Redbooks

ibm.com/redbooks

# IBM Wave for z/VM Installation, Implementation, and Exploitation

Lydia Parziale

Edi Lopes Alves

George Azevedo

Igor Bajo Dutra

Roland Trauner

Cloud

z Systems

IBM

Redbooks

International Technical Support Organization

# IBM Wave for z/VM Installation, Implementation, and Exploitation

December 2015

SG24-8192-01

**Note:** Before using this information and the product it supports, read the information in “Notices” on page vii.

**Second Edition (December 2015)**

This edition applies to IBM Wave for z/VM v1.2.0 (PID number 5648-AE1 1.1) Version 6, Release 3 of z/VM (product number 5741-A07), SUSE Enterprise Linux Server 11 SP4, and Red Hat Enterprise Linux release 6.5.

# Contents

# 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.*

**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.

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 measurements 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.

This information contains 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 contains 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.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

| | | |
|---|---|---|
| AIX® | IBM z Systems™ | WebSphere® |
| DB2® | IBM z13™ | z Systems™ |
| DirMaint™ | PR/SM™ | z/OS® |
| ECKD™ | RACF® | z/VM® |
| HiperSockets™ | Redbooks® | z10™ |
| IBM® | Redbooks (logo) ® | z13™ |
| IBM z™ | System z10® | zEnterprise® |

The following terms are trademarks of other companies:

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

Microsoft, Windows, and the Windows logo 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.

UNIX is a registered trademark of The Open Group in the United States and other countries.



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

# Find and read thousands of IBM Redbooks publications

- Search, bookmark, save and organize favorites
- Get up-to-the-minute Redbooks news and announcements
- Link to the latest Redbooks blogs and videos

**Get the latest version of the Redbooks Mobile App**

iOS **Download Now** Android

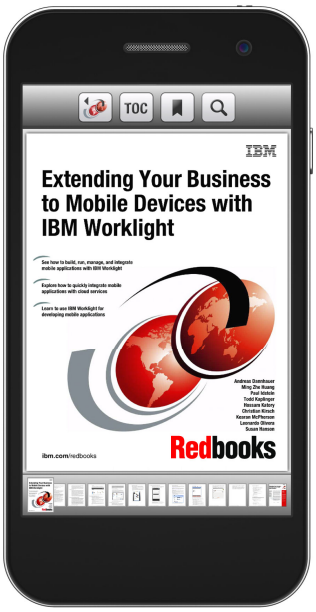


# Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!


**ibm.com/Redbooks**
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

# Preface

IBM® Wave for z/VM® (IBM Wave) is a virtualization management solution for IBM z/VM and Linux on z Systems™ systems. This virtualization management software provides a simplified and cost-effective way for companies to harness the consolidation capabilities of the IBM z™ Systems platform and its ability to host the workloads of tens of thousands of commodity servers. IBM Wave is a complete management solution for z Systems based virtual server farms.

This IBM Redbooks® publication provides a guide to understanding IBM Wave by providing information about the IBM Wave architecture and how it fits into the cloud. This publication also provides a planning and design guide that is based on common scenarios. This publication also provides installation and configuration task information and how to manage and operate the environment.

The intended audience for this publication is IT Architects who are responsible for planning their IBM Wave environments and IT Specialists who are responsible for implementing them.

**Note:** The terms *virtual machine*, *virtual server*, *guest*, and *z/VM user* all refer to a virtual machine. These terms are used interchangeably throughout this publication.

## Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

**Lydia Parziale** is a Project Leader for the ITSO team in Poughkeepsie, New York, with domestic and international experience in technology management, including software development, project leadership, and strategic planning. Her areas of expertise include Linux on z Systems and database management technologies. Lydia is a Certified IT Specialist with an MBA in Technology Management and has been employed by IBM over 25 years in various technology areas.

**Edi Lopes Alves** is a Senior IT Specialist in the IBM z Systems™ portfolio for the GTS team in São Paulo, Brazil. She has +25 years of experience working as a z/VM and Linux on z Systems specialist. Edi is a certified z Systems Technical Specialist with a master degree in e-Business at ESPM/São Paulo. She supports z/VM and Linux on z Systems systems at Banco do Brazil and its cloud initiatives. She has supported IBM Global Accounts (IGA) for several years as part of IBM Green and IBM Blue Harmony projects. Edi co-authored three IBM Redbooks publications: *Using z/VM v 6.2 Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8039, *Linux on IBM System z: Performance Measurement and Tuning*, SG24-6926, and *Introduction to the New Mainframe: z/VM Basics*, SG24-7316.

**George Azevedo** is a Team Manager for mainframe operational systems at Banco do Brazil in Brazil. He holds a degree in Accounting, and has 25 years of experience in the IT Infrastructure field. George has worked at Banco do Brazil for 27+ years. His areas of expertise include IBM AIX®, Solaris, SCO UNIX, Linux for all architectures, IBM z/OS®, and z/VM. He has led a large project about planning and implementing enhancements of the Banco do Brazil z/VM environment since 2011.

**Igor Bajo Dutra** is an Operacional Systems Specialist for Banco do Brazil in Brazil. He holds a degree in Computer Science and is a Certified IT Specialist in Computer Forensics. His areas of expertise include programming in COBOL and Natural language for Mainframes, z/VM, and Linux on z Systems systems.

**Roland Trauner** is a Senior IT specialist in Germany working in the global IBM Wave Tiger Team. The Wave Tiger Team is a pre-sales support team that is focused on z/VM and Linux on z Systems environments running together with IBM Wave. Roland worked as a systems programmer and in technical support and other technical jobs on the mainframe. He also worked in the ITSO as a project manager, where he wrote and produced several books. Roland also worked as a brand manager for z Systems hardware. He joined IBM in 1979.

Thanks to the following people for their contributions to this project:

Rich Conway
**International Technical Support Organization, Poughkeepsie Center**

Evren O. (Oz) Baran
**IBM Wave for z/VM development team**

Don E. Kallberg
**IBM Wave for z/VM development team**

Daniel S. Bokser
**IBM Wave for z/VM development team**

Erin Farr and the IBM Wave for z/VM development team
**IBM**

Eric Marins
**IT Architect on IBM Brazil**

Ana Paula Paiva de Oliveira
**IT System Specialist on IBM Brazil**

Thanks to the authors of the previous editions of this book.

- Authors of the first edition, IBM Wave for z/VM Installation, Implementation, and Exploitation, published in June 2014, were:

  Lydia Parziale, Susan Greenlee, Ernest Horn, Richard Lewis, Sean McBride, Steve McGarril, and Gene Ong

## Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

**ibm.com**/redbooks/residencies.html

## Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

# Summary of changes

This section describes the technical changes that are made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-8192-01
for IBM Wave for z/VM Installation, Implementation, and Exploitation
as created or updated on December 8, 2015.

## December 2015, Second Edition

We revised the order of several chapters or parts to produce the following flow:

- Architecture
- Planning
- Installing
- Customizing
- Using

This revision includes the following new and changed information.

### New information

- You can do a bare metal installation from Red Hat Enterprise Linux Servers.
- There is more information about IBM DirMaint™, IBM RACF®, single-system image (SSI), and emulated device (EDEV).
- There are two new kinds of reports.
- There is a new chapter (Chapter 8, “IBM Wave for z/VM parameters” on page 261) that describes IBM Wave / BTS parameters that might influence performance and resource usage.
- We added an appendix (Appendix A, “Miscellaneous” on page 269) with miscellaneous information. This appendix also contains some flow charts.

### Changed information

In the Change IBM Wave User Preferences window, there is a new External SSH Program Parameters field to configure the parameters of the SSH program that is associated with IBM Wave.

# Part 1

# Architecture and installation

This part describes IBM Wave for z/VM (IBM Wave) and its business value. It also describes the basic architecture of the IBM Wave environment and the lab environment that was used in the creation of this publication.

This part also describes the planning and design decisions that must be made before the installation is done, and the basic installation tasks.

The part includes the following chapters:

1

# Introduction to IBM Wave for z/VM

This chapter provides an introduction to IBM Wave for z/VM (IBM Wave) and includes the following topics:

- What is IBM Wave
- Benefits of using IBM Wave

## 1.1 What is IBM Wave

IBM Wave is a provisioning and productivity management solution that was originally acquired with the IBM acquisition of CSL International. This product provides a graphical user interface (GUI) tool that simplifies the management and administration of both z/VM and Linux environments on IBM z Systems systems. It also allows some basic monitoring, and reporting that allows also auditing and documentation of the environment. The scope of IBM Wave can span multiple LPARs, single-system image (SSI) clusters, and servers.

IBM Wave provides a visualization that represents virtual resources and their relationships. It provides a customizable graphical or tabular view to help administrators understand the status of the entire system or a subset, such as system, network, or storage topology in a single view. The status of servers, networks and network devices, and storage is automatically detected by IBM Wave. Changes in resources and their relationships are monitored and reflected in displays. Many of the IBM Wave displays feature advanced filters, tagging, layout, and layer selection to provide customized views.

This simplified GUI can be used to manage virtual servers and resources that otherwise must be performed with z/VM system commands by a skilled z/VM systems programmer. With IBM Wave, the user can point and click to complete the following tasks:

- Perform basic z/VM guest actions, such as activate, deactivate, recycle, pause, and resume
- Provision virtual resources, such as z/VM guests, network, and storage
- Capture and clone virtual servers across LPARs
- Create and configure vSwitch and guest LANs
- Connect virtual servers to virtual networks
- Install Linux on a virtual guest
- Relocate virtual guests with live guest relocation
- Display and monitor page and SPOOL and add and remove disks
- Provision and track storage or free OSA and IBM HiperSockets™ devices that use Device Pools
- Manage storage by using the Manage Storage wizard that enables management at the z/VM level, such as dedicating devices, adding minidisks, and enabling management on the Linux level, such as the creation and expansion of LVM volume groups, regular partitions, and logical volumes, and adding FCP attached storage to z/VM guests

## 1.2 Benefits of using IBM Wave

IBM Wave has many benefits for today's z/VM environment. It simplifies management and administration of a z/VM environment by abstracting z/VM resources and providing a GUI to automate z/VM functions. With a simplified GUI and its point and click functions, there is much less need to understand extensive z/VM command syntax. This is true for Linux administrators that are assigned projects. Using IBM Wave instead of z/VM or Linux command-line syntax also prevents typographical errors and errors. IBM Wave mostly verifies the feasibility of an action and either does not allow it by not offering the option or by visualization. IBM Wave then indicates by a color change to red that the action is not possible.

A unique capability of IBM Wave is its ability to create a project, which is a set of virtual resources that consist of any combination of servers, networks, and storage devices across the enterprise. The definition of projects can enable different constituencies, such as operations, systems programming, application development, project management, or users to exercise their authority to manage appropriate aspects of the z/VM environment.

The scope and permissions a Linux administrator (which is also called a regular IBM Wave user) has within a project can be uniquely assigned. *Scope* refers to what a Linux administrator can view or see. *Permissions* refers to what actions a Linux administrator can perform within the assigned scope. A Linux administrator can have different scopes and permissions for different projects. Likewise, different Linux administrators can have different scopes and permissions for the same project.

The delegation of administrative capability can also be assigned based on defined roles, such as project administration, network administration, storage administration, and operator. The flexibility of setting up Linux administrators and assigning unique projects, scopes, and permissions, and the use of defined roles makes it feasible to distribute administration responsibility to specific individuals across organizations that are facilitating project self-administration. IBM Wave provides an environment that allows a level of autonomy and self-service without extensive z/VM skills.

IBM Wave improves the ease of use of z/VM and provides the following other functions and flexibility to manage a virtualized environment:

- Most actions can be carried out against multiple objects, whether they are on a single or multiple LPARs.
- Intelligent icons help the user understand the status of the system and its resources. For example, an icon can indicate whether a virtual guest is running, whether it is running a Linux distribution from SUSE® or Red Hat, or whether it is a Conversational Monitoring System (CMS) guest, a system service machine, or an IBM Wave internal virtual server.
- Definition and use of function types ensures ordered activation and deactivation of z/VM guests. For example, a database server can be started followed by a transaction server and a web server with deactivation occurring in the reverse order.
- Custom attributes and their associated values can be defined and associated with virtual guests. Multiple attributes can be assigned to guests and then used for grouping and filtering in any combination to act on a set of guests. For example, an attribute that is called Server Use might be defined with possible values Production, QA, and Development. Another attribute that is called Server Type might be defined with possible values Database, Transaction, and Web.
- Ability to broadcast a message to a set of Linux users. This can be finely filtered with the use of custom attributes.
- Intelligent Active Notes (IANs) allow free-form text to be associated with an object and displayed when you hover a cursor over text. Starting an action on the object causes the IAN to open.
- At-a-glance views of overall system storage utilization with the ability to drill down to the volume level for site-level administrators and for Linux administrators views of storage groups to which they are uniquely permitted.
- Simplified mechanism of storage assignment and storage management, which allows the IBM Wave user to interact with DASD Volumes, DASD Groups, and DirMaint regions.

- Performance view of all managed systems includes current CPU utilization, page, and spool space and virtual to real storage ratios. Access to the Linux guest level is also supported.
- Generate concise reports for selected objects with the ability to use complex filters and customization.

2

# IBM Wave for z/VM architecture

This chapter provides an overview of the basic IBM Wave for z/VM (IBM Wave) architecture and how it fits into your enterprise, whether small or large. It also describes the foundational components of IBM Wave within each tier of the architecture.

