Status** The status of the DASD volume. **EDEV** Indicates that the volume is an emulated device. **Created By** Indicates who created the DASD volume. **Create Time** The time stamp when the DASD volume was updated. **Updated By** Indicates who updated the DASD volume. **Update Time** The time stamp when the DASD volume was updated.

FCP Connections Report

The IBM Wave Fibre Channel Protocol (FCP) Connections Report generates a list of all the FCP devices that are defined in IBM Wave.

The **Report Parameters** pane contains the following **Filter Parameters** that you can add to the report:

FCP Device
The name of the FCP device.

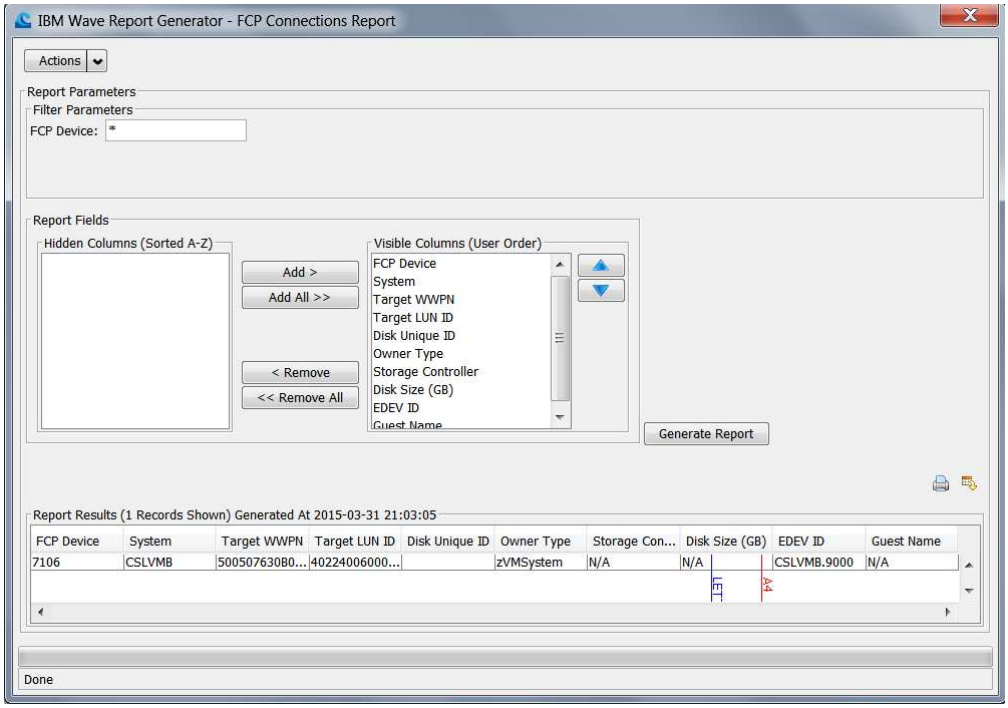

Figure 156. IBM Wave Report Generator: FCP Connections report

FCP Connections Report

Table 21. FCP Connections Report

The Report Fields in the *Hidden Columns (Sorted A - Z)***:**	**The Report Fields in the** *Visible Columns (Use Order)* **are as follows:**
None	**FCP Device** The ID of the FCP device. **System** The z/VM system to which the FCP device belongs. **Target WWPN** The WWPN assigned to the FCP device. **Target LUN ID** The LUN ID assigned to the FCP device. **Disk Unique ID** The unique ID for the disk. **Owner Type** The owner type. **Storage Controller** The storage controller ID. **Disk Size** The disk size in GB. **EDEV Address** The address for the emulated device (EDEV). **Guest Name** The guest name.

Flagged z/VM Object Report

The Flagged z/VM Object Report generates a list of all the objects that require attention.

The **Report Parameters** pane can contain the following **Filter Parameters**. By default, all of the objects appear in the printed report.

CPCs The name of the CPCs that requires attention.

Systems
The name of the z/VM systems that require attention.

Users The name of the users that require attention.

Prototypes
The name of the prototypes that require attention.

Virtual Networks
The name of the virtual networks that require attention.

DASD Groups
The name of the DASD groups that require attention.

DASD Volumes
The name of the DASD volumes that require attention.

Device Pools
The name of the device pools that require attention.

Real Devices
The name of the real devices that require attention.

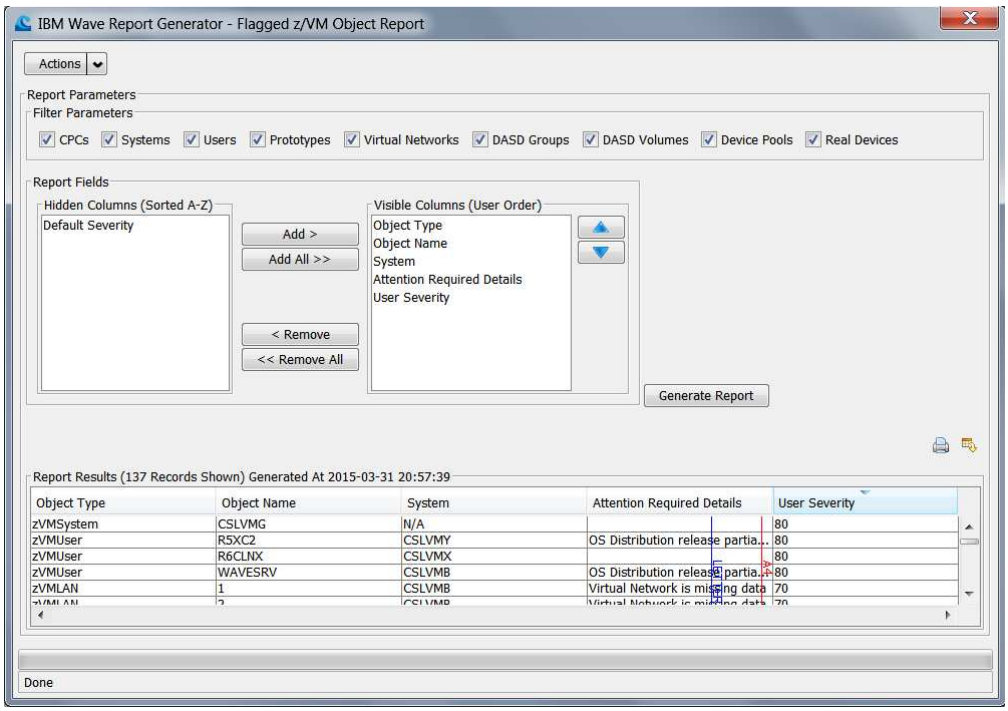

Figure 157. IBM Wave Report Generator: Flagged z/VM Object report

Flagged z/VM Objects Report

Table 22. Flagged z/VM Object report

The Report Fields in the Hidden Columns (Sorted A - Z) are as follows:	The Report Fields in the Visible Columns (Use Order) are as follows:
Default Severity Indicates the default severity for the object.	**Object Type** The type of object that is flagged. **Object Name** The name of the object that is flagged. **System** The z/VM system to which the flagged object belongs or N/A. **Attention Required Details** The text that describes all the conditions that led the object to be flagged. **User Severity** The highest severity of all the errors that led the object to be flagged.

IBM Wave Custom Attribute Report

Use the **IBM Wave Custom Attribute Report** to generate a report for the custom attributes that are classified in the **Custom Attribute Manager**.

The **Report Parameters** pane contains the following **Filter Parameters** that you can add to the report:

Attribute Name
The name of the custom attribute.