## 2.1 IBM Wave basic architecture

IBM Wave implements the following three-tier architecture, as shown in Figure 2-1:

- A graphical user interface (GUI) client
- The WAVESRV Linux guest server
- The Target Virtualization Platform (TVP) application programming interface (API)

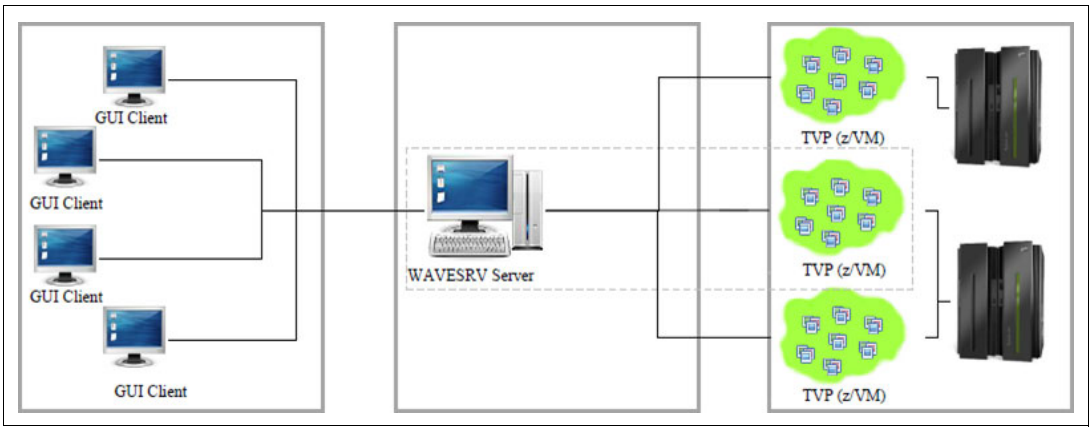


*Figure 2-1 IBM Wave three-tier architecture*

This section describes each tier and how they communicate with each other. A high-level overview of this communication is shown in Figure 2-2.

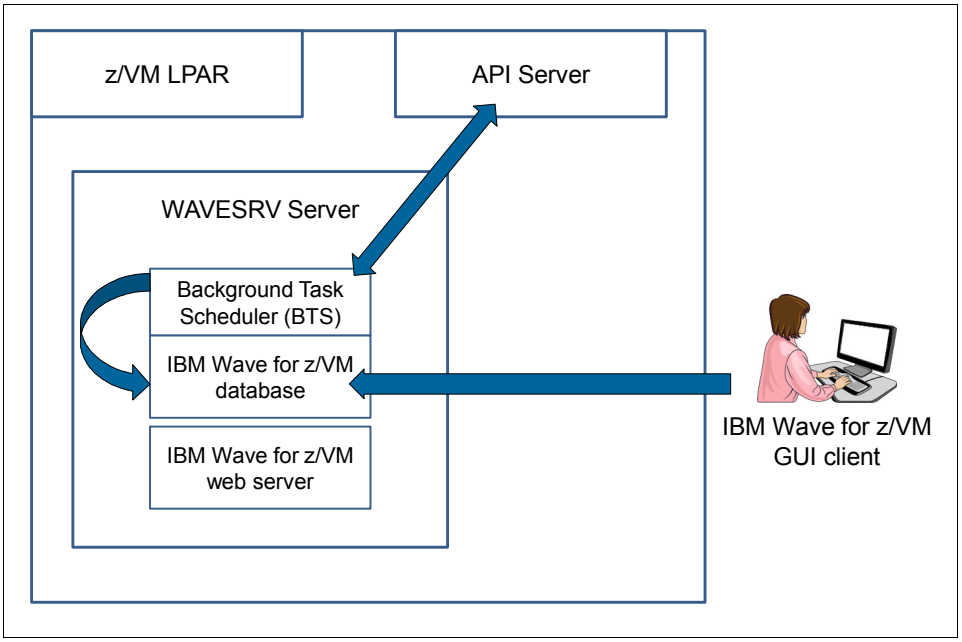


*Figure 2-2 Communication between tiers of IBM Wave*

### 2.1.1 GUI client

The GUI client controls the running, attributes, and behavior of the Linux virtual servers and the rest of the managed z/VM resources according to the scope and permissions of the IBM Wave. It provides a graphical interpretation of the knowledge base that the user can use to interact with the TVP by using point-and-click and drag operations. It is responsible for the facilitation of all interaction and viewing of the objects that are managed by IBM Wave. The GUI engine also facilitates all of the interaction with the z/VM LPAR.

The GUI client is a Java application that is deployed by using Java WebStart technology and is supported on Microsoft Windows. The Apache web server on the WAVESRV Linux guest server contains a link that downloads and installs the IBM Wave GUI client on the workstation. For MS Windows users, it also creates a link on the desktop and in the Start menu.

After the application is installed, starting the application can be done by double-clicking the icon or shortcut. The Java Web Start facility includes an automatic search for updates. This means that after applying IBM Wave maintenance (which updates the GUI client), it is not necessary to reinstall IBM Wave. The application is automatically updated the next time that it is started.

The IBM Wave GUI engine is in constant contact with the Background Task Scheduler (BTS) on the WAVESRV Linux guest server to ensure that the GUI display represents the most current state of the environment. All communications between the GUI, BTS, and IBM Wave service machines are encrypted.

### 2.1.2 WAVESRV Linux guest server

The WAVESRV Linux guest server runs within a z/VM LPAR or as a native Linux for z Systems partition on IBM PR/SM™. The IBM Wave software level that is used in the project for this book is IBM-Wave-1.2.0.00.5 (1.2 fix pack 0.5). It requires Red Hat Enterprise Linux Version 6 Servers, or SUSE Linux Enterprise Server 11. The project uses Red Hat Enterprise Linux 6.5 Servers and SUSE Linux Enterprise Server 11 SP4. IBM Wave can manage Red Hat Enterprise Linux 7 Servers and SUSE Linux Enterprise Server 12 guests, but, at the time of writing, IBM Wave itself must be installed on these levels. IBM intends to support the installation and running of the IBM Wave server (IBM WAVESRV) on RHEL 7 and SLES 12 distributions in a future service update.

The WAVESRV Linux guest server is a virtual server and contains the foundational components for IBM Wave: the application knowledge database (or knowledge base), a web server, and the BTS, as shown in Figure 2-2 on page 8.

#### Application knowledge base

The application knowledge database that is used with IBM Wave is a MySQL database. It is on the WAVESRV Linux guest server and contains information about all z/VM elements that are managed by IBM Wave and the IBM Wave Regular Log, the Wave Users, and so on.

The following information is stored in the IBM Wave application knowledge base:

- Objects that are managed by IBM Wave: Central processor complexes (CPCs), z/VM systems, Linux for z Systems virtual servers and their z/VM hosts, Virtual Networks and Virtual Networks connections' definitions, Prototypes, IBM Wave logs, IBM Wave users, Projects, Storage Groups, Disks, and their attributes.
- For each unique directory that IBM Wave detects, a directory record is saved in the knowledge base, and any z/VM system that uses that directory is associated with the directory record automatically. IBM Wave monitors the directory periodically; therefore, if the unique identification of the directory changes, these changes are reflected in the knowledge base and the GUI clients.
- When a z/VM system is added to IBM Wave management, the auto-detection process starts automatically. The auto-detect mechanism accesses the z/VM LPAR and queries various aspects of the z/VM system. The results of these queries populate the IBM Wave knowledge base.
- IBM Wave automatically detects z/VM Accounts that are used in managed z/VM systems and stores information about them in the IBM Wave knowledge base.

The IBM Wave knowledge base is on the `/var` file system. Configure this file system as a logical volume that is managed by the Logical Volume Manager (LVM).

### Background Task Scheduler

The BTS updates the knowledge base with live data from all managed and active z/VM LPARs. The BTS, along with the z/VM API server, provides information from the knowledge base to the GUI client. Therefore, even if changes are made outside of IBM Wave (for example, by a z/VM systems programmer that uses a 3270 emulator), they are reflected in the GUI client.

The BTS handles requests by using worker threads through a proprietary communication protocol that is based on TCP/IP.

### Web server

The WAVESRV Linux guest server also runs a minimal Apache web server, which is installed to enable the download of the GUI client by using the Java Web Start application.

## 2.1.3 Target Virtualization Platform

The Target Virtualization Platform (TVP) is a mediation layer that provides the interface into the managed environment. IBM Wave uses the system management application programming interface (SMAPI) to mediate requests from the BTS and the GUI Client.

The TVP represents the hypervisor that hosts the virtual guests. The BTS and the GUI clients use the TVP API to query and perform changes to the TVP and hosted virtual guests.

Specific functions that are run on the z/VM system use the following service machines:

- The Short Service Machine runs various compiled REXX EXECs to interact with the z/VM environment (WAVEWRKS).
- The Long Service Machine provides another thread of execution to run longer scripts or runs some directory manager commands (WAVEWRKL).
- This Cross-System Cloning Service Machine is used to stream minidisks from a source z/VM system to a target z/VM system during cloning actions (WAVEWRKC).

The names of the service machines that are used here are default names that can and should be configured with different names within IBM Wave. Different names are used in multi-system environments and single-system image (SSI) cluster.

As part of the auto-detection process, when a new z/VM system is added to IBM Wave management, these three service machines are created and started on the z/VM system automatically.

IBM Wave performs various queries and runs commands on the TVP by using these service machines and SMAPI. These interactions are performed through the BTS. The communication between the service machines and the BTS can optionally be encrypted by using SSL.

The service machines run in any z/VM LPAR that is managed by IBM Wave and must be kept running. IBM Wave monitors these service machines and issues alerts if there are any errors.

### 2.1.4 IBM Wave command-line interface

IBM Wave provides a command-line interface (CLI) to run specific IBM Wave actions by using Windows batch files or Linux shell scripts. This CLI is not required for IBM Wave, but is available if it is wanted.

In the version of IBM Wave that is used in this book (IBM Wave V1.2 Fix Pack 0.5), the following command functions are supported by the CLI:

- Activate a specific z/VM guest
- Deactivate a specific z/VM guest
- Display details about a specific z/VM guest
- Retrieve the status of a BTS workunit ID
- Signals to the BTS that the activation of the specific z/VM guest is complete

3

# Planning and design

This chapter provides guidance about planning your IBM Wave for z/VM (IBM Wave) installation. It also provides some design guidelines and worksheets with examples that help with the installation of IBM Wave in your own environment.

This chapter includes the following topics:

- Planning for installation
- Preparing your own worksheets

# 3.1 Planning for installation

Consider the sizing for memory and your file systems, including the file systems that are used for logging and backup, before the installation is done. Also, plan and verify the appropriate levels of software on z Systems systems and on your workstations. In a multi-z/VM environment, you must also consider the location of the WAVESRV Linux server. Only one VM is required to manage many z/VM systems. So, plan on which system you will install and run `WAVESRV`.

In our environment, we installed IBM Wave 1.2 Fix Pack 0.5 (`IBM-Wave-1.2.0.00.5`). So, all planning in this book is based on this release and level.

A flowchart for planning can be found in Figure A-4 on page 286.

## 3.1.1 Hardware requirements

IBM Wave requires one of the following IBM z Systems servers:

- IBM z13™
- IBM zEnterprise® BC12 or EC12
- zEnterprise 114 or 196
- IBM System z10® BC or EC

In our environment, we installed IBM Wave on an IBM z13.

## 3.1.2 Server software requirements

IBM Wave requires one of the following operating systems:

- z/VM V5.4
- z/VM V6.2
- z/VM V6.3

For our environment, we ran IBM Wave on z/VM V6.3 on a four-system single-system image (SSI) cluster without RACF and s on a stand-alone machine with RACF.

IBM Wave also requires SMAPI, DirMaint, and TCP/IP to be configured and active. Telnet access must be enabled and an FTP server must be running. Although it is not a requirement, the Performance Toolkit for VM should be configured and active.

## 3.1.3 Workstation requirements

The following prerequisites must be met on every workstation running the IBM Wave GUI client:

- Microsoft Windows with Java 1.7 installed (64-bit installation is recommended.)
- Firefox or Internet Explorer
- PuTTY or an equivalent telnet/SSH client installed

In our environment, we found that Java 1.6 and Java 1.8 worked. (The last release we used was 1.8.51.) Any workstation that can run those Java client versions can run the IBM Wave GUI, which includes Linux and other workstations. Also, other browsers such as Chrome work well, but are not supported.

### 3.1.4 Wave server requirements

The WAVESRV Linux guest server runs within a z/VM LPAR or as a native Linux for z Systems partition on PR/SM. The IBM Wave software level that we use in our environment is IBM-Wave-1.2.0.00.5. It requires Red Hat Enterprise Linux 6 Servers, or SUSE Linux Enterprise Server 11. We used Version 6.5 with Red Hat Enterprise Linux and Version 11 SP4 for SUSE Linux Enterprise Server. Even though IBM Wave can manage RHEL 7 and SUSE's V12 guests, it must be installed on these levels.

If you install IBM Wave for the first time, you must install the base IBM Wave RPM (`IBM-Wave-1.10-1.s390x.rpm`) and then install the test fixes for IBM Wave V1.2. If you already installed IBM Wave and you are on Version 1.1 with Fix Pack 10, then you can install IBM Wave V1.2 on top. For more information, see the readme file that comes with the software and "IBM Wave for z/VM fix pack installation" on page 270.

The Linux system can run as a z/VM guest or an LPAR. It requires Java SE Runtime 1.7 or later, MySQL (minimum Version 12.22 dist 4.0.8) or later and Apache (HTTPd) to be installed.

For SUSE, mysql (5.0.67), mysql-MAX (5.0.67), mysql-client (5.0.67), apache-prefork, and nfs-kernel-server (provides nfa-utils) must be installed

For RHEL: mysql-5.0.45-7.el5, mysql-server-5.0.45-7.el5, httpd-2.2.3-11.el5_1.3, nfs-utils-1.0.9-33.el5, and nfs-utils-lib-1.0.8-7.2.z2 must be installed.

For an enabled firewall, you must allow the following services:

- SSH (TCP port 22)
- Web server (TCP port 80)
- Web server SSL (TCP port 443)
- MySQL (TCP port 3306)
- NFS (See your firewall manual on how to configure it to allow NFS.)
- IBM Wave communication between the IBM Wave GUI and BTS (TCP Port 3300).

The installation includes the following minimum requirements:

- 1.5 GB RAM (2 GB is preferable.)
- /: (root) 3 GB (4 GB is preferable.)
- `/boot`: 100 MB
- `/ var`: 3 GB

By default, the logs are stored in `/var`, so you should configure this file system as a logical volume under logical volume management (LVM) so that it can be extended when needed. This setting is necessary because the IBM Wave Knowledge Base is on this file system and it might need to be expanded later, depending on various factors, such as the number of historical data sets to be kept and the number of z/VM systems added to IBM Wave management. Make `/var/log/wave` a separate file system because if you run IBM Wave in debug mode, the amount of data that stored in the log file can be large.

The WAVESRV Linux server can be implemented as a virtual server within a z/VM LPAR. Set the WAVESRV virtual machine (VM) storage size at logon to 2 GB.

*Example 3-1 Sample directory entry with comments for the WAVESRV Linux server*

```
*************************************************************
USER WAVESRV <Password Here> 2G 4G GC
CPU 00
IPL 150             - IPL from the "/" mini disk -
MACHINE ESA 4
```

```
OPTION QUICKDSP     - careful - set absolute share -
CONSOLE 0009 3215
NICDEF 0800 TYPE QDIO LAN SYSTEM <GLAN/VSWITCH Name>
SPOOL 000C 3505 A
SPOOL 000D 3525 A
SPOOL 000E 1403 A
LINK MAINT 0190 0190 RR     - Links to Maint can be removed after installation -
LINK MAINT 019D 019D RR     - to make WAVESRV eligible for VMRELOCATE          -
LINK MAINT 019E 019E RR     - for SSI installations only                       -
MDisk 0191 3390 1 3 <DASD Volume Name>    - This is a CMS minidisk
MDisk 0150 3390 1 3000 <DASD Volume Name> - minidisk will be used for "/"
MDisk 0151 3390 1 4500 <DASD Volume Name> - minidisk will be used for "/var"
MDisk 0152 3390 1 200 <DASD Volume Name>  - minidisk will be used for swap
************************************************************
```

Because VMs are added to the dispatch list immediately when setting the quick dispatch option, be careful when setting **OPTION QUICKDSP**. In addition, set the absolute share (the percentage of system resources for the user) of WAVESRV to 1 or 2%.

## IBM Wave server backup planning

If you start backups in IBM Wave, the backup location in IBM Wave is in the `/usr/wave/DBBackup` directory. Backups require some disk space, and every time that you install a new IBM Wave test fix RPM, it clears the `/usr/wave` directory, which is why you should move a backup to another location immediately after the backup. Also, make sure that you do not have too many administrators running the backup command. There is no further step or dialog box in the administration window; after you select **Backup IBM Wave Database** from the GUI, it starts immediately.

Details about the Backup function are described in 5.10, “Backup option for IBM Wave” on page 153

## Preparation for SSI Installations

In our environment, we have a four-member SSI cluster. We define the first member (ITSOZVM1) to host the IBM Wave Server (WAVESRV). The IBM Wave Server can be relocated, but exists in one member of the SSI cluster. In fact, it mostly exists in only one z/VM server in the whole environment because you want to manage all z/VM systems from that IBM Wave Server.

The IBM Wave Service Machines require a different setup. IBM Wave Service Machines must run in every z/VM server that must be managed by IBM Wave. For more information, see 3.1.5, “IBM Wave Service Machine requirements” on page 19. There are two possible setup scenarios:

- The first scenario uses the same name and definition for the service machines on each member of the SSI cluster. This configuration requires many manual definitions. You must create identities for each service machine on each member of the cluster.
- The second scenario defines different IBM Wave Service Machines with different names on every member of the SSI cluster. The advantage is that you do not need to define any identities; IBM Wave defines all the service machines during the Auto Detect step.

The first scenario looks like Figure 3-1.

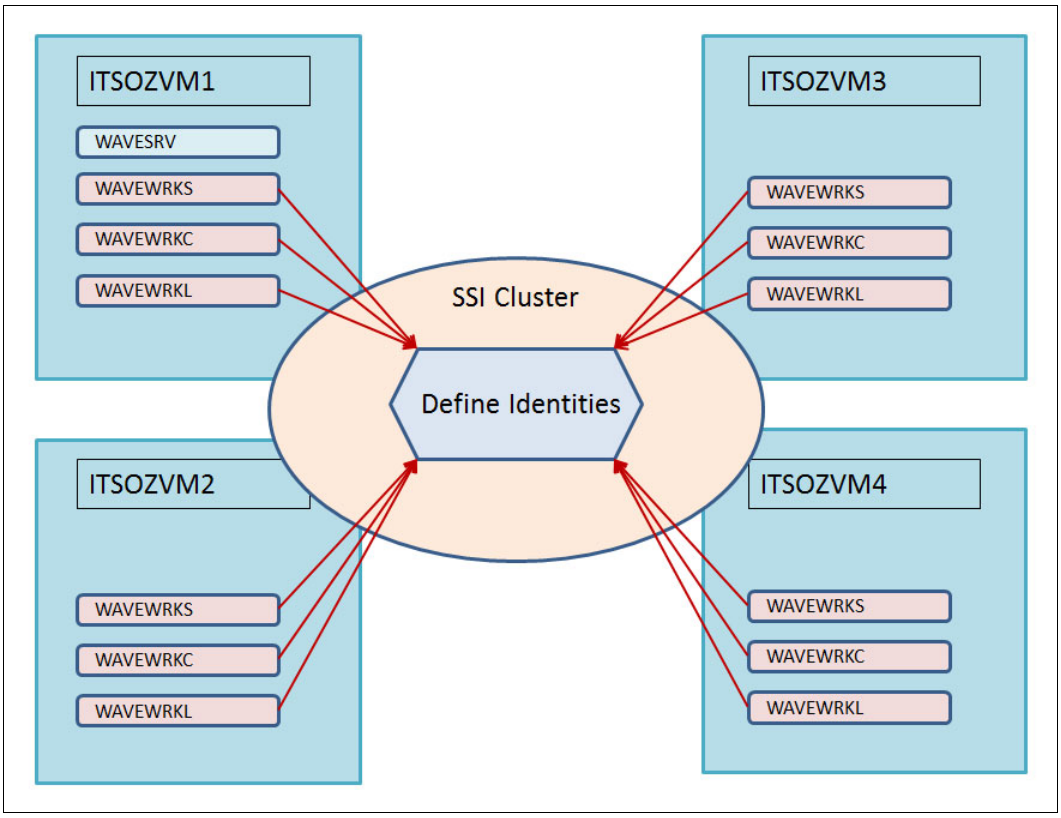


*Figure 3-1 IBM Wave SSI Cluster - Service Machines as Identities*

The second scenario looks like Figure 3-2.

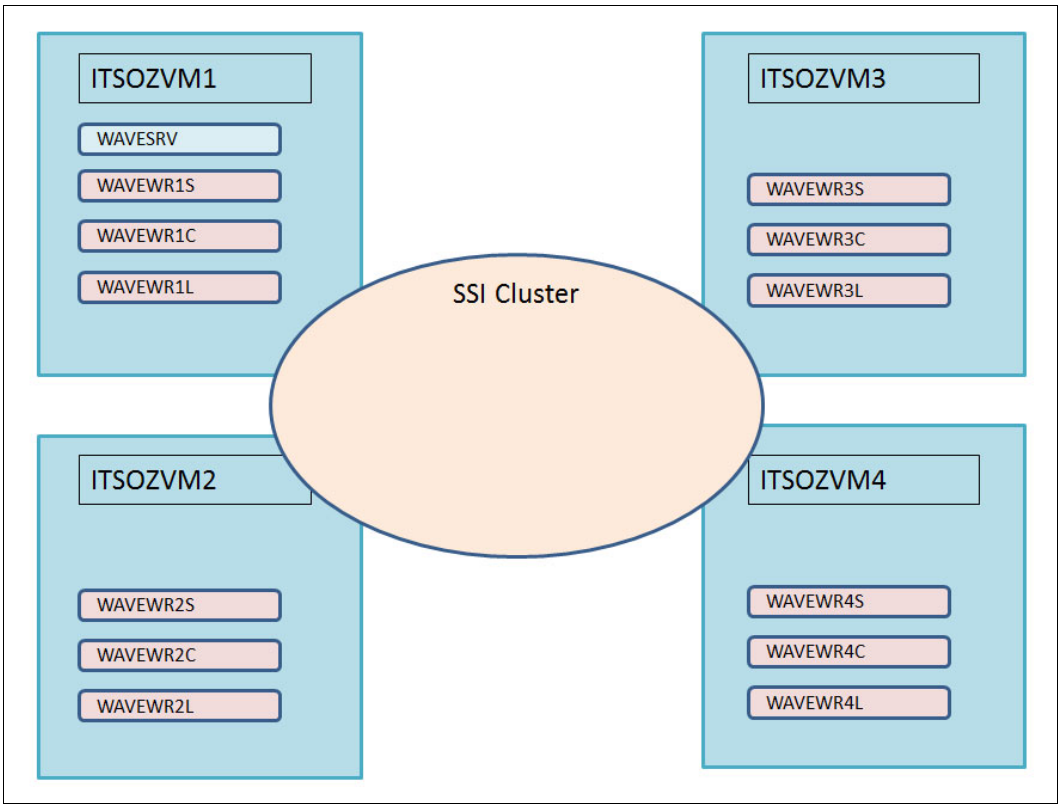


*Figure 3-2 IBM Wave SSI Cluster - service machines with different names*

For an IBM Wave Cluster with Identities for the IBM Wave Service Machines, described in Figure 3-1 on page 17, you must follow the definitions that are described in 4.1.7, “Single-system image definitions” on page 39.

For an IBM Wave Cluster without Identities for the IBM Wave Service Machines, described in Figure 3-2, you must plan the names for the IBM Wave Service Machines. In our environment, we use the names WAVEWR1S, WAVEWR1C, and WAVEWR1L for the first member (ITSOZVM1). Then, we used WAVEWR2S, WAVEWR2C, and WAVEWR2L for the second member (ITSOZVM2), and so on. Because we use different names for the service machines, we might need to authorize and define them in DirMaint, TCP/IP, Performance Toolkit, and in your external security management system, such as RACF.

## Preparation for installation with RACF or other security facilities

If the environment is not protected by RACF or another security facility, IBM Wave creates all IBM Wave Service Machines in the Auto Detect step automatically. In a “protected” environment, those service machines must be predefined and authorized. IBM Wave then detects them in the Auto Detect step and skips the creation of the service machines. User IDs must be predefined and access to minidisks must be granted. For more information, see 4.1.2, “Preinstallation configuration tasks for z/VM” on page 28, 4.1.8, “RACF definitions” on page 47, and “Customization with other security managers” on page 282.

### 3.1.5 IBM Wave Service Machine requirements

The IBM Wave Service Machines are simple z/VM Conversational Monitoring System (CMS) service machines that run REXX scripts to run commands and obtain information for the IBM Wave Server. The IBM Wave Server requires three IBM Wave Service Machines on every managed VM server. Every service machine requires 64 - 128 MB of memory, 130 MB of dedicated disk space, and some shared disk space. The disk space requirement varies when configuring as dedicated or stand-alone, SSI shared identities, or shared directories. If you use dedicated service machines, then IBM Wave defines them for you by using the Auto Detect wizard when adding z/VM systems to manage. SSI or shared directory configurations require manual configuration steps before running the Auto Detect wizard.

Example 3-2, Example 3-3, and Example 3-4 on page 20 show sample directory entries for the IBM Wave Service Machines when working with a shared directory. A shared directory requires two more z/VM minidisks of about 65 MB each, which are linked by using **SYSAFFIN** statements.

*Example 3-2 Sample directory entry for the WAVEWRKS service machine working with a shared directory*

```
USER WAVEWRKS <Password Here> 64M 128M ABCDEFG
IPL CMS
MACHINE ESA 4
OPTION LINKOPAS DIAG88
CONSOLE 0009 3215
SPOOL 000C 3505 A
SPOOL 000D 3525 A
SPOOL 000E 1403 A
LINK MAINT 0190 0190 RR
LINK MAINT 019D 019D RR
LINK MAINT 019E 019E RR
LINK TCPIP 0592 0592 RR
LINK MAINT 0193 0193 RR
*CSLTAG01: WAVE-INTERNAL
```

*Example 3-3 Sample directory entry for the WAVEWRKL service machine working with a shared directory*

```
USER WAVEWRKL <Password Here> 64M 128M ABCDEFG
IPL CMS
MACHINE ESA 4
OPTION LINKOPAS DIAG88
CONSOLE 0009 3215
SPOOL 000C 3505 A
SPOOL 000D 3525 A
SPOOL 000E 1403 A
LINK MAINT 0190 0190 RR
LINK MAINT 019D 019D RR
LINK MAINT 019E 019E RR
LINK TCPIP 0592 0592 RR
LINK MAINT 0193 0193 RR
LINK WAVEWRKS 0191 0191 RR
LINK WAVEWRKS 0399 0399 RR
*CSLTAG01: WAVE-INTERNAL
```

*Example 3-4 Sample directory entry for the WAVEWRKC service machine working with a shared directory*

```
USER WAVEWRKC <Password Here> 64M 128M ABCDEFG
IPL CMS
MACHINE ESA 4
OPTION LINKOPAS LNKE DIAG88
CONSOLE 0009 3215
SPOOL 000C 3505 A
SPOOL 000D 3525 A
SPOOL 000E 1403 A
LINK MAINT 0190 0190 RR
LINK MAINT 019D 019D RR
LINK MAINT 019E 019E RR
LINK TCPIP 0592 0592 RR
LINK MAINT 0193 0193 RR
LINK WAVEWRKS 0191 0191 RR
LINK WAVEWRKS 0399 0399 RR
*CSLTAG01: WAVE-INTERNAL
```

### 3.1.6 IBM Wave Command-Line Interface preparation

The IBM Wave Command-Line Interface (CLI) is optional and not required for IBM Wave and the IBM Wave GUI to operate. However, if you want to use it, you must set up a method of sending the required files to the destination server or workstation. In our environment, we set up an anonymous FTP server in our WAVESRV server that allows access to those files.

## 3.2 Preparing your own worksheets

Throughout the process of installing and configuring IBM Wave, you need numerous pieces of information that are related to your z/VM environment and how you want to configure IBM Wave. The worksheets in this section are provided to help you record this information before you start the installation process to allow for easy reference later.

#### Linux guest for use as an IBM Wave server

A major prerequisite for installing IBM Wave is to create the Linux guest that serves as the IBM Wave Server. Complete the information that is shown in Table 3-1 so that you can access it by using an SSH application and install the IBM Wave RPM.

*Table 3-1 IBM Waver Server information*

| Distribution | IP address | Root password |
|---|---|---|
| | | |

#### IBM Wave Administrator ID

When you start the IBM Wave webstart application for the first time, you are immediately prompted to create the IBM Wave Administrator. Complete the information that is shown in Table 3-2 on page 21.

In our environment, we use the following information. You may pick whatever Security Question and Security Question Answer that you want, but we show our information here.

- User Name: WAVEADM
- Security Question: Name of first cat
- Security Question Answer: Garfield

*Table 3-2 Wave Administrator information*

| | |
|---|---|
| **User name** | |
| **Password** | |
| **Security Question** | |
| **Security Question Answer** | |

## Central processor complex details

After you install IBM Wave, you define your central processor complex (CPC).