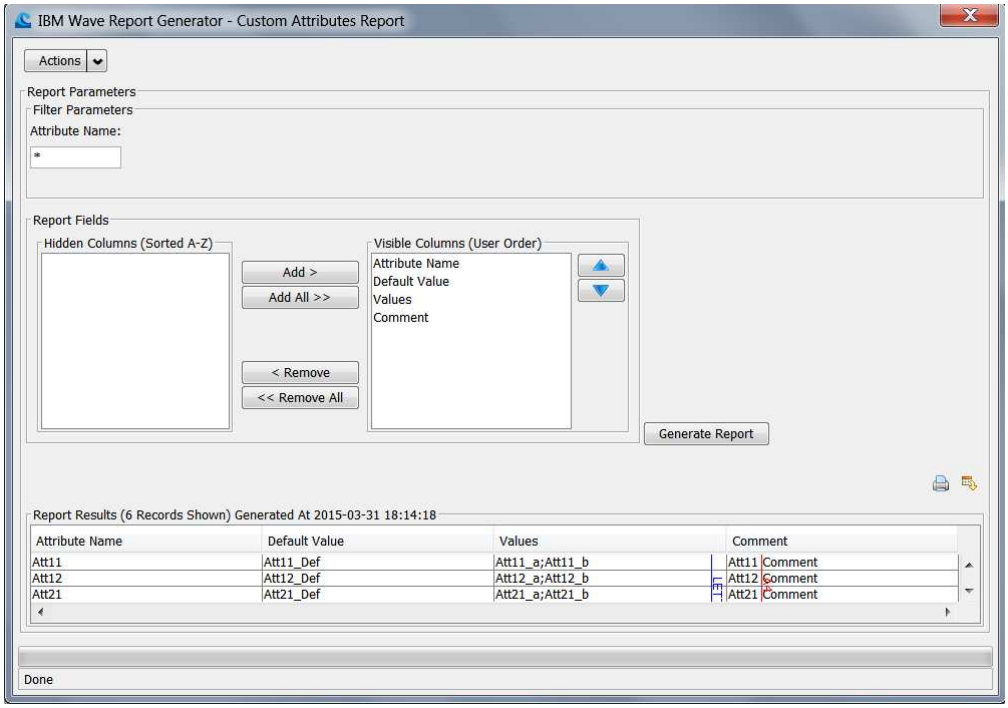

Figure 158. IBM Wave Report Generator: Custom Attribute report

- For custom attribute behaviors in a report, see "Custom attribute report behavior" on page 220.
- For information about using the Custom Attribute Manager, see the topic about the "Custom Attribute Manager" in .

Table 23. Custom Attribute Report

The Report fields in the *Hidden Columns (Sorted A - Z)***:**	**The Report fields in the** *Visible Columns (Use Order)***:**
None.	**Attribute Name** The name of the attribute that you define. **Default value** The default value for the custom attribute. **Values** The values that you define for the custom attributes. **Comment** Comment about the attribute.

Custom attribute report behavior

The behavior of the custom attribute report.

The following behaviors apply to the custom attributes that appears in the **Report Fields**:

- A custom attribute is added to IBM Wave, as a hidden column in the Custom Attribute report.
- When you delete one or more custom attributes from the **Custom Attribute Manager** while the **z/VM Guest** report is open, and the following conditions apply:
  - If one or more deleted attributes are hidden, the attribute is removed from the hidden list.
  - If a single deleted attribute is visible, you receive a warning message that indicates an attribute was removed from IBM Wave.
  - If more than one visible deleted attribute is removed, you receive a warning message that indicates the names of all the attributes that were removed from IBM Wave.
- When a saved report has visible columns that were removed, you receive a warning message when the report is opened that indicates:
  - A single custom attribute was removed from IBM Wave. The message indicates that the custom attribute name will be removed from the report when the report is saved.
  - The names of multiple custom attributes that were removed from IBM Wave. The message indicates that the custom attribute names will be removed from the report when the report is saved.

IBM Wave User Report

The IBM Wave User Report generates a list of details about the defined IBM Wave users and objects.

The **Report Parameters** pane contains the following **Filter Parameters**:

User Name
The name of the IBM Wave user.

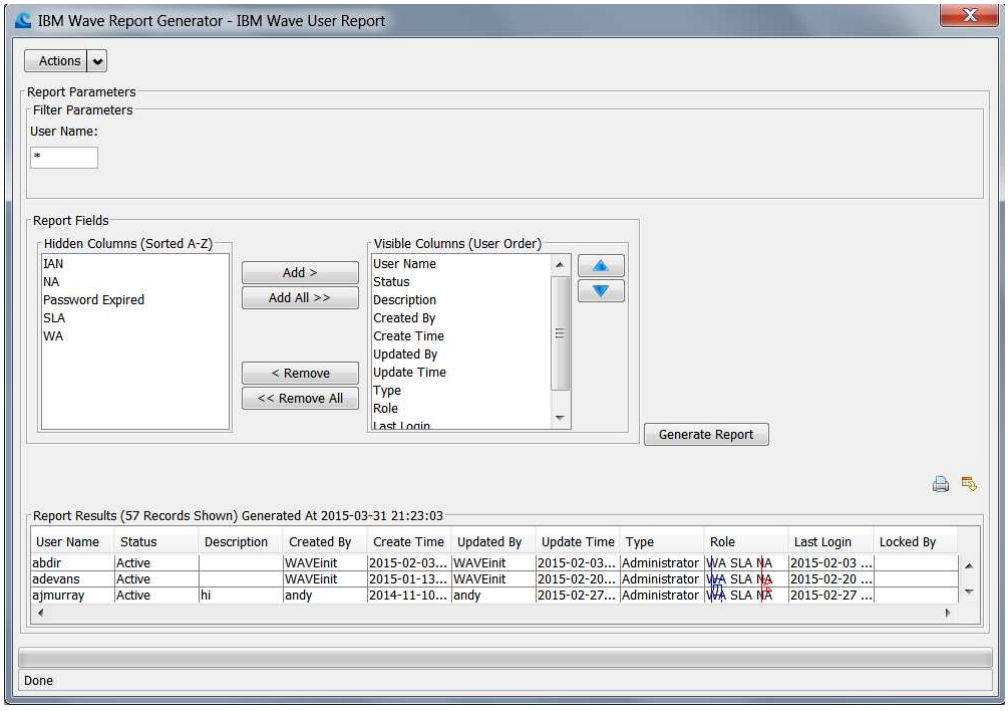

Figure 159. IBM Wave Report Generator: IBM Wave User report

IBM Wave User Report

Table 24. IBM Wave User Report

The Report Fields in the *Hidden Columns (Sorted A - Z)*:	**The Report Fields in the** *Visible Columns (Use Order)*:
IAN Indicates the name of the user who created has an intelligent active note (IAN). **NA** Indicates that the user has IBM Wave network administrator (NA) permissions. **Password Expired** A yes or no field that indicates if the password expired. **SLA** Indicates that this user has IBM Wave site level administrator (SLA) permissions. **WA** Indicates that the user has IBM Wave administrator (WA) permissions.	**User Name** The name of the IBM Wave user. **Status** The status of the IBM Wave user (active or suspended). **Description** An optional description of the user that is added when the user is created or modified. **Created By** The name of the user who the created the IBM Wave user. **Create Time** The timestamp when the user was created. **Update By** The last program or user to modify the user ID. **Update Time** The timestamp of the update. **Type** The type of the IBM Wave user (administrator or regular). **Role** The permissions of the IBM Wave user (WA, SLA, NA). **Last Login** The timestamp for the user's last login. **Locked By** The name of the user who holds the lock.

IBM Wave User Permission report

The IBM Wave User Permission report can be used to determine the access level permissions and scopes for IBM Wave users. For example, you can generate a report that lists all of the IBM Wave users that have access to a specific system, project, LAN, and DASD group.

The **Report Parameters** pane contains the following **Filter Parameters** that you can customize and add to your report. Filter parameters are added to the hidden column.

User Name
The user name of the IBM Wave user that you want the report to use as a filter.