You provide a name to identify uniquely your CPC, the model of your CPC (for example, 2827 / zEC12 or 2964 / z13), and the CPU ID of your system.

This CPU ID is used to validate that your system is licensed to run IBM Wave. To obtain this information, run the CP command **`QUERY CPUID`**.

This command returns a response in the following format:

```
CPUID = aasssssscccdddd
```

The `ssssss` string (third through eighth hexadecimal digits) is the value that you must enter in IBM Wave when you are defining a CPC. Record this information in Table 3-3.

In our environment, we use the following information:

- Name: ITSOZVM1
- Model: 2964
- CPU ID: 09DA87

*Table 3-3 Central processor complex information*

| Name | Model | CPU ID |
|---|---|---|
| | | |

## z/VM system details

After you define your CPC, you can add z/VM systems. This process includes the use of the z/VM system name to create an External Entity, and then using other values to complete the Create New z/VM system window.

The Name value is used to identify uniquely your z/VM system in IBM Wave. A preferred choice for this name might be your z/VM system ID as it is defined in the z/VM `SYSTEM CONFIG` file.

The z/VM Version and z/VM Service Level values identify the level of z/VM that is running on this z/VM system. These values can be obtained by running the CP command **`QUERY CPLEVEL`**.

The IPv4 Address, IPv6 Address, and Host Name fields define how the IBM Wave server can establish a TCP/IP connection to the z/VM system.

The 3270 Port # and TLS/SSH values reflect how your z/VM environment is configured to allow TN3270 terminal sessions. In most environments, the 3270 Port # is 23 or 3270.

The Performance Toolkit and TCP/IP service machine names are the names of the z/VM users that are running these service machine functions in your z/VM environment. The Performance Toolkit service machine often is named PERFSVM and the TCP/IP service machine is named TCPIP. You can check these names by running the CP command **QUERY NAMES**.

The TCP/IP service machine contains a minidisk that is used to store the `TCPIP PROFILE` file. By default, this file is stored on a minidisk at the 592 address. You can confirm this information by accessing the 592 minidisk of TCPIP and looking for the `PROFILE` file.

Use the worksheet that is shown in Table 3-4 to record the information about your z/VM system.

In our environment, we use the following information:

- Name: ITSOZVM1
- z/VM Version: 6.3
- z/VM Service Level: 1501
- IPV4 address: 9.12.xx.xx
- IPV6 address: n.a.
- Host name: ITSOZVM1
- 3270 Port #: 23
- 3270 through TLS/SSH: n.a.
- Performance Toolkit Service Machine Name: PERFSVM
- Address of TCPIP Minidisk containing Profile: 198

*Table 3-4 z/VM system information*

| | |
|---|---|
| **Name** | |
| **z/VM Version** | |
| **z/VM Service Level** | |
| **IPv4 Address** | |
| **IPv6 Address** | |
| **Host Name** | |
| **3270 Port #** | |
| **3270 through TLS/SSH?** | |
| **Performance Toolkit Service Machine Name** | |
| **TCPIP Service Machine Name** | |
| **Address of TCPIP Minidisk containing PROFILE** | |

## Directory manager details

To use IBM Wave, your z/VM environment must run a directory manager, such as DirMaint.

To interact with the directory manager in your z/VM environment, IBM Wave requires the user name and password of a z/VM user who is authorized to run SMAPI commands. In the case of DirMaint, this user requires specific privilege classes and the ability to run commands without entering a password. This user name and password can be any z/VM user, including the actual IBM Wave server.

In z/VM environments that are running DirMaint, IBM Wave creates a Dummy Region of one cylinder, which is used as a placeholder because DirMaint cannot create empty storage groups. The Name and VOLID values are used to create this Dummy Region during the initial Auto-Detection process for adding a z/VM system. Although this region must be defined on a real DASD volume, this Dummy Region is never allocated. Enter a valid VOLID and a name for this Dummy Region, such as DUMMY.

Use the worksheet that is shown in Table 3-5 to record the information about the directory manager. In our environment, we use the following information:

- Type of Directory Manager: DirMaint
- User Name: DirMaint
- Password: <password>
- Volid for Wave Dummy Region: DUMMY
- Name of Wave Dummy Region: DUMMY

*Table 3-5 Directory manager information*

| | |
|---|---|
| **Type of Directory Manager** | |
| **User Name** | |
| **Password** | |
| **VOLID for Wave Dummy Region** | |
| **Name of Wave Dummy Region** | |

## IBM Wave service machine details

Although IBM Wave uses the Systems Management API (SMAPI), it also uses three CMS service machines. The Short Service Machine runs various compiled REXX EXECs to interact with the z/VM environment on behalf of the z/VM GUI. The Long Service Machine performs the same function, but it provides another thread of execution to run longer scripts in a way that does not prevent the running of the numerous short IBM Wave scripts that are constantly running against your z/VM environment. The Cross-System Cloning Service Machine is used to stream minidisks from a source z/VM system to a target z/VM system during cloning actions. These service machines often are named WAVEWRKS, WAVEWRKL, and WAVEWRKC. In addition to providing the names of these service machines, you must provide a storage group to allow IBM Wave to allocate various minidisks for these service machines.

Use the worksheet that is shown in Table 3-6 to record the information about the IBM Wave service machines. In our environment, we use the following information:

- Short Service Machine Name: WAVEWRKS
- Long Service Machine Name: WAVEWRKL
- CSC Service Machine Name: WAVEWRKC
- Storage Group: WAV009

*Table 3-6 IBM Wave service machine information*

| | |
|---|---|
| **Short Service Machine Name** | |
| **Long Service Machine Name** | |
| **CSC Service Machine Name** | |
| **Storage Group for IBM Wave Service Machines' minidisks** | |

## vSwitch details

After IBM Wave completes the auto-detection process for your z/VM system, it creates metadata objects for each of your vSwitches. However, IBM Wave flags each of your vSwitches as "Objects that require attention" in the System Status tab of your z/VM system's Current System View. This flag occurs because the metadata objects of your vSwitches lack defined Default NIC values. When IBM Wave is used to connect z/VM guests to vSwitches, IBM Wave adds a virtual NIC to each of these guests with an address that is equal to the Default NIC value of the vSwitch to which the guest is connecting. Because a virtual NIC is composed of three OSA addresses, the Default NIC value is only the first of three consecutive addresses. For example, if a vSwitch has a Default NIC value of AC00, this NIC uses AC00, AC01, and AC02.

When you are planning the Default NIC values for your vSwitches, it is important to ensure that the Default NIC values do not overlap. Therefore, if one vSwitch uses AC00, another vSwitch cannot use AC01 or AC02 because this value overlaps. The second vSwitch must be AC03 to prevent an overlap. You can get this information from your network administrator.

Use the worksheet that is shown in Table 3-7 to record the information about vSwitch.

*Table 3-7 vSwitch information*

| vSwitch Name | Wave Default NIC |
|---|---|
| | |
| | |
| | |
| | |
| | |

## Virtual Network Segment details

IBM Wave automatically generates Virtual Network Segments that are based on the assumption that all segments are IPv4 Class C blocks. It is likely that this assumption is invalid, which requires you to define manually your Virtual Network Segments correctly. In many environments, you might get this information from your network administrator.

Use the worksheet that is shown in Table 3-8 to record the information about Virtual Network Segments.

*Table 3-8 Virtual Network Segment information*

| Name | Network IP | Gateway IP | Netmask | Broadcast IP | VLAN ID |
|---|---|---|---|---|---|
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |

4

# Installation and configuration

This chapter provides step-by-step instructions to guide you in the installation of IBM Wave of z/VM (IBM Wave). It starts with preinstallation tasks for your z/VM environment and Linux on z Systems guests.

Next, this chapter guides you in the installation of IBM Wave in your environment. Finally, it provides instructions about post-configuration tasks that are required before you can start detecting your environment.

This chapter includes the following topics:

- Preinstallation tasks
- IBM Wave installation
- IBM Wave Command-Line Interface installation

# 4.1 Preinstallation tasks

To install successfully IBM Wave, you must ensure that you have all of the prerequisite software installed and configured on z/VM and the Linux system that is used as the WAVESRV Linux server. The list of the software prerequisites for z/VM, The IBM Wave server, and your local workstation are described in Chapter 3, “Planning and design” on page 13. Before completing the product installation, all prerequisites must be met.

## 4.1.1 Managed Linux system prerequisites

For every Linux on z Systems virtual machine (VM) that is managed by IBM Wave, the SSH, VMCP, and CMSFS packages must be installed.

## 4.1.2 Preinstallation configuration tasks for z/VM

Before you install the IBM Wave product into the WAVESRV system, you must complete the following z/VM configuration tasks, which are also described in Chapter 3, “Planning and design” on page 13.

## 4.1.3 z/VM system management application programming interface

IBM Wave uses the system management application programming interface (SMAPI) to perform various tasks on the z/VM system. Therefore, it is necessary to make sure that SMAPI is configured and running before IBM Wave is installed. The exact configuration steps for SMAPI differ based on the version and release level of z/VM.

For more information about the SMAPI, see *The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3*, SG24-8147.

Based on the example environment that is running z/VM V6.3, to configure SMAPI for use with IBM Wave, complete the following steps:

1. Modify the `DMSSISVR NAMES` file to uncomment the directory manager definition and the dump handler definition. Use the VMSES/E **`LOCALMOD`** command to change this file. You use the MAINT630 user ID, so it already has access to minidisk 5E6, which contains the **`LOCALMOD`** exec.

```
q disk
LABEL  VDEV M  STAT   CYL TYPE BLKSZ   FILES  BLKS USED-(%) BLKS LEFT  BLK TOTAL
MNT191 191  A   R/W   175 3390 4096       77        605-02      30895      31500
MNT5E6 5E6  B   R/W     9 3390 4096      134       1304-80        316       1620
MNT2CC 2CC  C   R/W    10 3390 4096        4        189-11       1611       1800
MNT51D 51D  D   R/W    26 3390 4096      308       1915-41       2765       4680
PMT551 551  E   R/W    40 3390 4096       12        133-02       7067       7200
MNT190 190  S   R/O   207 3390 4096      698      20272-54      16988      37260
MNT19E 19E  Y/S R/O   500 3390 4096     1124      30408-34      59592      90000
```

*Figure 4-1 Maint630 minidisks that are accessed*

For more information about the **`LOCALMOD`** command, see *VMSES/E z/VM Service Guide,* GC24-6232.

Complete the following steps:

a. From the MAINT630 user ID, enter the command that is shown in Figure 4-2.

```
LOCALMOD CMS DMSSISVR NAMES
```

*Figure 4-2 VMSES/E update DMSSISVR names*

b. Enter `1` when you are prompted to continue. The system places you into an XEDIT session for the `DMSSISVR NAMES` file.

c. Scroll to the lines that are shown in Figure 4-3.

```
*:server.DIRMAINT
*:type.DMGR
```

*Figure 4-3 Default DIRMAINT section*

d. Modify the lines that are shown in Figure 4-4.

```
:server.DIRMAINT
:type.DMGR
```

*Figure 4-4 Uncommented DIRMAINT section*

e. Scroll to the lines that are shown in Figure 4-5.

```
* Dump Handler
*:server.OPERATNS
*:type.WORKER
*:short.DMPH
```

*Figure 4-5 Default Dump Handler section*

f. Modify the lines that are shown in Figure 4-6.

```
* Dump Handler
:server.OPERATNS
:type.WORKER
:short.DMPH
```

*Figure 4-6 Uncommented Dump Handler section*

g. Run **`file`** to save the changes. The system creates an override file that is named `DMSSISVR NAML0002` with the contents that you modified on the Conversational Monitoring System (CMS) LOCALMOD minidisk (MAINT630 3C4).

h. Enter the command that is shown in Figure 4-7 to rebuild the file, including the new modifications.

```
SERVICE CMS BUILD
```

*Figure 4-7 Rebuild modified DMSSISVR NAMES File*

i. Enter the command that is shown in Figure 4-8 to move the new `DMSSISVR NAMES` file from the alternative build disk to the runtime disk (MAINT 193).

```
PUT2PROD
```

*Figure 4-8 Put updated DMSSISVR NAMES into production*

**Note:** If you have an single-system image (SSI) environment, you must run **PUT2PROD** on all SSI members.

2. Shared File System (SFS) definition

   The SMAPI authorized users file is named `VSMWORK1 AUTHLIST`, which is in the `VMSYS:` file pool under the root directory for VSMWORK1 (`VMSYS:VSMWORK1.`). To access it, run the command that is shown in Figure 4-9.

```
acc VMSYS:VSMWORK1. z
filelist * * z
```

*Figure 4-9 Default AUTHLIST file*

The default file that is included with z/VM V6.3 has the data that is shown in Figure 4-10.

```
DO.NOT.REMOVE
DO.NOT.REMOVE
MAINT                                                    ALL
ZHCP                                                     ALL
```

*Figure 4-10 Default AUTHLIST File*

In our environment, we used the MAINT user ID for authorization within IBM Wave. If you want to add a user ID to this file, copy an existing line and then replace only the user ID text on the duplicated line.

3. As many of the SMAPI calls are carried out by the directory manager VM (DIRMAINT), you must authorize the SMAPI worker VMs to use DirMaint. IBM Wave also includes worker VMs that must be authorized to use DirMaint directly. You must update the Configuration override file (`CONFIGxx DATADVH`) and Authorization file (`AUTHFOR CONTROL`) to include SMAPI and the IBM WAVE VMs. This task is described in 4.1.4, “DirMaint” on page 30.

## 4.1.4 DirMaint

DirMaint provides organized and secure interactive facilities for maintaining the z/VM system directory. Directory management is simplified by the DirMaint command interface and automated facilities. DirMaint supports all the z/VM directory statements. Most DirMaint directory commands have the same names and format as the VM directory statements that they support.