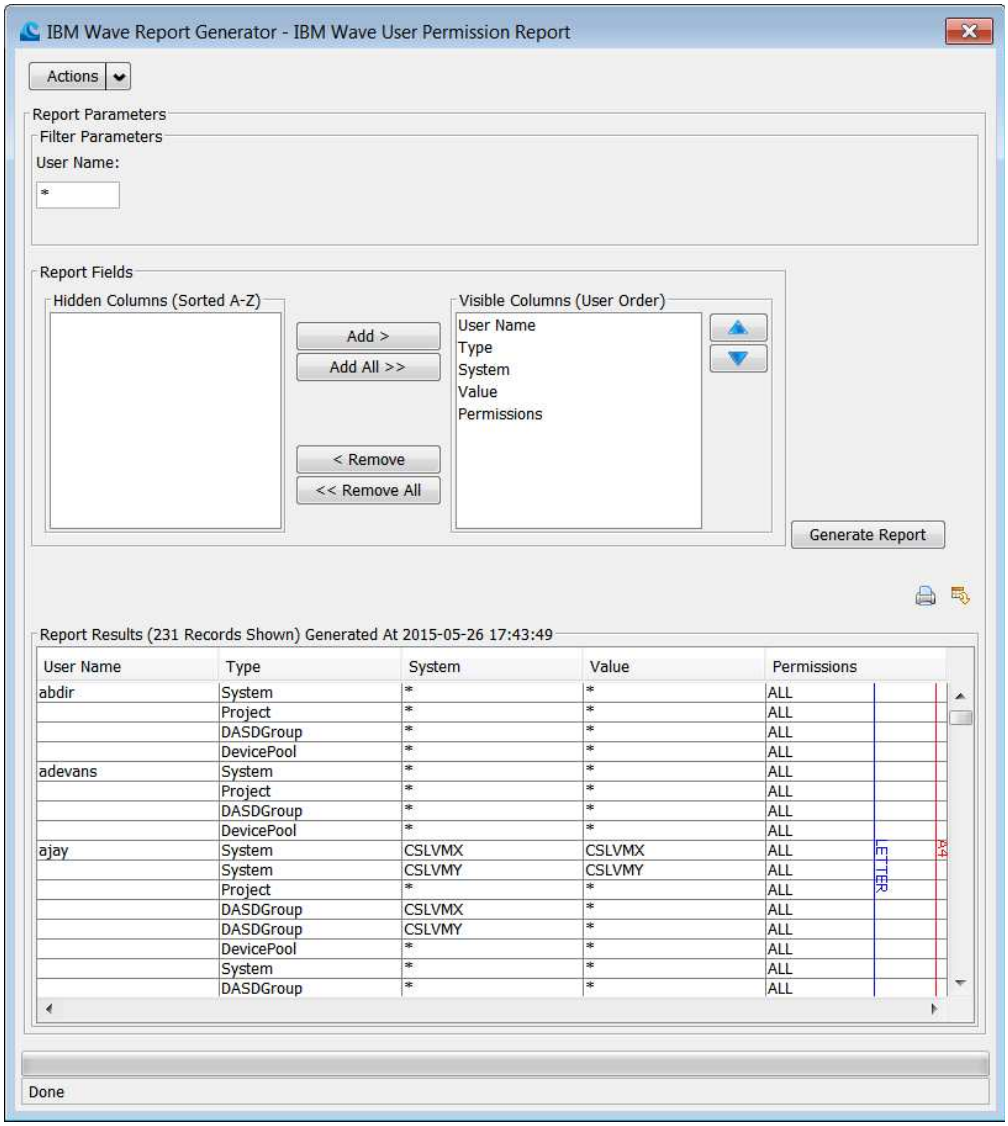

Figure 160. IBM Wave Report Generator: User Permission report

IBM Wave User Permission Report

Table 25. IBM Wave User Permission report

The fields in the *Hidden Columns (Sorted A - Z)***:**	**The Report fields in the** *Visible Columns (Use Order)***:**
None.	The report fields are: **User Name** The name of the IBM Wave user. **Type** The type of permission for System, Project, DASD group, and Device Pool. **System** The name of the z/VM system to which this scope belongs. **Value** The value of the scope. **Permissions** The permissions in this scope.

Chapter 4. Script management subsystem

IBM Wave contains a script management subsystem. The script subsystem helps you (with the correct permissions) create, delete, edit, and save a Linux shell scripts. You can then run the script on one or more virtual servers by using the **Execute Script** task.

The descriptive data and the scripts are stored in two locations:

- The descriptive data is stored in the IBM Wave database.
- The scripts are stored in a directory on the WAVESRV virtual server.

When you run a script, the first thing that happens is the script directory is mounted on the Network File System (NFS) for the target virtual server. This process allows scripts to call other scripts (because all the scripts are saved in one directory).

A script can be private or global. A global script is visible to a site-level administrator (SLA), which is useful for scripts that are common to all projects. A private script is visible to only the IBM Wave user who created it. Only an SLA has permission to change the global status by using the check box on the script editor. For all other users, the global field is disabled.

Tip: The action that you can take against a script depends on your scopes and permissions. When you hover over a script, a tooltip indicates your permission status.

Each script is assigned to a category for management. The categories are dynamically created, and need not be predefined.

IBM Wave User Script Manager

To open the **Script Manager**, from the IBM Wave main menu, select **Users Tasks** > **IBM Wave User Script Manager**.

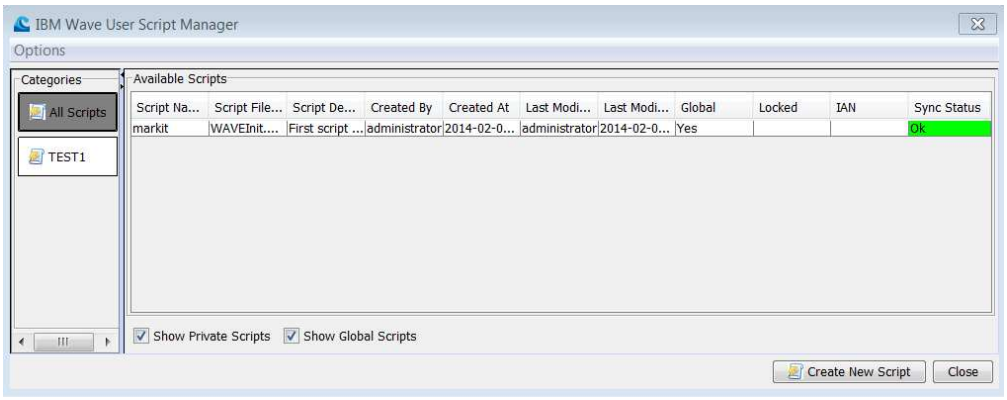

Figure 161. IBM Wave User Script Manager

The **Categories** pane of the window contains script categories and functions as a filter. Clicking one of the category filters the list of scripts in the **Available Scripts**

© Copyright IBM Corp. 2007, 2018

pane to include only the scripts that are assigned to the category. The "All Scripts" option disables the filter and shows all of the scripts in the system.

The right pane contains a set of fields describing the script:

Script Name
The name of the script.

Script File Name
The file name under which the script is saved on the WAVESRV server in the `/usr/wave/scripts` directory. The script file name must be unique, and must not contain spaces. The file name is generated automatically when a script is saved. If any spaces exist in the script name, the space are replaced with underscores ("_").

Script description
This is an optional description for the script.

Created By
The IBM Wave user who created the script.

Created On
The date and time when the script was created.

Last Modified By
The last IBM Wave user to save the script.

Last Modified On
The date and time when this script was last saved.

Global
Indicates whether the script is global or private.

Locked
Indicates whether this script is locked, and if so, by whom.

IAN Indicates whether this script has an IAN attached to it, and if so, the contents of the IAN are described.

Sync Status
Indicates whether the script is successfully synced to all defined NFS servers.

Creating scripts

To open the **IBM Wave User Script Manager** , from the **IBM Wave Main Menu** click **User Tasks > IBM Wave User Script Manager**.

To create a new script, click **Create New Script**. The **Create IBM Wave User Script** window opens (as shown in Figure 162 on page 227).

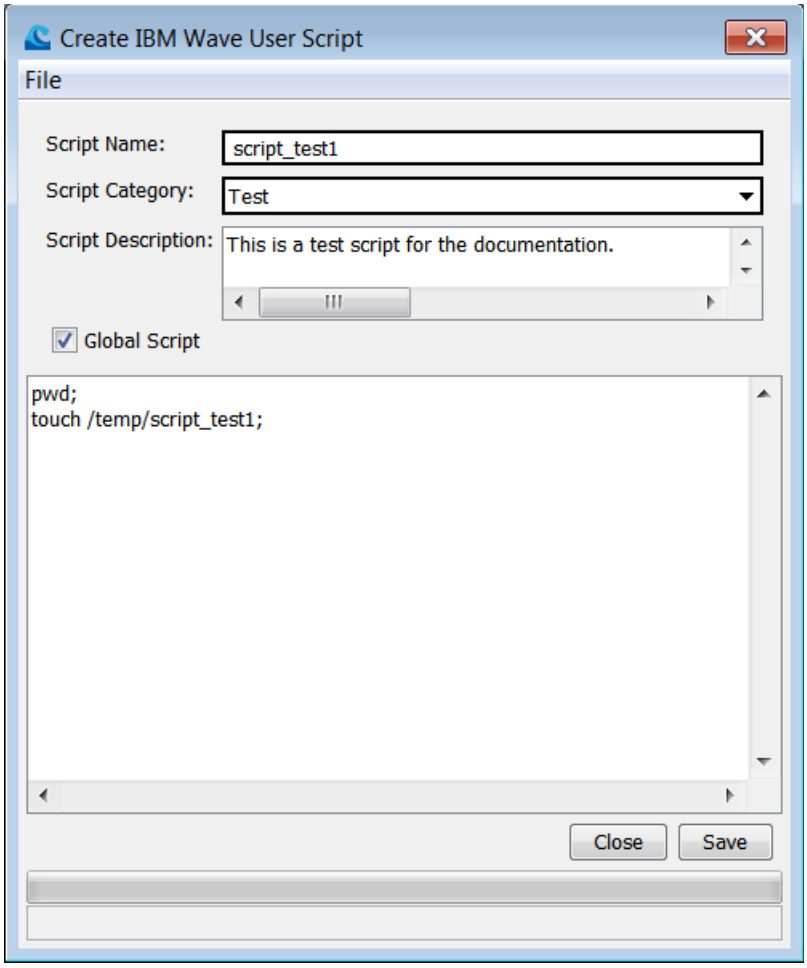

Figure 162. Create IBM Wave User Script

The **Create IBM Wave User Script** window contains the following fields:

Script Name
The script name field is mandatory.

Script Category
The category to which to assign this script. The drop down field contains all available categories. To create a new category, type over the text in the field.

Script description
An optional description for the script. The description typically contains basic information about what the script does.

Global Script
The option to create a global script that is available to all IBM Wave users. Otherwise, the script is private.

Script contents
This text input area where the script is written.

To create a new script, type it into the script contents text-input area. To save the script, click **Save**, or from the menu select **File** > **Save**. Saving the file creates it on the WAVESRV server in the `/usr/wave/scripts` directory with an automatically generated file name.

The file name includes the name of the user who created the script and a variation of the script name. Files that contain spaces makes script management difficult. If the script name contains a space, it is replaced with underscores ("_"). When the **Save** is complete, the window switches to the script editing window, and reloads the script from the script directory on the WAVESRV server. You can review the saved script contents before closing.

Note: If more NFS servers are defined, after you click **Save**, IBM Wave attempts to sync the script to all defined NFS servers. If one or more NFS servers fails the sync process, the sync status indication for the script indicates the failure.

Editing scripts

As a regular user, you have edit permissions for your own private scripts. You can perform the following tasks.

- Edit the script contents.
- Change the descriptive data related to the script.

Reminder: Only the site level administrator (SLA) can edit global scripts. To quickly check the permissions of a script, hover over the script to see the status.

To edit an existing script, double-click on its corresponding entry in the **Script Manager**, and the **Edit Script** pane opens. For a complete description of the fields in this pane, see “Creating scripts” on page 226.

When you are ready to save, click **Save**, or from the menu select **File** > **Save**.

Note: If additional NFS Servers are defined, after you **Save**, IBM Wave attempts to sync the script to all defined NFS Servers. If one or more NFS Servers fails the sync process, the “Sync Status” indicates the failure.

Deleting scripts

When you have the proper permissions to delete, you can delete a script from the **Script Manager**.

To delete an existing script, right-click on the script entry in the **Script Manager**, and select **Delete Script** from the menu.

The delete action removes the descriptive data for the script and the script file from the WAVESRV virtual server.

Syncing scripts

If one or more additional NFS Servers are defined, you can request that a particular script to be resynced to all NFS Servers, or a global resync for all scripts. To resync a particular script, right click the script and select "Sync This Script". For a global resync request, from the **Options** menu, select **Sync Scripts**.

Note: The re-Sync request is sent to the BTS. If during the resync one or more of the NFS Servers fails, this will be indicated in the **Sync Status** for the particular script. Also, IBM Wave does not perform a comparison of the script contents on the WAVESRV server against the contents on the NFS Server. Rather, it relies on its own sync indication. Manipulating the script outside of IBM Wave is not supported and can lead to undesirable results.

Running scripts from guests

About this task

You can run a shell script against one or more Linux guests in parallel. Site level administrators can edit a script before its run.

Procedure

1. Right-click on one or more guests, and then select **Execute Script**. The **Execute Script** menu opens as shown in Figure 163.

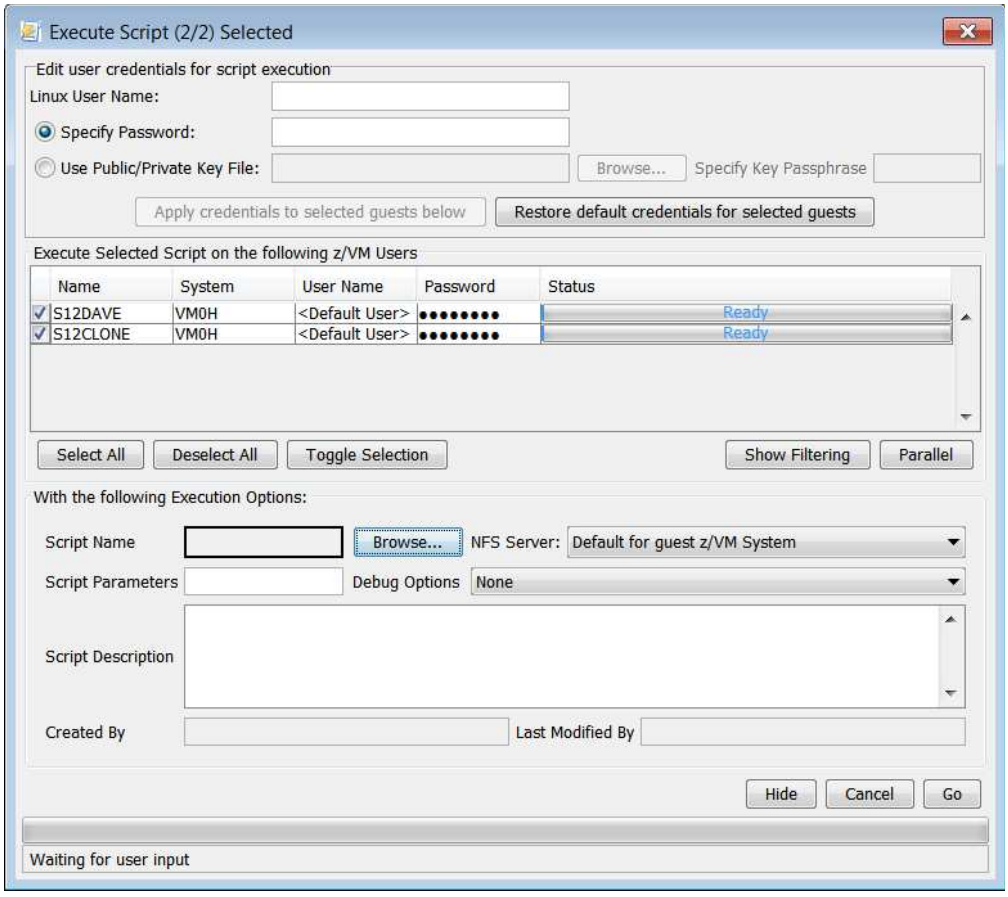

Figure 163. Running shell scripts from guests

2. Click **Browse...** to locate the correct script, and then double-click the script to load it into the "Script Name" field. The script is loaded.
3. If necessary, change any fields in the **Execute Script** menu, such as "NFS Server".
4. To run the script, click **Go**.

Results

The shell script is run against one or more z/VM Linux guests that you selected.

Chapter 5. Command line interface

IBM Wave provides a command line interface (CLI) to run specific IBM Wave actions using Linux shell scripts or Windows batch files. The IBM Wave CLI is built on Java™ technology.

CLI package

The IBM Wave CLI package is a compressed archive, either in `tar` file format for Linux deployments, or `zip` file format for Windows deployments. Installing the IBM Wave RPM on the Linux BTS generates the compressed CLI packages with the required scripts and places them into the `/usr/wave/WAVECLI` directory.

Note: The CLI requires software to unpack .zip and .tar files.

If either the `.tar` file or `.zip` file cannot be found in the `/usr/wave/WAVECLI` directory, you can build them using the `packWAVECLI` script, also located in the `/usr/wave/WAVECLI` directory.