**Note:** This section assumes that DirMaint is installed and is active in the environment.

This configuration can be done from the MAINT630 user ID. You must configure the DirMaint files that are described in this section. Complete the following steps:

1. Edit the CONFIGAA DATADVH file.
   a. When you configure DirMaint for use, you create a configuration override file to change to the CONFIG DATADVH file that is supplied by IBM. In our environment, we named that file CONFIGAA DATADVH.
   b. Enter the command that is shown in Figure 4-11 for DirMaint to send you the configuration file.

```
DIRM SEND CONFIGAA DATADVH
```

*Figure 4-11 Get CONFIGAA DATADVH file*

   c. Receive the spool file that was sent by DirMaint by using the **CMS RECEIVE** command. The configuration file (CONFIGAA DATADVH) is placed on the MAINT630 191 minidisk.
   d. Use XEDIT to modify the file by adding the lines that are shown in Figure 4-12.

```
ALLOW_ASUSER_NOPASS_FROM= VSMGUARD *
ALLOW_ASUSER_NOPASS_FROM= VSMWORK1 *
ALLOW_ASUSER_NOPASS_FROM= VSMWORK2 *
ALLOW_ASUSER_NOPASS_FROM= VSMWORK3 *
ALLOW_ASUSER_NOPASS_FROM= WAVEWRKS *
ALLOW_ASUSER_NOPASS_FROM= WAVEWRKC *
ALLOW_ASUSER_NOPASS_FROM= WAVEWRKL *
ASYNCHRONOUS_UPDATE_NOTIFICATION_EXIT.TCP= DVHXNE EXEC
ASYNCHRONOUS_UPDATE_NOTIFICATION_EXIT.UDP= DVHXNE EXEC
```

*Figure 4-12 Updates to CONFIGAA DATADVH*

**Note:**

- SMAPI users can use the ALLOW_ASUSER_NOPASS_FROM lines to issue commands to the Directory Manager by using the **ASUSER** modifier and the password of that user.
- The ASYNCHRONOUS_UPDATE_NOTIFICATION_EXIT lines activate an exit that notifies SMAPI of changes that are made to the user directory.
- If privacy of residual data is a concern on your system, use DISK_CLEANUP= YES.
- The ONLINE= IMMED line sets your changes to be made immediately.
- The RUNMODE= OPERATIONAL line sets directory changes to be made. This line can be set to TESTING and the changes will not be made.

e. If you did not activate DATAMOVE when you originally configured DirMaint for use, you should do so now because IBM Wave requires that DATAMOVE be available. Add the statements that are shown in Figure 4-13 to the `CONFIGAA DATADVH` file.

```
DISK_CLEANUP= YES
DATAMOVE_MACHINE= DATAMOVE * *
ONLINE= IMMED
RUNMODE= OPERATIONAL
RACF_RDEFINE_VMBATCH_DEFAULTS=
```

*Figure 4-13 Activate DATAMOVE*

**Note:**

- If privacy of residual data is a concern on your system, use `DISK_CLEANUP= YES`.
- Specify the name of the DATAMOVE VM. If you have an SSI environment, you have at least one DATAMOVE per SSI member.
- The `ONLINE= IMMED` line sets your changes to be made immediately.
- The `RUNMODE= OPERATIONAL` line sets directory changes to be made. This line can be set to `TESTING` and the changes will not be made.
- The `RACF_RDEFINE_VMBATCH_DEFAULTS=` line does not create a VMBATCH-specific resource entry. Otherwise, DirMaint creates a VMBATCH resource for this user ID with this line as a default. If you are not installing RACF, this line can be omitted.

f. Run **`file`** to save the changes.

g. Run the commands that are shown in Figure 4-14 to send the modified file back to DirMaint and to have DirMaint reprocess the contents of the file.

```
DIRM FILE CONFIGAA DATADVH A
DIRM RLDDATA
DIRM RLDCODE
```

*Figure 4-14 Send And activate the configuration file*

**Note:** If you are in an SSI environment, you must to create another `CONFIGXX DATADVH` file to specify the other DirMaint related machines that must be defined in the SSI members. For more information, see “The DirMaint configuration in an SSI environment” on page 35.

2. Edit the `AUTHFOR CONTROL` file.

   a. To direct DirMaint to send you the `AUTHFOR CONTROL` file, run the command that is shown in Figure 4-15.

```
DIRM SEND AUTHFOR CONTROL
```

*Figure 4-15 Get AUTHFOR file*

   b. You receive the spool file that was sent by DirMaint by running **`CMS RECEIVE`**. The configuration file is placed on the MAINT630 191 minidisk.

c. Run **XEDIT** to modify the file by adding the lines that are shown in Figure 4-16.

```
ALL VSMGUARD * 140A ADGHMOPS
ALL VSMGUARD * 150A ADGHMOPS
ALL VSMWORK1 * 140A ADGHMOPS
ALL VSMWORK1 * 150A ADGHMOPS
ALL VSMWORK2 * 140A ADGHMOPS
ALL VSMWORK2 * 150A ADGHMOPS
ALL VSMWORK3 * 140A ADGHMOPS
ALL VSMWORK3 * 150A ADGHMOPS
ALL WAVEWRKS * 140A ADGHMOPSZ
ALL WAVEWRKS * 150A ADGHMOPSZ
ALL WAVEWRKL * 140A ADGHMOPSZ
ALL WAVEWRKL * 150A ADGHMOPSZ
ALL WAVEWRKC * 140A ADGHMOPSZ
ALL WAVEWRKC * 150A ADGHMOPSZ
```

*Figure 4-16 Updates to AUTHFOR CONTROL*

**Note:** When giving access into the authorization list, include records for both 140A and 150A command levels for each target ID/authorized user pair to keep the compatibility with both DirMaint command levels.

d. Run **file** to save the changes.

e. To send the modified AUTHFOR CONTROL file back to DirMaint and to cause DirMaint to reprocess the content of the changed file, run the command that is shown in Figure 4-17.

```
DIRM SEND AUTHFOR CONTROL A
DIRM RLDDATA
DIRM RLDCODE
```

*Figure 4-17 Send and activate AUTHFOR changes*

f. You can ensure that your AUTHFOR updates complete by running the **DIRMAINT** command, as shown in Figure 4-18.

```
DIRM FOR ALL AUTHFOR ?
```

*Figure 4-18 Check the AUTHFOR changes*

The command response should show that VSMGUARD, VSMWORK1 - 3, and WAVEWRKS, WAVEWRKL, and WAVEWRKC are all authorized.

g. To support IBM Wave, the DIRMAINT VM must run with the C and E privilege classes (PRIVCLASS), in addition to the default classes of BD and G. The command that is shown in Figure 4-19 adds privilege classes C and E to the DIRMAINT VM.

```
DIRM FOR DIRMAINT PRIV +CE
```

*Figure 4-19 Add privilege classes to DIRMAINT*

The new privilege classes do not take effect until the DIRMAINT VM is logged off and then logged back on. You can accomplish this task by running the **DIRM SHUTDOWN** and then running **XAUTOLOG DIRMAINT** after the DIRMAINT machine fully logs off from the system.

3. Edit the EXTENT CONTROL file.

   The EXTENT CONTROL file defines disks (volumes) to DirMaint for minidisk allocation. It also contains system and device default values that are used during allocation operations. There are two main sections that should be populated:

   – *Regions* define the disks and their sizes to DirMaint. The **AUTOR** keyword can be used in user directory entries to take space from the regions. As a preferred practice, the region name and volume label are always identical.

   – *Groups* define pools of disks so the **AUTOG** keyword can be used to take space from the pools, and not from specific disks.

   Because IBM Wave allocates DASD space, it is necessary to configure the DirMaint EXTENT CONTROL file before IBM Wave is installed. At a minimum, you should have one region that is defined. When IBM Wave is installed, a dummy region of one cylinder is created as a placeholder for “empty” storage groups. Although this region is defined on a real volume, it is never physically allocated. In our environment, we created a group with the name WAV009 for use by IBM Wave when VMs are created, as shown in Figure 4-20.

```
:REGIONS.
*RegionId  VolSer    RegStart      RegEnd  Dev-Type  Comments
...
DUMMY      DUMMY            1           1  3390-09
LXD81A     LXD81A           1         END  3390-09
LXD81B     LXD81B           1         END  3390-09
LXD819     LXD819           1         END  3390-09
LXD818     LXD818           1         END  3390-09
LXD81C     LXD81C           1         END  3390-09
...
:END.
:GROUPS.
*GroupName RegionList
*
WAV009     LXD020   LXD021   LXD022   LXD023   LXD024   LXD026   LXD027
WAV009     LXD028   LXD029   LXD02A   LXD121   LXD122   LXD123   LXD124
WAV009     LXD125   LX991E   LX991F   LX9920   LX9921   LX9922   LX9923
WAV009     LXD81B   LXD819   LXD81A   LXD818   LXD81C
...
:END.
```

*Figure 4-20 Excerpt of the EXTENT CONTROL file*

To avoid problems during the IBM Wave installation, you want to ensure that the EXCLUDE section (Figure 4-21 on page 35) is configured to prevent any problems because of full pack overlap allocations for IBM VMs, such as MAINT, PMAINT, SYSDUMP, or any other that you created in your installation.

```
:EXCLUDE.
* ENTRY_NAME ADDRESS
MAINT* 012*
MAINT* 013*
PMAINT 013*
PMAINT 014*
SYSDUMP1 012*
SYSDMP* 012*
:END.
```

*Figure 4-21 Exclude definition*

To double-check the region name and volume name that you use during the IBM Wave installation, run the commands that are shown in Figure 4-22 and ensure that a return code of zero is returned. In our environment, we used the region name and volume name LXD81C. You must substitute a value for LXD81C that matches the information in your `EXTENT CONTROL` file.

```
DIRM DASD QUERY REGION LXD81C
DIRM DASD QUERY VOLUME LXD81C
```

*Figure 4-22 Check the DASD that is defined in the Region section*

If you do not receive a return code of 0 from the commands that are shown in Figure 4-22, you must fix any problems in your `EXTENT CONTROL` file before you proceed with the installation of IBM Wave.

Anytime you modify the `EXTENT CONTROL` file, you must reload so that DirMaint can access the new configuration, as shown in Figure 4-23.

```
DIRM RLDE
```

*Figure 4-23 DirMaint command after changing EXTENT CONTROL*

Anytime you modify configuration or authfor files, you must reload them so that DirMaint can access the new configuration, as shown in Figure 4-24.

```
DIRM RLDC
DIRM RLDD
```

*Figure 4-24 DirMaint command after changing other DirMaint files*

## The DirMaint configuration in an SSI environment

When running DirMaint in an SSI environment, the USER DIRECT source is shared among all SSI members. DIRMAINT and DATAMOVE user IDs run in one member only, DIRMSATx and DATAMOVx user IDs run in the other SSI members.

For more information about a DirMaint installation, see *The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3*, SG24-8147.

You must follow the steps in this publication to enable the product and put the code in production in each z/VM of the SSI environment.

**Note:** This publication assumes that DirMaint is installed and is active in the SSI environment with one member only.

DIRMSAT user IDs are the VMs that communicate with DirMaint directly when receiving commands from the other z/VM members of SSI. In each z/VM, there are DIRMAINT, DIRMSAT, and DATAMOVE user IDs.

In our SSI environment, DirMaint is running in ITSOZVM1, DIRMSAT2 in ITSOZVM2, DIRMSAT3 in ITSOZVM3, and DIRMSAT4 in ITSOZVM4. ITSOZVM1 also contains DATAMOVE and is defined as DATAMOV2, DATAMOV3, and DATAMOV4 on the respective SSI members.

To configure DirMaint and the related VMs to support other members of SSI, complete the following steps.

1. Log on as MAINT.
2. Verify that the user IDs DIRMSAT2, DIRMSAT3, DIRMSAT4, DATAMOV2, DATAMOV3, and DATAMOV4 are defined with an **AUTOONLY** password by running **dirm for <name> get** for each user ID. DirMaint sends the directories of each machine to your reader and you must run **readerlist** to check each directory.
3. If they have **AUTOONLY** as the password, run **receive** on each file from your reader to write the directories to your 191 minidisk.
4. Run **XEDIT** command to edit each directory file and modify the password to a password that is used in your environment.
5. Run **file** to each directory file to save the changes.
6. Send the modified directories to DirMaint by running **dirm for <name> replace**.

All DIRMSATx and DATAMOVEx users are now set up and can be used in an SSI environment.

Now, you must verify that CONFIGSS DATADVH has entries for the DIRMSAT and DATAMOVE user IDs. Run **DIRM SEND CONFIGSS DATADVH** to get this file, and check and include the definitions for each SSI member (z/VM), as shown in Figure 4-25.

```
*********************************************************
*  CONFIGSS DATADVH
*  Created at install time to define a Satellite Server
*  for each member
*********************************************************
SATELLITE_SERVER= DIRMSAT ITSOZVM1
SATELLITE_SERVER= DIRMSAT2 ITSOZVM2
SATELLITE_SERVER= DIRMSAT3 ITSOZVM3
SATELLITE_SERVER= DIRMSAT4 ITSOZVM4
DATAMOVE_MACHINE= DATAMOVE ITSOZVM1 *
DATAMOVE_MACHINE= DATAMOV2 ITSOZVM2 *
DATAMOVE_MACHINE= DATAMOV3 ITSOZVM3 *
DATAMOVE_MACHINE= DATAMOV4 ITSOZVM4 *
```

*Figure 4-25 DirMaint definition of DIRMSAT and DATAMOVE user IDs in an SSI environment*

After you modify the `CONFIGSS DATADVH` file and save it by running **`file`**, send this configuration file to the DIRMAINT user ID by running **`DIRM FILE CONFIGSS DATADVH`**. To reload all of the DirMaint configurations, run the command that is shown in Figure 4-26.

```
dirm rldc
dirm rldd
```

*Figure 4-26 Load the DirMaint configuration*

Now, you can bring all DIRMSATs and DATAMOVEs user IDs online in each SSI member. To do this task, log on to the MAINT user ID in each SSI member (ITSOZVM2, ITSOZVM3, and ITSOZVM4) and run the following commands:

1. **`xautolog DIRMSAT2`** in ITSOZVM2
2. **`xautolog DIRMSAT3`** in ITSOZVM3
3. **`xautolog DIRMSAT4`** in ITSOZVM4

You do not need to do this task in ITSOZVM1 because that is where DirMaint is running.

Now, you have DirMaint up and running in all SSI members. You can test it by running the following DirMaint command in each z/VM:

```
DIRM FOR MAINT REVIEW
```

For more information about DirMaint, see *The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3*, SG24-8147.

### 4.1.5 TCP/IP for z/VM

Because IBM Wave can establish 3270 sessions from within the IBM Wave GUI, the `PROFILE TCPIP` file must be updated for the z/VM TCP/IP stack to authorize the IBM Wave worker machines, WAVEWRKS (short) and WAVEWRKL (long), to use the **`obeyfile`** command. You also want to ensure that the internal client and FTP server are active.