- To package the JAR and script files as a `.tar` file use `/usr/wave/WAVECLI/packWAVECLI linux`. The script creates the **wavecli.tar** file within the `/usr/wave/WAVECLI`.
- To package the JAR and script files as a `.zip` file use`/usr/wave/WAVECLI/packWAVECLI windows`. The script creates the **wavecli.zip** file within the `/usr/wave/WAVECLI`.

Tip: To generate the **wavecli.zip** file, a package capable of creating a ***.zip** archive must be installed. Otherwise, the **packWAVECLI** script ends with an error condition.

Deploying the CLI package

To use the IBM Wave command-line interface (CLI), the IBM Wave CLI package must be deployed on the workstation or server that is running the CLI commands. The package must be stored in a dedicated directory.

These files will be included after installation of the package:

.jar file
The Java program and library files that are required to run the IBM Wave CLI.

wavecli.scr
The Linux **wavecli.scr** file is a front end for starting the CLI Java classes to run the IBM Wave CLI commands. The file supplied is an example that you can modify which requires your system PATH variable to specify the location for the Java virtual machine (JVM).

wavecli.bat
The Windows **wavecli.bat** file is a front end for starting the CLI Java classes to run the IBM Wave CLI commands. The file supplied is an example that you can modify which requires your system PATH variable to specify the location for the Java virtual machine (JVM).

© Copyright IBM Corp. 2007, 2018

CLI syntax

The IBM Wave command-line interface (CLI) must be used with only the supplied **WAVECLI.scr** or **WAVECLI.bat** files.

WAVECLI.scr ?
The syntax that displays a help message with all the available parameters and commands.

WAVECLI.scr -u user -bts ip [optional parameters] command <parameter parameter ...>
The syntax that is constructed to run the command.

The IBM Wave CLI syntax is as follows:

- Angle brackets for required parameters: **<-u>**
- Square brackets for optional parameters: **[-p] <password>**
- Vertical bars for choice of items: **deactivate <Guest Name> <IMMED | WITHIN 1-32767>**

CLI parameters

The CLI parameters are order-sensitive, which means the parameters must be written in a specific order for the work unit to be successful. For example, you must enter the required **[-tvp] <z/VM System>** parameter values before you enter the **activate** command. Instead of providing the required parameter values every time you run the CLI commands, you can edit the **WAVECLI.scr** and **WAVECLI.bat** files to always contain the correct values for the **[-u]**, **[-bts]**, and **[-tvp]** parameters.

The following are the required and optional parameters values for the CLI.

[-u] <IBM Wave User>
This required parameter specifies the IBM Wave user through which the IBM Wave CLI commands are issued. All generated BTS Work Units and log output is attributed to the specified user.

[-bts] <BTS IP Address>
This required parameter specifies the IP address of the WAVESRV where the BTS resides.

[-tvp] <z/VM System Name>
The required parameter for several IBM Wave CLI commands that specifies the z/VM System name to which the IBM Wave CLI commands refer.

Note: The z/VM System Name is the name given to the z/VM System when added to IBM Wave management. This name is not mandatory to match with the real z/VM system name/LPAR Name.

[-p] <IBM Wave User Password>
This optional parameter specifies the IBM Wave user's password. If this parameter is not specified, the IBM Wave CLI prompts interactively for the users password to be entered in a secure manner.

[-f] <filename>
This optional parameter for the file name specifies a full path of the local file, which contains a list of IBM Wave CLI commands that are run consecutively. For more information, see the topic about ""External CLI command file" on page 236".

[-forceLogin]
Use this optional parameter if a previous CLI call fails because the Wave resource serialization (WRS) entry for the CLI session was retained in the IBM Wave database.

[-ignoreIAN]
This optional parameter specifies the command is run regardless of whether an IAN exists for the specified object.

[-v]
An optional parameter, which specifies to use verbose output.

[-port]
This optional parameter specifies the port to use to connect to the BTS. If left unspecified, the default port of 3300 is used.

[-noCOR[
This optional parameter does not display the COR entry.

[-forceTLS]
This optional parameter establishes a connection with the BTS if TLS encryption is available and enabled.

Entering CLI commands

Each time the WAVE command-line interface (CLI) is run, the following events occur:

1. A syntax check is run against the parameters provided.
2. The IBM Wave CLI session is opened for the specified BTS.
3. The user credentials are verified with the BTS.
4. For each CLI command:
   a. Construct the BTS work unit.

      Note: When running the CLI file without the **-f** flag, there is only one CLI command, and therefore only one BTS work unit is constructed during the CLI session. When the **-f** flag is specified, a different BTS work unit is constructed for each command in the file.
   b. Submit the work unit to the BTS.
   c. Wait for the BTS work unit to finish processing.
   d. Print the output about the BTS work unit.
5. Close the IBM Wave CLI session with the BTS.
6. Print final return code.

When an IBM Wave CLI command is entered, an initial syntax check is run. If the syntax of the command is valid, the new BTS work unit is constructed and submitted to the BTS specified in the CLI parameters. The BTS work unit contains a BTS request that corresponds to the command entered. The BTS work unit can also be tracked from the IBM Wave GUI client.

If the **-v** flag is not set, IBM Wave CLI will output the COR entry for the BTS request and the return code of the BTS work unit by default. If the **-noCOR** flag is set, IBM Wave CLI does not output the COR entry for the BTS request.

After the BTS work unit finishes processing, IBM Wave CLI will end with a command return code and an exit return code. The return code of the CLI command will depend on the single BTS work unit submission ran during the CLI

session. If the **-f** flag was used, and the specified command file contains more than one CLI command, the exit code is the highest return code from any of the BTS work units that were submitted.

CLI commands

The following IBM Wave commands are supported for the command line interface (CLI):

The following CLI commands require you to set the **-tvp** flag to specify the z/VM System for the specific z/VM guest:

- **activate**
- **deactivate**
- **pause**
- **recycle**
- **resume**
- **signalActivationDone**

activate <*Guest Name*>
Activate the specified z/VM Guest.

deactivate <*Guest Name*> [IMMED|WITHIN 1-32767]
Deactivate the specified z/VM Guest. Using the **IMMED** flag causes the specified z/VM Guest to be **FORCED**. Using the **WITHIN** flag causes a z/VM SIGNAL to be sent to the z/VM Guest.

Important: You must take different approaches when you deactivate z/VM Guests depending on the operating system that is running on the z/VM Guest.

- When Linux is run on the z/VM guest, do not use the **IMMED** flag.
- When CMS is run, the **WITHIN** flag cannot shut down a CMS guest. By default, CMS guests cannot trap and handle z/VM SIGNAL, so **WITHIN** has no affect.

details <*Guest Name*>
Details about a specified z/VM Guest in the `.csv` file format. Without the **-tvp** flag, details are returned for the guest after its z/VM system is found. If systems are sharing a directory, only the active guest details are shown. Include the **-dir** flag to narrow the search to a particular directory. You can use **-dir** with or without specifying the **-tvp** flag. To specify a directory name that contains spaces, surround the name with double quotation marks. For example, **-dir "my dir name"**.

display <*systems*|*lpars*|*guests*> <*regex*> <*active*>
display <*directories*>
Outputs the specified data from the IBM Wave database in `.csv` file format.

- Use the **active** parameter to return only the active guests.
- Use the **regex** parameter to filter the search results for **guests** or **systems** by name (**lpars** is a synonym for systems). The regex parameter is a Java Regular Expression that you can use to either:
  - Select a specific object (for example, “display systems LINXVM01”).
  - Filter to select a group of objects (for example, “display guests 'APPDB.*' ”).

Notes: To use a regular expression that contains special characters for your shell, you need to escape them (the set of escape characters and how to escape them is specific to your choice of shell, so they are not explained here in detail).

- If the expression contains a space, you might have to quote the expression to ensure that it is passed as a single argument to the CLI.
- When you use special characters, you might need to use quotation marks or backslashes to prevent shell expansion. For example, to return all active guests that contain WAVE in the name:

  ```
  WAVECLI.scr -u user -bts 9.9.9.9 display guests 'WAVE.*' active
  ```

- When you want to match literal special characters, such as $ in the name, you might need to escape them twice in the regular expression. For example, to match a guest that is named $GUEST you might need to use a double backslash (\\) to escape the literal $ in a guest name resulting in the regex parameter of '\\$GUEST'.

For a summary of the Java regular expression constructs, see http://docs.oracle.com/javase/7/docs/api/java/util/regex/Pattern.html.

pause *<Guest Name>*
Pauses the guest by putting the guest in the z/VM **CP READ** state. To pause, the IBM Wave service machines issue the **CP STOP** command. To resume, run the **resume** <*Guest Name*> command.

recycle *<Guest Name>*
Recycles the guest by running the **deactivate** and **activate** commands. The **recycle** command disconnects the z/VM guest. When **recycle** completes, the changes to the z/VM definitions become active.

resume *<Guest Name>*
Resume the guest from the **pause** state. To resume, the IBM Wave service machines issue the **CP BEGIN** command.

status *<BTS Workunit ID>*
Retrieve status of a BTS work unit ID.

signalActivationDone *<Guest Name>*
Signals to the BTS that the activation of the specified z/VM guest is complete.

For more information about the activate action, see the topics "Overview of multiple task actions" on page 105 and "Activate" on page 110.

Receiving output from a CLI command

By default, for each IBM Wave command line interface (CLI) command, the executable prints out the COR entry generated by the corresponding BTS request, and finally the return code of the BTS Work unit. If the **-noCOR** flag is set, COR output is suppressed. If the **-v** flag is set, additional verbose information is printed.

CLI batch mode processing

Each time the IBM Wave CLI executable file is invoked with the **-f** flag, the CLI performs the following tasks:

1. Verify that a CLI command was not also specified from the command line.
2. Read each line of the input file, ignoring any empty lines, and then:
   a. Strip leading and trailing white space.

b. Consolidate all contiguous white space characters into a single white space character.
c. Perform a syntax check for the CLI command and its parameters.
d. End if the command requires CLI arguments that were not specified at invocation such as with **-tvp**.

Results: After the batch commands are read and the syntax is verified, processing continues as described in the topic "Entering CLI commands" on page 233.

External CLI command file

Starting the IBM Wave CLI with the **-f** (filename) parameter, and the path to the external command file causes the CLI to operate in batch mode. When you are in CLI batch mode, the IBM Wave CLI cannot accept any of the commands as program arguments. Instead, specify the commands in an external file in your file system.

For a list of commands, see the topic "CLI commands" on page 234.

External command file format

Each line of the external command file is treated as a separate IBM Wave CLI command together with its parameters. The first token of each line contains an IBM Wave CLI command and the remaining tokens are the command arguments. For each line, the following rules apply:

1. Each line can hold one command at most.
2. Empty lines are ignored.
3. Use white space characters to separate tokens.

```
DETAILS  WAVEWRKS
ACTIVATE  DANFITZ
DEACTIVATE        BADGUEST IMMED
DEACTIVATE  LINGUEST  WITHIN   10
```

Figure 164. Example: External CLI command file

Chapter 6. IBM Wave API

You can run IBM Wave actions such as activate, deactivate, recycle, and more under the protection of IBM Wave user's scope and permissions with the IBM Wave API. The IBM Wave API can be used with any REST client.

Starting and Stopping the Liberty Server

The IBM Wave RESTful API is hosted on an embedded WebSphere® Application Server Liberty installation. To use the API, you must first verify that the IBM Wave API server is running.

Stopping, starting, and querying the server status is done by running the following commands on the WAVESRV command line:

- To query the status of the server issue `/usr/wave/websphere/wlp/bin/server status`. If the server is running, the command outputs the following message:

  `Server defaultServer is running with process ID xxxxxxx`
- To start the server issue `/usr/wave/websphere/wlp/bin/server start`.
- To stop the server issue `/usr/wave/websphere/wlp/bin/server stop`.

You can start running API calls if the server is running and you completed the directions for customizing the API server. For more information, see .

Note: IBM Wave does not install scripts to start or stop the Liberty server as part of the WAVESRV boot process. You can embed the start and stop commands in your Linux startup or shutdown scripts.

HTTP Parameters

The REST APIs use the following HTTP parameters:

Header: Authorization: Basic XXXXXXXXXXXXXXXXX

- The IBM Wave user credentials are passed in the Authorization header of API call. The supported masking is Base64.

Header: X-CSRF-TOKEN: noCSRF

Cookie: CSRF-TOKEN=noCSRF

- The combination of these two parameters prevents CSRF attacks. The value in the cookie must match the value in the header.

Header: Content-type: application/json

- Mandatory when sending a JSON body to the API. If this parameter is omitted the API returns an HTTP response of 415.

Header: If-Match: *ETAG-VALUE*

- Mandatory when issuing a POST API call. Each POST call must have the if-match header with the ETag value of the specific resource the action is related to.

© Copyright IBM Corp. 2007, 2018

Optional Header: X-Requested-With: XMLHttpRequest
- This parameter is used to help prevent CSRF attacks.

Supported HTTP Verbs and Responses

IBM Wave API include the following HTTP verbs:
- GET
- POST

API calls retrieve or take action upon a specified resource. Successful action or retrieval of a resource depends on two aspects:

1. IBM Wave user scope and permissions - Is the IBM Wave user who submitted the API call allowed to view the resource he requested?
   - If an API call is submitted for a resource that is out of the IBM Wave user's scope, a response of 404 – resource not found - is returned.
   - If IBM Wave user is not permitted to perform an action a response of 401 – not authorized - is returned.
2. The state of the resource. The API follows the same limitations as the IBM Wave GUI. This means that if it is not possible to view or take action on a resource in the IBM Wave GUI because of the state of a resource, then it is not possible to perform that action from the API. For example, trying to activate a guest on a system that is suspended results in a response of 409 – resource conflict - with a reason that explains why the user receives the response.

General summary of HTTP responses:

200 OK. Returned for GET APIs. The body will contain the resource representation. The Etag header in the response contain the etag value used to the POST actions.

202 Accepted. Response from a successful POST command. A work unit has been submitted. The Location Header in the response contains the work unit Id.

304 Resource not modified since last request.

400 Bad request, Malformed or incomplete parameters.

401 Not authorized. Either the login failed or the IBM Wave user is not permitted to perform the action.

403 Forbidden, incorrect use of headers.

404 Resource does not exist or is out of the IBM Wave user scope.

409 Conflict in current resource state.

412 Precondition failure. The Etag validation failed.

500 Internal server error.

503 Unable to connect to the IBM Wave BTS.

Error responses 400 and 409 will also contain a body with additional information as to why the API call failed.

The body contains a JSON with the following format:

```
{
  "errorList": [
    {
      "field": the field with the error,
      "message": detailed message for the failure
    }
  ],
  "verificationResponseRC": Representation of the RC
  "workunitID": for future use.
  "errorExists": true – always.
  "errorListAsString": the error list from above in string format
  "submitRCMessage": General message regarding the failure.
}
```

If more then one error exists the **errorList** element will contain more items.

Example: Trying to GET the guest details of guest on a system that is suspended using the following command:

```
curl -H "X-CSRF-TOKEN: noCSRF" --cookie "CSRF-TOKEN=noCSRF" -H
"Authorization: Basic YmVuYzp0cG5zMjAxMw==" -k https://IP-ADDRESS:9443/
WaveAPIServer/ibm/wave/systems/SYSTEM-NAME/guests/GUEST-NAME/
```

Returns the following output:

```
{
  "errorList": [
    {
      "field": "guestId",
      "message": "The system is in Suspend state"
    }
  ],
  "verificationResponseRC": "RESOURCE_STATE_CONFLICT",
  "workunitID": null,
  "errorExists": true,
  "errorListAsString": "The system is in Suspend state ",
  "submitRCMessage": "The specified resource is not eligible for this action while in its current state"
}
```

HTTP Header Responses

A successful GET API response contains two areas of information.

1. The body - The body contains the resource representation in JSON format.
2. The headers - In the headers you can find the ETag header.

Note: The ETag value must be used when performing a POST call.