By default, the `PROFILE TCPIP` file is created on the TCPMAINT VM's 198 minidisk. Add the VMs to the `OBEY` section of that file, as shown in Figure 4-27.

```
OBEY
OPERATOR TCPMAINT MAINT MPROUTE REXECD SNMPD SNMPQE LDAPSRV MAINT630
WAVEWRKS WAVEWRKL
ENDOBEY
```

*Figure 4-27 Updates to the OBEY section*

In Figure 4-27, the first line that follows the keyword **`OBEY`** is part of the default definitions when z/VM V6.3 is installed and the IPWIZARD utility is used to configure an initial TCP/IP connection. In our environment, we added the WAVEWRKS and WAVEWRKL VMs to this list.

Ensure that the PORT section of your PROFILE TCPIP file contains the uncommented lines, as shown in Figure 4-28.

```
PORT
  20   TCP FTPSERVE  NOAUTOLOG ; FTP Server
  21   TCP FTPSERVE            ; FTP Server
  23   TCP INTCLIEN            ; TELNET Server
```

*Figure 4-28 Review the PORT section*

Finally, ensure that the FTPSERVE VM is automatically started by the TCPIP VM. You should see an AUTOLOG section similar to the example that is shown in Figure 4-29.

```
AUTOLOG
FTPSERVE X
ENDAUTOLOG
```

*Figure 4-29 Review the AUTOLOG section*

After the changes are made to the PROFILE TCPIP file, you must restart your TCPIP VM to pick up these changes.

**Note:** Take care when restarting TCPIP if you are accessing z/VM remotely through a 3270 emulation session, such as when using IBM Personal Communications. This session is unavailable during this process and you must access z/VM from the Hardware Management Console (HMC) by opening a 3270 window.

## 4.1.6 Performance Toolkit

IBM Wave uses IBM Performance Toolkit to display performance information for individual VMs. For more information about setting up and configuring Performance Toolkit, see *z/VM Performance Toolkit Reference*, SC24-6210.

For IBM Wave to use Performance Toolkit, the short duration worker machine (WAVEWRKS) must be authorized to gather data from Performance Toolkit. This authorization is accomplished by adding WAVEWRKS to the FCONRMT AUTHORIZ file that is maintained on the 191 minidisk of the PERFSVM VM. To update this file, log off from the PERFSVM VM, link to its 191 minidisk, and then modify the file (which means that performance monitoring is temporarily suspended). You also can log on to the PERFSVM VM and update the file while it is still running. If you choose to update the file while PERFSVM is still running, you must run **FCONTROL RELOAD AUTHORIZ** after the file is changed.

To authorize the IBM Wave short duration worker machine, update the FCONRMT AUTHORIZ file with the entry that is shown in Figure 4-30.

```
WAVEWRKS    *    DATA CMD
```

*Figure 4-30 Update the FCONRMT AUTHORIZ file*

IBM Wave requirements for Performance Toolkit include making sure that the VMCF interface is turned on, which is accomplished by uncommenting the FC MONCOLL VMCF ON statement in the FCON $PROFILE file that is shown in Figure 4-31 on page 39.

```
FC MONCOLL VMCF ON
```

*Figure 4-31 FCON $PROFILE update*

It is also necessary to configure the CP Monitor to enable collection of sample records from the processor, storage, and user (all) domains. This task is accomplished by making certain that the `PROFILE EXEC` file for PERFSVM includes the commands that are shown in Figure 4-32 uncommented.

```
'CP MONITOR SAMPLE ENABLE PROCESSOR'
'CP MONITOR SAMPLE ENABLE STORAGE'
'CP MONITOR SAMPLE ENABLE USER ALL'
```

*Figure 4-32 CP Monitor sample domains*

These configuration steps are done when you are setting up Performance Toolkit for general system use. However, it is good to review the `FCONRMT PROFILE` and the `PERFSVM PROFILE EXEC` files to be certain that the necessary configuration steps are complete.

### 4.1.7 Single-system image definitions

A z/VM SSI cluster is a multisystem environment in which the z/VM systems can be managed as a single resource pool and guests can be moved from one system to another while they are running. With the IBM z/VM SSI feature, a running Linux VM can be relocated non-disruptively from one member system to any other member by a process that is known as Live Guest Relocation (LGR). This process provides application continuity across planned z/VM and hardware outages.

IBM Wave fully supports z/VM systems that are members of SSI clusters. Such systems are grouped in the Hardware Viewer, in a similar fashion to z/VM systems sharing a directory. The color of the group in the Hardware Viewer (light blue) indicates that they are a part of an SSI cluster, as shown in Figure 4-33, and the name of the group reflects the name of the SSI cluster as it appears in the output of the **`QUERY SSI`** CP command.

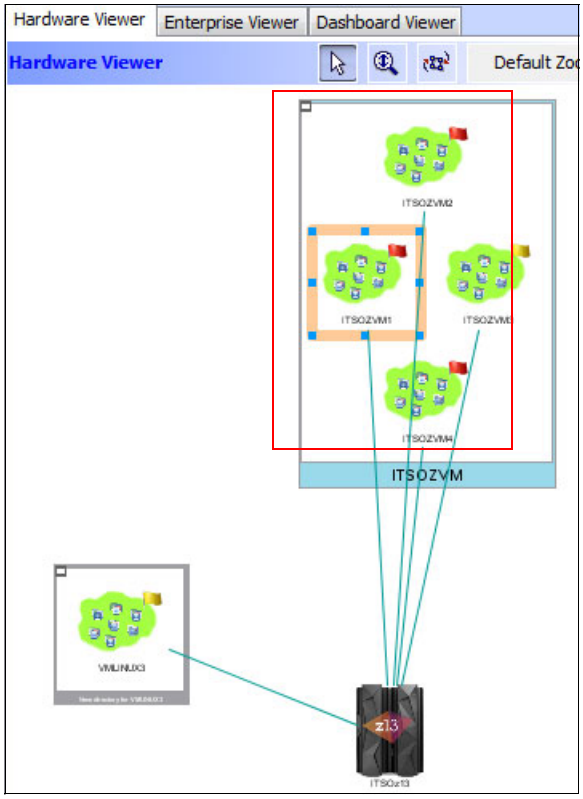


*Figure 4-33 Hardware viewer of an SSI cluster of four z/VMs*

## Live Guest Relocation

LGR is fully supported, under the same conditions, provisions, and limitations that are imposed by z/VM V6.2 and above. LGR can be run against one or more guests, as a multiple task action, either by selecting the appropriate action from the menu when right-clicking the guests or by dragging the guests in the Enterprise Viewer when the “Group By” setting is set to “z/VM system”. The LGR process generates a BTS Workunit that can be tracked from the BTS Workunit Viewer.

**Note:** To drag a guest, you must select the guest icon from the z/VM LPAR and press the CTRL key at the same time, moving the guest icon to the z/VM LPAR target.

You can do an LGR of WAVESRV itself. To make WAVESRV eligible to be relocated, you must complete the following steps:

1. Certify that the directory definition of this server has the option **`CHPIDV ONE`**.
2. Remove the CMS definitions, such as 190, 19e, 19d, and 191 minidisks from the WAVESRV user directory. You can remove them by either deleting the definitions or comment out the definitions by putting an asterisk in front of the definitions, as shown in Figure 4-34 on page 41.

```
USER WAVESRV LNX4ITSO 2G 4G GC
   COMMAND DEFINE STORAGE 2G STANDBY 2G
   COMMAND SET RUN ON
   CPU 00
   IPL 152
   MACHINE ESA 4
   OPTION QUICKDSP CHPIDV ONE
   CONSOLE 0009 3215
   NICDEF 0800 TYPE QDIO LAN SYSTEM VSWITCH1
   SPOOL 000C 3505 A
   SPOOL 000D 3525 A
   SPOOL 000E 1403 A
*  LINK MAINT 0190 0190 RR
*  LINK MAINT 019D 019D RR
*  LINK MAINT 019E 019E RR
*  MDISK 0191 3390 9325 50 LXD022 MR READ WRITE MULTI
   MDISK 0150 3390 1625 3000 LXD022 MR READ WRITE MULTI
   MDISK 0151 3390 4625 4500 LXD022 MR READ WRITE MULTI
   MDISK 0152 3390 9125 200 LXD022 MR READ WRITE MULTI
   MDISK 0153 3390 1 10015 LXD028 MR READ WRITE MULTI
   MDISK 0154 3390 1 10015 LXD029
```

*Figure 4-34 WAVESRV directory definition to be eligible for relocation*

If you do not modify the WAVESRV directory definition, when you try to do the server *relocation test*, you will see the message that is shown in Figure 4-35.

```
vmrelo test wavesrv to itsozvm2
HCPRLH1940E WAVESRV is not relocatable for the following reason(s):
HCPRLE1956I WAVESRV: Single path CHPID virtualization is not enabled
HCPRLI1996I WAVESRV: Virtual machine device 0190 is a link to a local minidisk
HCPRLI1996I WAVESRV: Virtual machine device 019D is a link to a local minidisk
HCPRLI1996I WAVESRV: Virtual machine device 019E is a link to a local minidisk
```

*Figure 4-35 VMRELOCATE test command failed*

With the correct definition, WAVESRV can be relocated to the other z/VM members of the SSI cluster, as shown in Figure 4-36.

```
vmrelocate test wavesrv to itsozvm2
User WAVESRV is eligible for relocation to ITSOZVM2
```

*Figure 4-36 VMRELOCATE test command with success*

After the server relocation test is successfully completed, you can now do the *relocation move* by running the command that is shown in Figure 4-37.

```
vmrelocate move wavesrv to itsozvm2
Relocation of WAVESRV from ITSOZVM1 to ITSOZVM2 started
User WAVESRV has been relocated from ITSOZVM1 to ITSOZVM2
```

*Figure 4-37 VMRELOCATE move command from z/VM member 1 to member 2*

After WAVESRV relocation is complete, you can run commands from the IBM Wave application to verify that everything works correctly.

To relocate the server back to its original position, log on to the MAINT user ID in the z/VM member and do the relocation move back to the original z/VM system, as shown in Figure 4-38.

```
vmrelo move wavesrv to itsozvm1
11:40:40 Relocation of WAVESRV from ITSOZVM2 to ITSOZVM1 started
11:40:42 User WAVESRV has been relocated from ITSOZVM2 to ITSOZVM1
11:40:42 USER DSC   LOGOFF AS  WAVESRV  USERS = 35    FORCED BY SYSTEM
```

*Figure 4-38 VMRELOCATE move command from z/VM member 2 to member 1*

## SSI user ID definitions

If the guest is defined in an SSI cluster, there are two types of VMs:

- Single-configuration VM definition

  A user ID that is defined by a single-configuration VM definition (the traditional type of user definition) can be logged on to any member of the SSI cluster, but only one member at a time. Single-configuration VMs are eligible for guest relocation. Use this type of VM for your Linux guests.

- Multiconfiguration VM definition

  A user ID that is defined by a multiconfiguration VM definition can be logged in to concurrently on multiple members of the SSI cluster. The instances of the user ID on the members have common attributes, but can also be configured to access different resources. Multiconfiguration VMs are not eligible for guest relocation. Use this type of VM for system support, such as the MAINT user ID and server VMs, such as TCP/IP.

### *IBM Wave service machines that are defined as Identities*

When you define service machines within an SSI cluster, IBM Wave supports the definition of the service machines as Identities with Subconfigurations. However, this implementation must be done manually by the z/VM system programmer.

**Note:** IBM Wave supports the service worker machines WAVEWRKS, WAVEWRKL, and WAVEWRC defined as User or Identities or a mix of definitions. If you want to use them as User, IBM Wave creates them during the AUTO DETECT process.

To define the three IBM Wave Service Machines that use Identities and Subconfigurations, complete the following steps. In our environment, there are four z/VM systems in the ITSOZVM SSI Cluster that are named ITSOZVM1, ITSOZVM2, ITSOZVM3, and ITSOZVM4:

**Note:** This procedure uses DirMaint commands.

1. Define the Identity for the Short Service Machine (WAVEWRKS):
   a. Log in to the MAINT user ID.
   b. Create a file that is named `WAVEWRKS DIRECT` on the A minidisk (191) by running the following commands:

      ```
      IDENTITY WAVEWRKS <PASSWORD> 64M 128M ABCDEFG
      IPL CMS
      MACHINE ESA 4
      ```

```
OPTION LNKNOPAS DIAG88
CONSOLE 0009 3215
SPOOL 000C 3505 A
SPOOL 000D 3525 A
SPOOL 000E 1403 A
LINK MAINT 0190 0190 RR
LINK MAINT 019D 019D RR
LINK MAINT 019E 019E RR
LINK TCPIP 0592 0592 RR
LINK MAINT 0193 0193 RR
```

<PASSWORD> is a password that is chosen by you for the three service machines.

c. Add the Identity by running the following command:

```
DIRM ADD WAVEWRKS
```

2. Define a SUBCONFIG for each system in the SSI cluster by completing the following commands. In our environment, we created four SUBCONFIGs, one for each z/VM LPAR (member of SSI).

   a. Log in to the MAINT user ID.

   b. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

```
X WAVWKS-1 DIRECT
```

   c. Paste the following content into the file, then save and exit:

```
SUBCONFIG WAVWKS-1
*CSLTAG01: WAVE-INTERNAL
AMD 0191 3390 AUTOG 100 <WAV009> RR PW <READ> <WRITE> <MULTI>
AMD 0399 3390 AUTOG 100 <WAV009> RR PW <READ> <WRITE> <MULTI>
```

   <WAV009> is the name of the DirMaint group pool name, and <READ> <WRITE> <MULTI> are the passwords that are chosen to access the minidisk.

   d. Add the SUBCONFIG by running the following command:

```
DIRM ADD WAVWKS-1 BUILD ON ITSOZVM1 IN WAVEWRKS
```

   e. Repeat steps a to c, changing the names of the SUBCONFIG as follows:

      i. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

```
X WAVWKS-2 DIRECT
```

      ii. Paste the following content into the file, then save and exit:

```
SUBCONFIG WAVWKS-2
*CSLTAG01: WAVE-INTERNAL
AMD 0191 3390 AUTOG 100 <WAV009> RR PW <READ> <WRITE> <MULTI>
AMD 0399 3390 AUTOG 100 <WAV009> RR PW <READ> <WRITE> <MULTI>
```

      <WAV009> is the name of the DirMaint group pool name, and <READ> <WRITE> <MULTI> are the passwords that are chosen to access the minidisk.

      iii. Add the SUBCONFIG by running the following command:

```
DIRM ADD WAVWKS-2 BUILD ON ITSOZVM2 IN WAVEWRKS
```

      iv. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

```
X WAVWKS-3 DIRECT
```

v. Paste the following content into the file, then save and exit:

```
SUBCONFIG WAVWKS-3
*CSLTAG01: WAVE-INTERNAL
AMD 0191 3390 AUTOG 100 <WAV009> RR PW <READ> <WRITE> <MULTI>
AMD 0399 3390 AUTOG 100 <WAV009> RR PW <READ> <WRITE> <MULTI>
```

<WAV009> is the name of the DirMaint group pool name, and <READ> <WRITE> <MULTI> are the passwords that are chosen to access the minidisk.

vi. Add the **SUBCONFIG** by running the following command:

```
DIRM ADD WAVWKS-3 BUILD ON ITSOZVM3 IN WAVEWRKS
```

vii. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

```
X WAVWKS-4 DIRECT
```

viii. Paste the following content into the file, then save and exit:

```
SUBCONFIG WAVWKS-4
*CSLTAG01: WAVE-INTERNAL
AMD 0191 3390 AUTOG 100 <WAV009> RR PW <READ> <WRITE> <MULTI>
AMD 0399 3390 AUTOG 100 <WAV009> RR PW <READ> <WRITE> <MULTI>
```

<WAV009> is the name of the DirMaint group pool name, and <READ> <WRITE> <MULTI> are the passwords that are chosen to access the minidisk.

ix. Add the SUBCONFIG by running the following command:

```
DIRM ADD WAVWKS-4 BUILD ON ITSOZVM4 IN WAVEWRKS
```

3. Define the Identity for the long Service Machine (WAVEWRKL) by completing the following steps:

   a. Log in to the MAINT user ID.

   b. Create a file that is named WAVEWRKL DIRECT on the A minidisk (191) by using the following process:

```
IDENTITY WAVEWRKL <PASSWORD> 64M 128M ABCDEFG
IPL CMS
MACHINE ESA 4
OPTION LNKNOPAS DIAG88
CONSOLE 0009 3215
SPOOL 000C 3505 A
SPOOL 000D 3525 A
SPOOL 000E 1403 A
LINK MAINT 0190 0190 RR
LINK MAINT 019D 019D RR
LINK MAINT 019E 019E RR
LINK TCPIP 0592 0592 RR
LINK MAINT 0193 0193 RR
LINK WAVEWRKS 0191 0191 RR
LINK WAVEWRKS 0399 0399 RR
```

<PASSWORD> is a password that is chosen by you for the three service machines.

   c. Add the Identity by running the following command:

```
DIRM ADD WAVEWRKL
```

4. Define a SUBCONFIG for each system in the SSI cluster by completing the following steps. In our environment, we create four subconfigurations, one for each z/VM LPAR (member of SSI).
   a. Log in to the MAINT user ID.
   b. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

      ```
      X WAVWKL-1 DIRECT
      ```

   c. Paste the following content into the file, then save and exit:

      ```
      SUBCONFIG WAVWKL-1
      *CSLTAG01: WAVE-INTERNAL
      ```

   d. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

      ```
      X WAVWKC-1 DIRECT
      ```

   e. Paste the following content into the file, then save and exit:

      ```
      SUBCONFIG WAVWKC-1
      *CSLTAG01: WAVE-INTERNAL
      ```

   f. Add the SUBCONFIG by running the following commands:

      ```
      DIRM ADD WAVWKL-1 BUILD ON ITSOZVM1 IN WAVEWRKL
      DIRM ADD WAVWKC-1 BUILD ON ITSOZVM1 IN WAVEWRKC
      ```

   g. Repeat steps a to f, changing the names of the SUBCONFIG as follows:
      i. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

         ```
         X WAVWKL-2 DIRECT
         ```

      ii. Paste the following content into the file, then save and exit:

         ```
         SUBCONFIG WAVWKL-2
         *CSLTAG01: WAVE-INTERNAL
         ```

      iii. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

         ```
         X WAVWKC-2 DIRECT
         ```

      iv. Paste the following content into the file, then save and exit:

         ```
         SUBCONFIG WAVWKC-2
         *CSLTAG01: WAVE-INTERNAL
         ```

      v. Add the SUBCONFIG by running the following commands:

         ```
         DIRM ADD WAVWKL-2 BUILD ON ITSOZVM2 IN WAVEWRKL
         DIRM ADD WAVWKC-2 BUILD ON ITSOZVM2 IN WAVEWRKC
         ```

      vi. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk by running the following command:

         ```
         X WAVWKL-3 DIRECT
         ```

      vii. Paste the following content into the file, then save and exit:

         ```
         SUBCONFIG WAVWKL-3
         *CSLTAG01: WAVE-INTERNAL
         ```

      viii. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk, by running the following command:

         ```
         X WAVWKC-3 DIRECT
         ```

ix. Paste the following content into the file, then save and exit

```
SUBCONFIG WAVWKC-3
*CSLTAG01: WAVE-INTERNAL
```

x. Add the SUBCONFIG by running the following commands:

```
DIRM ADD WAVWKL-3 BUILD ON ITSOZVM3 IN WAVEWRKL
DIRM ADD WAVWKC-3 BUILD ON ITSOZVM3 IN WAVEWRKC
```

xi. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk, by running the following command:

```
X WAVWKL-4 DIRECT
```

xii. Paste the following content into the file, then save and exit:

```
SUBCONFIG WAVWKL-4
*CSLTAG01: WAVE-INTERNAL
```

xiii. Create a file that is named <SUBCONFIG_Name> DIRECT on the A minidisk, by running the following command:

```
X WAVWKC-4 DIRECT
```

xiv. Paste the following content into the file, then save and exit:

```
SUBCONFIG WAVWKC-4
*CSLTAG01: WAVE-INTERNAL
```

xv. Add the SUBCONFIG by running the following commands:

```
DIRM ADD WAVWKL-4 BUILD ON ITSOZVM4 IN WAVEWRKL
DIRM ADD WAVWKC-4 BUILD ON ITSOZVM4 IN WAVEWRKC
```

5. Format the 191 and 399 minidisks of WAVEWRKS by completing the following steps. This task should be done in every system in the SSI cluster (in our environment, this procedure must be done in each member of the SSI cluster (ITSOZVM1, ITSOZVM2, ITSOZVM3, and ITSOZVM4):
   a. Log in to the MAINT user.
   b. Link to the 191 minidisk of WAVEWRKS with WRITE permissions by running the following command:

      ```
      LINK WAVEWRKS 191 1191 WR
      ```

   c. Format the minidisk by running the following command:

      ```
      FORMAT 1191 J
      ```

   d. When asked for a label, enter WAV191.
   e. Detach the disk by running the following command:

      ```
      rel J (DET
      ```

   f. Link to the 399 minidisk of WAVEWRKS with WRITE permissions by running the following command:

      ```
      LINK WAVEWRKS 399 1399 WR
      ```

   g. Format the minidisk by running the following command:

      ```
      FORMAT 1399 J
      ```

   h. When asked for a label, enter WAV399.
   i. Detach the disk by running the following command:

      ```
      rel J (DET
      ```

   j. Repeat steps a to i for each member in the SSI cluster.

All IBM Wave service machines identities are created, formatted, and ready to be used.

For more information about the SSI, see *Using z/VM v 6.2 Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8039 and *The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3*, SG24-8147.

## 4.1.8 RACF definitions

IBM RACF for z/VM provides a mechanism for the control of resources that are associated with the z/VM environment. These external controls expand z/VM security management and auditing capabilities in scope and granularity.

**Note:** This publication assumes that RACF for z/VM is installed and is active in the environment. Furthermore, it assumes that the environment being protected is a single z/VM system that uses a single RACF database, and that any pertinent product configuration steps to enable RACF control are correctly followed.

RACF commands can either be issued from a RACF session, or by using the RACF command line. The commands that are used in this section are in the format of the RACF command-line version.

### Configuring SMAPI to work with RACF

To configure SMAPI to work with RACF, complete the following steps:

1. Access your system through a 3270 emulator.
2. Log on as MAINT on the first SSI member.
3. Give VSMWORK1 the CONTROL authority to the z/VM minidisk (VMMDISK) that contains the `SYSTEM CONFIG` file (PMAINT CF0) by running the following commands:
   - **`RAC PERMIT PMAINT.CF0 CLASS(VMMDISK) ACC(CONTROL) ID(VSMWORK1)`**
   - **`RAC PERMIT MAINT.CF1 CLASS(VMMDISK) ACC(CONTROL) ID(VSMWORK1)`**
4. Give VSMWORK1 the CONTROL access to the generic class VMBATCH by running the following command:

   ```
   RAC PERMIT ** class(vmbatch) id(vsmwork1) access(control)
   ```

5. Configure SMAPI workers to read the TCPMAINT 198 disk by running the following commands:
   - **`RAC PERMIT tcpmaint.198 class(vmmdisk) acc(read) id(vsmguard)`**
   - **`RAC PERMIT tcpmaint.198 class(vmmdisk) acc(read) id(vsmwork1)`**
   - **`RAC PERMIT tcpmaint.198 class(vmmdisk) acc(read) id(vsmwork2)`**
   - **`RAC PERMIT tcpmaint.198 class(vmmdisk) acc(read) id(vsmwork3)`**
6. Change the default password expiration to your security standard by running the following command. In our environment, we used 186 days.

   ```
   RAC setropts password(interval(186))
   ```

   **Note:** If you already defined some of the RACF steps by using your own specific definition, skip step 6, or at least compare its values to the commands that are shown in the examples.

## Enabling RACROUTE

To enable RACROUTE, complete the following steps:

1. Enable the SMAPI service machines VSMREQI6, VSMREQIN, VSMREQIU, VSMEVSRV, DTCSMAPI, VSMWORK1, VSMWORK2, and VSMWORK3 to use RACROUTE services by running the commands that are shown in Figure 4-39.

```
RAC SETROPTS CLASSACT(FACILITY)
RAC RDEFINE FACILITY ICHCONN UACC(NONE)
RAC PERMIT ICHCONN CLASS(FACILITY) ID(VSMREQI6) ACCESS(UPDATE)
RAC PERMIT ICHCONN CLASS(FACILITY) ID(VSMREQIN) ACCESS(UPDATE)
RAC PERMIT ICHCONN CLASS(FACILITY) ID(VSMREQIU) ACCESS(UPDATE)
RAC PERMIT ICHCONN CLASS(FACILITY) ID(VSMEVSRV) ACCESS(UPDATE)
RAC PERMIT ICHCONN CLASS(FACILITY) ID(DTCSMAPI) ACCESS(UPDATE)
RAC PERMIT ICHCONN CLASS(FACILITY) ID(VSMWORK1) ACCESS(UPDATE)
RAC PERMIT ICHCONN CLASS(FACILITY) ID(VSMWORK2) ACCESS(UPDATE)
RAC PERMIT ICHCONN CLASS(FACILITY) ID(VSMWORK3) ACCESS(UPDATE)
RAC SETROPTS RACLIST(FACILITY)
```

*Figure 4-39 RACF commands*

2. Make the DTCSMAPI, VSMWORK1, VSMWORK2, and VSMWORK3 VMs exempt by running the commands that are shown in Figure 4-40.

```
RAC SETROPTS CLASSACT(VMXEVENT)
RAC RDEFINE VMXEVENT USERSEL.DTCSMAPI
RAC RALTER VMXEVENT USERSEL.DTCSMAPI ADDMEM(FOR.C/NOCTL)
RAC RALTER VMXEVENT USERSEL.DTCSMAPI ADDMEM(LINK/NOCTL)
RAC SETEVENT REFRESH USERSEL.DTCSMAPI
RAC RDEFINE VMXEVENT USERSEL.VSMWORK1
RAC RALTER VMXEVENT USERSEL.VSMWORK1 ADDMEM(FOR.C/NOCTL)
RAC RALTER VMXEVENT USERSEL.VSMWORK1 ADDMEM(LINK/NOCTL)
RAC SETEVENT REFRESH USERSEL.VSMWORK1
RAC RDEFINE VMXEVENT USERSEL.VSMWORK2
RAC RALTER VMXEVENT USERSEL.VSMWORK2 ADDMEM(FOR.C/NOCTL)
RAC RALTER VMXEVENT USERSEL.VSMWORK2 ADDMEM(LINK/NOCTL)
RAC SETEVENT REFRESH USERSEL.VSMWORK2
RAC RDEFINE VMXEVENT USERSEL.VSMWORK3
RAC RALTER VMXEVENT USERSEL.VSMWORK3 ADDMEM(FOR.C/NOCTL)
RAC RALTER VMXEVENT USERSEL.VSMWORK3 ADDMEM(LINK/NOCTL)
RAC SETEVENT REFRESH USERSEL.VSMWORK3
```

*Figure 4-40 RACF commands*

RACF should now allow SMAPI to do its job.

## Configuring IBM Wave to work with RACF

This section describes how to configure IBM Wave to work with RACF. This section assumes that you completed the steps in “IBM Wave service machines that are defined as Identities” on page 42.