Example: To activate guest DEMO1 on system SYSTEM1 using an API call, the user must first get the ETag for the guest using the following call:**curl -H "X-CSRF-TOKEN: noCSRF" --cookie "CSRF-TOKEN=noCSRF" -H "Authorization: Basic YmVuYzp0cG5zMjAxMw==" -k https://IP-ADDRESS:9443/WaveAPIServer/ibm/wave/systems/SYSTEM1/guests/DEMO1/**. This get command returns a guest resource representation::

```
{
  "name": " DEMO1 ",
  "system": {
    "href": "https://IP-ADDRESS:9443/WaveAPIServer/ibm/wave/systems/SYSTEM1",
    "name": "SYSTEM1"
  },
  "defaultSystem": null,
  "distro": "Unassigned OS Distribution",
  "project": null,
  "sdg": {
    "href": "https:// IP-ADDRESS:9443/WaveAPIServer/ibm/wave/systems/ SYSTEM1/site-defined-groups/
```

```
    "name": "USER-LOCAL"
  },
  "status": "Inactive",
  "memoryMin": 512,
  "memoryMax": 1024,
  "cpus": 1,
  "lockedBy": "",
  "hostname": "",
  "description": "",
  "networks": [],
  "clusters": [],
  "initForWaveStatus": "NO",
  "customAttributes": [],
  "isConnectable": false,
  "directoryEntry": "https://IP-
ADDRESS:9443/WaveAPIServer/ibm/wave/systems/ SYSTEM1/guests/DEMO1/directory-entry",
  "selfUri": "https://IP-ADDRESS:9443/WaveAPIServer/ibm/wave/systems/ SYSTEM1/guests/DEMO1",
  "account": null
}
```

The response also contains the following headers:

```
Cache-Control →no-store
Content-Language →en-US
Content-Length →733
Content-Type →application/json;charset=utf-8
Date →Mon, 13 Feb 2017 16:51:07 GMT
ETag →"1872db1a556b33e34901e1fc2697fdef"
Pragma →no-cache
Strict-Transport-Security →max-age=604800; includeSubDomains
Vary →Origin
X-Content-Type-Options →nosniff
X-Powered-By →Servlet/3.1
```

The ETag header contains the ETag value. This value represents the current state of the resource when the GET API was performed. After the user receives the ETag he can do the following POST action. **curl -H "X-CSRF-TOKEN: noCSRF" --cookie "CSRF-TOKEN=noCSRF" -H "if-match: 1872db1a556b33e34901e1fc2697fdef " -H "Authorization: Basic YmVuYzp0cG5zMjAxMw==" -k https://IP-ADDRESS:9443/ WaveAPIServer/ibm/wave/systems/SYSTEM1/guests/ DEMO1 /actions/activate**. A response of 202 is returned with no body. The returned headers contain the work unit id similar to the example below.

```
Cache-Control →private,no-cache,no-store,no-transform,no-store
Content-Language →en-US
Content-Length →0
Date →Tue, 14 Feb 2017 12:23:16 GMT
Location →https://IP-ADDRESS:9443/WaveAPIServer/ibm/wave/workunits/07e1020e07171106b6
Pragma →no-cache
Strict-Transport-Security →max-age=604800; includeSubDomains
Vary →Origin
X-Content-Type-Options →nosniff
X-Powered-By →Servlet/3.1
```

Examples for using IBM Wave APIs

The IBM Wave REST APIs allow direct client calls to an API server. These calls are sent over HTTPS, as JSON objects. The API server in turn sends requests to the IBM Wave BTS, which runs the requested z/VM actions or returns data about the environment.

IBM Wave REST APIs can be called by any program that can communicate over HTTPS and supports cookies. When the IBM Wave API server is running in the

environment, it can receive calls from clients. HTTPS calls can be received by using Ajax, Perl, Java, Python, C, or a multitude of other scripts or sources.

The example offers a step-by-step explanation of how APIs can be used by writing a Linux Bash script, by using the program cURL and python to automate an API call through z/VM Operations Manager.

Currently, z/VM systems are managed on a per-system basis through Operations Manager. Operations Manager has an instance that is installed on each system and controls only its own system. However, IBM Wave has an environment-wide scope and can control guests on any z/VM system in the environment. Using Operations Manager schedules and actions, automation of z/VM tasks is made easier. Combining the IBM Wave API with Operations Manager allows you to use one Operations Manager schedule to manage the entire complex.

For example, by using this combination, it is possible to schedule a recycle of all the guests in a specific IBM Wave project for Friday at 11 PM.

The method for achieving that is by setting up a schedule in Operations Manager that runs a shell script on a Linux guest. The shell script uses the IBM Wave API and runs all the logic.

Abstract flow of the example

1. Schedule gets triggered.
2. Schedule starts action.
3. Action runs the CP SEND command with the Linux command or script name.
4. Script runs the IBM Wave API.
5. IBM Wave runs the IBM Wave actions.

The following example uses Operations Manager CMS commands that are entered from a CMS guest.

- Set up a Linux guest on the same z/VM system that is running Operations Manager. The Linux guest must have IP connectivity to the WAVESERV.
- Define an action in Operations Manager.

  ```
  gomcmd opmgrm1 defactn name(waveapi), command(CP SEND WSRV ./API.SH &p),
  env(svm)
  ```

  name(waveapi)
  Defines the actions name as `waveapi`.

  command(CP SEND WSRV ./API.SH &p)
  The action is to run the mentioned command.

  WSRV The name of the Linux guest that was set in the previous step.

  ./API.SH
  The name of the shell script that runs.

  &p Operations Manager parameter for sending parameters through scheduling.
- Define the schedule that runs the action. This example passes the project that is called `MYPROJECT` to the script.

  ```
  gomcmd opmgrm1 defschd name(waveschd), action(waveapi), parm(MYPROJECT),
  when(FRI-23:00)
  ```

name(waveschd)
The name of the new schedule.

action(waveapi)
The action that you want to run.

parm(MYPROJECT)
Parameter to send to the action. This parameter is passed to the Linux shell script.

when(FRI-23:00)
Time to run the action.

- Create a Linux shell script on the root file system called `API.SH`.

The Linux Shell Script

The sample script uses the curl command and Python to issue the API calls and work with the output.

Example cURL command

```
response=`curl -X "$1" -s -H "Content-Type: application/json"
-H 'X-Requested-With: XMLHttpRequest' -H 'X-CSRF-TOKEN: noCSRF' --cookie 'CSRF-TOKEN=noCSRF'
-k https://<user>:<pass>@<ip>:9443/ibm/wave"$2" 2>&1`
```

-s Provides "silent" execution so API calls can run without any excess output to the console. This parameter is optional.

-H This flag is used to provide HTTP headers.

-H "Content-Type: application/json"
Indicates that you want a JSON formatted response.

-H 'X-Requested-With: XMLHttpRequest'
Required HTTP header.

-H 'X-CSRF-TOKEN: noCSRF'
Required HTTP header.

--cookie 'CSRF-TOKEN=noCSRF'
Required cookie.

-k Indicates that the curl command is working in unsecure mode. No verification of SSL certificate is done.

Other Optional flags

-i Include the headers in the curl output.

-I Include only the headers in the curl output.

-w This parameter is used to print specific variables from the cURL output.

https://<user>:<pass>@<ip>:9443/ibm/wave"$2"
This parameter is the URL to call. HTTPS is used. Your Wave user name and password are provided along with the IP address of the API server. The port is 9443. All calls go the endpoint /ibm/wave followed by the call to run. The curl command transforms the user name and password into the authorization header with Base64 masking. You can also use the following header and omit the clear text user name and password.

-H "Authorization: Basic XXXX"
XXXX is the masked text. When this header is used, use `https://<ip>:9443/ibm/wave"$2"` as the URI.

2>&1 This parameter is a Linux redirection to standard output. This parameter enables the script to capture the output of the API call.

Parsing JSON with Python

The code is an example of using python from within a shell script.

```
getValueFromJSON() {
 JSONresult=`echo "$1" | python -c 'import json,sys;obj=json.load(sys.stdin);print('"$2"')'`
 echo "$JSONresult"
}
guestname=$(getValueFromJSON "$guests" 'obj["guests"]['"$g"']["name"]')
```

Any API call that returns a JSON response can use the `getValueFromJSON` method to retrieve specific values from the JSON.

The function accepts a JSON string or a variable that contains one, such as the output from the `callWaveAPIGET` function, as the first parameter. The JSON is used as input to the Python code, which has its own JSON parsing that are built in. This function is useful for extracting data values based on their property names or keys, such as the name of a system or the length of a JSON array.

You can parse the JSON to find out how many guests there are. In the script `obj` is the JSON object as defined in the Python program. The script is finding the `guests` property of the object, which returns an array of guests. To find the number of guests, the example gets the length of the array by using the `len()` method.

```
guestarrlen=$(getValueFromJSON "$guests" '(len(obj["guests"])-1)')
```

This example is retrieving the guest name from the first place in the guest collection. The guest collection is saved in the “guests” attribute.

```
guestname=$(getValueFromJSON "$guests" 'obj["guests"]['"0"']["name"]')
```

IBM Wave API Sample

The IBM Wave API example resides in `/usr/wave/samples/ibm_wave_api_sample.sh`.

Appendix A. Example for using IBM Wave

In some installations, it is common to have certain operational procedures for z/VM Systems and Linux virtual servers. IBM Wave provides many features and functions to manage and interact with the z/VM operating system and the Linux operating system. However, because most of the IBM Wave features and functions are "atomic", it can be useful to outline a procedure that uses "atomic"actions to achieve a procedural goal.

Relocating a guest stored on shared or remote copy storage

In some installations, it is common that the production z/VM Systems and the Disaster Recovery (DR) z/VM Systems use shared storage, or that the production and DR storage is remote copy. Linux guests might also be defined on the shared (or remote copy) DASD volumes, and can be relocated upon request between the z/VM Systems sharing the storage. Currently, no "Action" in IBM Wave provides the full functions, but some *atomic* actions can be used to achieve the same results.

For guest relocation, typically the z/VM system programmer must deactivate the guest on the source z/VM system, vary all of the DASD volumes that are associated with the guest offline, vary the DASD volumes online in the target z/VM System, and then activate the guest. IBM Wave provides an intuitive interface to do the same actions.

If both the source and the target z/VM Systems are managed with IBM Wave, it is possible to use IBM Wave to do the same tasks. IBM Wave can automatically detect the Linux guests and DASD volumes on every managed z/VM System. If, for example, the Linux guest is active on system A, IBM Wave can display the Linux guest as active and the DASD Volumes associated with it as online. If the same guest is defined in system B, IBM Wave can display the guest as disabled and its DASD volumes as offline.

To relocate a guest that is stored on shared, or remote copy storage by using IBM Wave, complete the following steps:

1. Customize the XVDSKOFF and XVDSKON user exits. These exits are called whenever a DASD Volume is varied offline or online from within IBM Wave. For more information about the XVDSKOFF and XVDSKON exits, see "*IBM Wave User Exits*" in *IBM Wave for z/VM: Administration and Customization*. It might be necessary to further customize the examples that are provided. If, for example, the storage is not shared, but instead uses remote copy. In this case, it is probably necessary to detach the remote copy link when varying the DASD volume offline, and change the DASD volume to read/write before varying the DASD volume online.
2. In the source z/VM System, complete the following actions:
   a. Select the guest or guests that need relocation, and run the "Deactivate" action on them. For more information, see "Deactivate z/VM Users" on page 118.
   b. Select the guest or guests that need relocation, and run the "Generate Disk Storage Map" action on them. For more information, see "Generate Disk Storage Map" on page 123.

© Copyright IBM Corp. 2007, 2018

c. Select all the relevant DASD Volume, and run the "Vary Offline" action on them. For more information about the "Vary Offline" action, see "Vary offline" on page 155.

3. In the Target z/VM System, complete the following steps:
   a. Select the guest or guests that need relocation, and run the "Generate DASD Map" action. For more information, see "Generate Disk Storage Map" on page 123.
   b. Select all the relevant DASD Volumes, and run the "Vary Online" action. For more information about the "Vary Online" action, see "Vary online" on page 155.
   c. Select the guest or guests that need relocation, and run the "Active" action. For more information about the "Activate" action, see "Activate" on page 110.

Appendix B. Sample CSV file for importing guest attributes

The following sample is a .csv file that you can use to import guest attributes. The sample assumes that there are two custom attributes defined “Application” and “Functionality”, and that each have the possible values preconfigured. Also existing are the Site Defined Groups (SDG): " **SDG1** ", " **SDG2** " and " **SDG3** ".

```
----------------- Start of CSV File -----------------------
intr_username,intr_system,intr_project, intr_SDG ,attr_Application,
attr_Functionality,attr_Importance
DEMO100,CSLZVM,Proj1, SDG1 ,App1,Func1,High
DEMO101,CSLZVM,Proj2, SDG2 ,App2,Func2,Medium
DEMO102,CSLZVM,Proj3, SDG3 ,App3,Func3,Low
------------------ End of CSV File -----------------------
```

© Copyright IBM Corp. 2007, 2018

Appendix C. IBM Wave shell script executor variables

When executing shell scripts using IBM Wave, the script is processed by a "script executor". The script executor sets up several variables that can be referenced from within the executed script.

```
WV_SCRIPT_NAME - The name of the shell script to be executed
WV_GUI_USER_VERSION - The version of the GUI client which is
executing the script
WV_GUI_USER_NAME - The IBM Wave User name logged into the
GUI client which is executing the script
WV_GUI_IP - The IP address of the workstation running the GUI
client which is executing the script
WV_SERVER_HOST - The NFS Host from which the executor will
attempt to mount the script
WV_SCRIPT_MOUNT_POINT - The mount point temporarily created
on the z/VM Guest on which the IBM Wave script repository is mounted
from the NFS server.
WV_LINUX_USER_NAME - The z/VM Guest name
WV_LINUX_HOST_NAME - The Linux hostname (fully qualified)
WV_LINUX_HOST_NAME1 - The Linux Hostname (without domain)
WV_LINUX_HOST_INTERFACE_LIST - The list of network interfaces
of the Linux Guest
```

For each network interface (<x> signifies the ETH number:

- `WV_LINUX_HOST_<x>_IP` - The IP Address of the interface.
- `WV_LINUX_HOST_<x>_BROADCAST` - The broadcast address of the interface.
- `WV_LINUX_HOST_<x>_MASK` - The netmask of the interface.
- `WV_LINUX_HOST_<x>_MTU` - The MTU of the interface.

© Copyright IBM Corp. 2007, 2018

Appendix D. IBM Wave by-id error resolution

IBM Wave does not support Linux `fstab` and `zipl/grub` disk location paths listed in by-id format. IBM Wave issues the following error when it detects the `by-id` format in use on a Linux guest.

```
2013-11-04 15:37:35 Error - IBM-WAVE does not support Linux guests with
UID pointers in the zipl/grub. ZIPL/Grub specifies root UID pointer on the
following line:

zipl: line 19: /dev/disk/by-id/ccw-IBM.750000000XG921.9241.1d.
00004ee300005ca30000000000000000-part1

grub: line 26: /dev/disk/by-id/ccw-IBM.750000000XG921.9241.1d.
00004ee300005ca30000000000000000-part1
```

To resolve this error, the administrator must modify the `fstab` and the `zipl` files to reflect device locations `by-path` instead of `by-id`. The `by-path` device locations are in the `/dev/disk/by-path/` directory.

After you locate the `by-path` file names in the `/dev/disk/by-path/` directory, use the following steps:

1. Modify `fstab`.
   a. `<editor> /etc/fstab`
   b. Replace all disk device location `by-id` paths with `by-path` locations.
   c. Save and exit.
2. Modify `zipl`.
   a. `<editor> /etc/zipl.conf` or `/boot/zipl/config` on SLES12.
   b. Replace all disk device location by-id paths with `by-path` locations.
   c. Save and exit.
   d. Recompile `zipl` by running the **`zipl`** command from the shell prompt.

When `fstab` and `zipl` are updated correctly, IBM Wave recognizes the device paths for each of the DASD disks (and the error does not reoccur).

© Copyright IBM Corp. 2007, 2018

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho,
Chuo-ku Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corp. 2007, 2018

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation Site Counsel
2455 South Road Poughkeepsie, NY
12601-5400 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information may contain sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see the IBM Online Privacy Policy at http://www.ibm.com/privacy and the IBM Online Privacy Statement at http://www.ibm.com/privacy/details, in particular the section entitled “Cookies, Web Beacons and Other Technologies”, and the IBM Software Products and Software-as-a-Service Privacy Statement at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com® 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 IBM trademarks is available on the web at IBM copyright and trademark information - United States(www.ibm.com/legal/us/en/copytrade.shtml).

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

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Accessibility

Accessible publications for IBM Wave

The publications for IBM Wave are available in IBM Knowledge Center and Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties using any of the information, use the Reader Comment Form in this publication to send a request. See “How to send your comments to IBM” on page xvii, and be sure to include the publication number and title.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Index

Special characters

Numerics

A

B

C

© Copyright IBM Corp. 2007, 2018

D

E

F

G

H

I

J

L

M

N

O

P

R

S

T

U

V

W

Z

Product Number: 5648-AE1

Printed in USA

SC27-6119-14