Because the service machines are populated by auto detect, you must grant all the necessary RACF permissions (shown in Figure 4-41 on page 49) before you attempt to auto detect the system.

```
RAC PERMIT WAVEWRKS.191 ID(WAVEWRKC WAVEWRKL WAVEWR5S) CLASS(VMMD) ACC (ALTER)
RAC PERMIT WAVEWRKS.399 ID(WAVEWRKC WAVEWRKL WAVEWR5S) CLASS(VMMD) ACC (ALTER)
RAC PERMIT DIRMAINT.1DF ID(WAVEWRKS) CLASS(VMMDISK) ACC(ALTER)
RAC PERMIT AUTOLOG1.191 CLASS(VMMDISK) ID(WAVEWRKS) ACCESS(ALTER)
RAC PERMIT AUTOLOG2.191 CLASS(VMMDISK) ID(WAVEWRKS) ACC(ALTER)
RAC PERMIT DIAG088 CLASS(VMCMD) ID(WAVEWRKS WAVEWRKL WAVEWRKC) ACCESS(READ)
RAC PERMIT WAVEWRKS CLASS(VMBATCH) ID(FTPSERVE) ACCESS(CONTROL)
RAC RDEFINE VMXEVENT USERSEL.WAVEWRKS
RAC RDEFINE VMXEVENT USERSEL.WAVEWRKL
RAC RDEFINE VMXEVENT USERSEL.WAVEWRKC
RAC RALTER VMXEVENT USERSEL.WAVEWRKS ADDMEM(FOR.C/NOCTL FOR.G/NOCTL)
RAC RALTER VMXEVENT USERSEL.WAVEWRKL ADDMEM(FOR.C/NOCTL FOR.G/NOCTL)
RAC RALTER VMXEVENT USERSEL.WAVEWRKC ADDMEM(FOR.C/NOCTL FOR.G/NOCTL)
RAC SETEVENT REFRESH USERSEL.WAVEWRKS
RAC SETEVENT REFRESH USERSEL.WAVEWRKL
RAC SETEVENT REFRESH USERSEL.WAVEWRKC
RAC PERMIT WAVEWRKS CLASS(VMBATCH) ID(FTPSERVE) ACCESS(CONTROL)
```

*Figure 4-41 RACF definitions for WAVE service machines*

Set the passwords for the three IBM Wave service machines by using your password, and as non-expiring, during the installation, as shown in Figure 4-42.

```
RAC ALU WAVEWRKS PASS(PASSWORD) NOEXP
RAC ALU WAVEWRKL PASS(PASSWORD) NOEXP
RAC ALU WAVEWRKC PASS(PASSWORD) NOEXP
```

*Figure 4-42 RACF changing the password and setting no expiring*

**Note:** After the installation, you must set up the specific password rules that are defined by your organization.

For more information about the RACF, see *The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3*, SG24-8147.

## 4.1.9 SCSI DASD definitions

z/VM provides native support for SCSI disks for paging, spooling, and other system data, and supports the SCSI disk LUNs to z/VM as FBA-512 byte block DASD. SCSI disks can be used directly by z/VM guests when an FCP subchannel is dedicated to each guest.

### Emulated devices on SCSI disks

FBA emulation is a powerful capability that gives system administrators the best of both worlds: the flexibility of FCP-based SAN disks, which are combined with the manageability and security of traditional mainframe disks.

A SCSI device is defined to the z/VM system by specifying an EDEVICE statement in the SYSTEM CONFIG file (in the PMAINT CF0 parm disk) for a permanent definition or running **CP SET EDEVICE** for a dynamic definition. In our environment, we set the definition in the SYSTEM CONFIG file, as shown in the Figure 4-43.

```
/*********************************************************************/
/*          FCP DISKS                                                */
/*********************************************************************/
EDEV 3000 TYPE FBA ATTR 2145,
FCP_DEV B800 WWPN 500507680110BB91   LUN 0000000000000000,
FCP_DEV B801 WWPN 500507680110BC24   LUN 0000000000000000,
FCP_DEV BA00 WWPN 500507680140BB91   LUN 0000000000000000,
FCP_DEV BA01 WWPN 500507680140BC24   LUN 0000000000000000
EDEV 3001 TYPE FBA ATTR 2145,
FCP_DEV B800 WWPN 500507680110BB91   LUN 0001000000000000,
FCP_DEV B801 WWPN 500507680110BC24   LUN 0001000000000000,
FCP_DEV BA00 WWPN 500507680140BB91   LUN 0001000000000000,
FCP_DEV BA01 WWPN 500507680140BC24   LUN 0001000000000000
...
```

*Figure 4-43 SYSTEM CONFIG file with EDEVICE definitions*

An example of dynamic commands that are run to define the EDEVICE 300F in our environment is shown in Figure 4-44.

```
set edevice 300F type fba attr 2145 fcp_dev b800 wwpn 5000507680110BB91 lun 000F000000000000
set edevice 300F type fba attr 2145 add path fcp_dev b801 wwpn 5000507680110BC21 lun 000F000000000000
set edevice 300F type fba attr 2145 add path fcp_dev ba00 wwpn 5000507680110BC21 lun 000F000000000000
set edevice 300F type fba attr 2145 add path fcp_dev ba01 wwpn 5000507680110BC21 lun 000F000000000000
```

*Figure 4-44 Dynamic definition of EDEVICE definitions*

The emulated device (EDEV) is assigned a real device number and becomes associated with a SCSI disk through an FCP device, which also has a real device number. The device number of an emulated FBA disk must not conflict with the device number of any real device in the z/VM system. All emulated FBA disks and all real devices must have unique real device numbers, even though the emulated FBA disks are not real devices.

IBM Wave supports EDEV management. You can use the user interface to create and modify an EDEV and differentiate between real device volumes and EDEV volumes. Figure 4-45 represents the icon of an EDEV volume, as shown in IBM Wave.

*Figure 4-45 EDEV icon in IBM Wave*

For more information about how to make EDEVICES available to IBM Wave, see 6.3, “Managing DASD” on page 196.

For more information about SCSI, see *Using z/VM v 6.2 Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8039.

## 4.2 IBM Wave installation

After all the prerequisites for z/VM are complete and a Linux system that matches the requirements is created, the installation of IBM Wave is easy. The process involves transferring the IBM Wave RPM file to the target Linux system and installing the RPM.

If you already installed IBM Wave V1.1 Fix Pack 10, go to 4.2.2, “IBM Wave V1.2 installation” on page 54 to install the IBM Wave V1.2 Fix Pack 0.5 on top of your current installation. Make a backup of your IBM Wave database first.

### 4.2.1 IBM Wave V1.1.0 base installation

In our example, we create a SUSE Enterprise Linux Server V11 SP4 Linux VM and install the IBM Wave V1.0 RPM. We transferred the RPM file in binary mode to the WAVESRV VM, and then ran the command that is shown in Figure 4-46.

```
rpm -ivh IBM-Wave-1.1.0.1_fp001.s390x.rpm
```

*Figure 4-46 Install IBM Wave V1.0*

The installation of IBM Wave has dependencies that must be satisfied, as shown in Figure 4-47.

```
wavesrv:/tmp # rpm -ivh IBM-Wave-1.1.0.1_fp001.s390x.rpm
Preparing...                ########################################### [100%]
Detected suse-11.4
which: no java in (/sbin:/bin:/usr/sbin:/usr/bin:/usr/X11R6/bin)
Error: Failed Dependencies:
         mysql >= 5.0.0 is required
         apache2-prefork >= 2.0.49 is required
         java >= 1.0.6 is required and must be in PATH
         nfs-kernel-server >= 1.2.0 is required and must be in PATH
Exiting...
error: %pre(IBM-Wave-1.10-1.s390x) scriptlet failed, exit status 1
error:   install: %pre scriptlet failed (2), skipping IBM-Wave-1.10-1
```

*Figure 4-47 IBM Wave installation dependencies*

You must install all prerequisites before trying to install the IBM Wave RPM. Because our example IBM Wave server is SUSE Linux Enterprise Server V11, the command that is used to install them is the **zypper** command, as shown in Figure 4-48.

```
zypper install mysql
zypper install apache2-prefork
zypper install nfs-kernel-server
zypper install java
```

*Figure 4-48 IBM Wave installation of the packages*

**Note:** If the IBM Wave server is Red Hat Enterprise Linux Server, you must use YUM to install the packages.

After the dependencies are installed, you can rerun the IBM Wave installation command, as shown in Figure 4-46 on page 51. After the command finishes, you should see a message that indicates a successful installation, as shown in Figure 4-49.

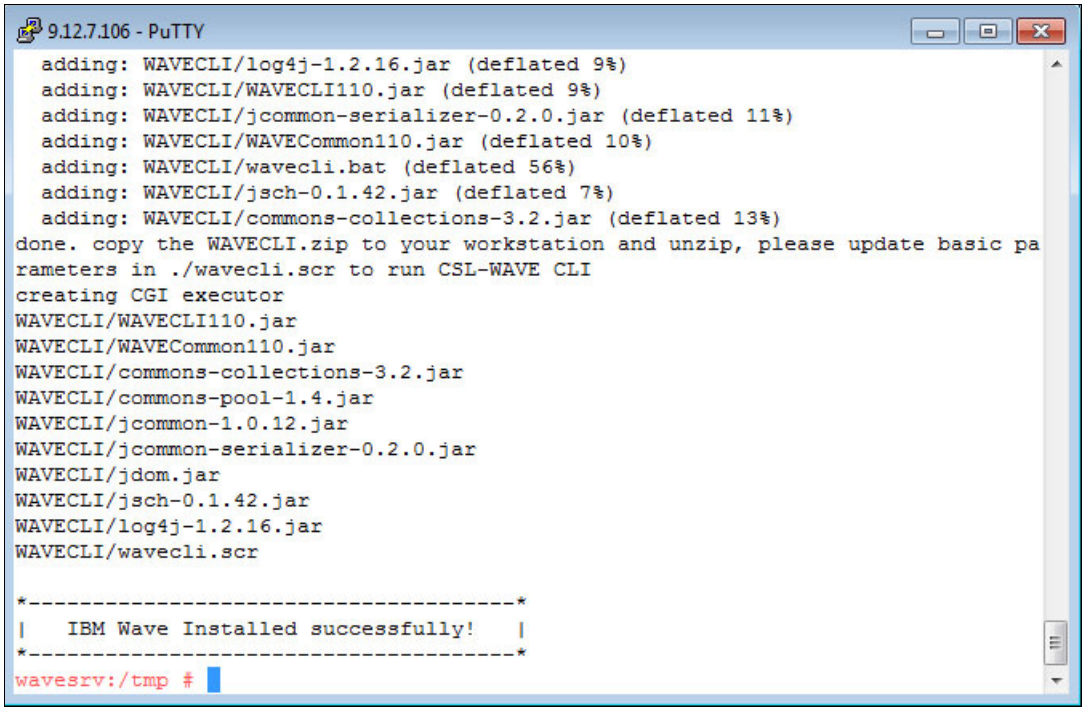


*Figure 4-49 RPM Installation complete*

You must now define the IBM Wave administrator by completing the following steps:

1. Start IBM Wave V1.1 in your browser by entering the IP address of the IBM Wave server in the location bar, as shown in Figure 4-50 on page 53.

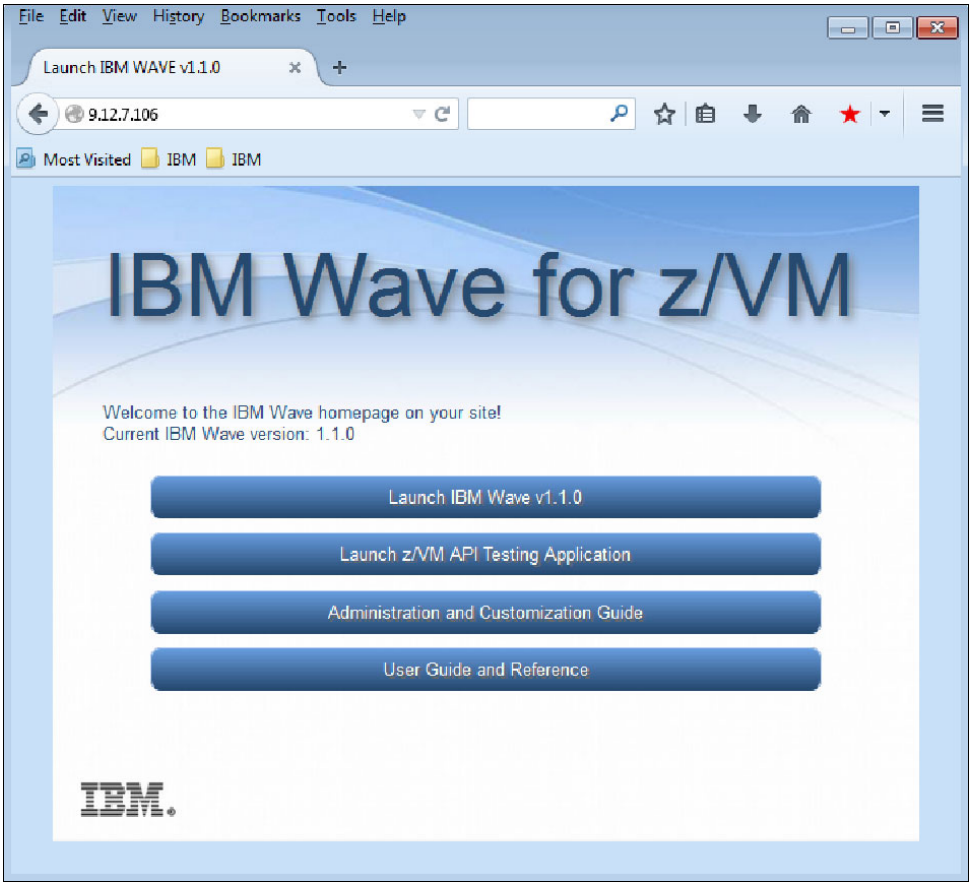


*Figure 4-50 IBM Wave V1.1 first start*

2. Save the file (`WAVE.jnlp`) that contains the Java WebStart application, as shown in Figure 4-51.

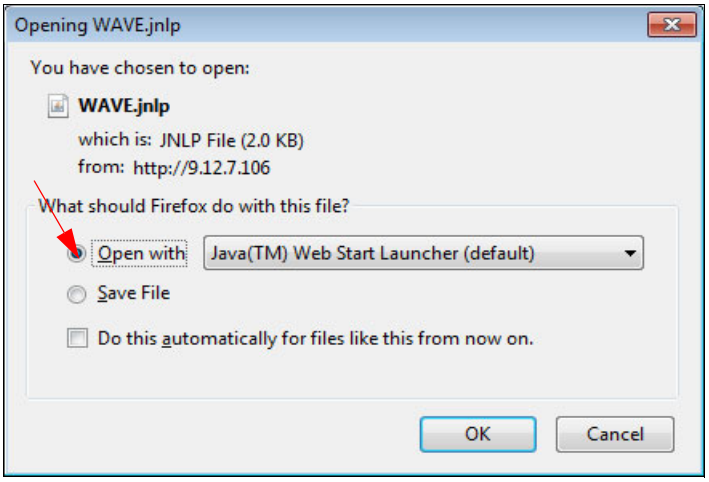


*Figure 4-51 WAVE.jnlp*

3. Accessing IBM Wave for the first time after a new installation opens the Welcome window, as shown in Figure 4-52. Use this window to create a superuser ID to perform all IBM Wave tasks. This ID is usually named waveadm.

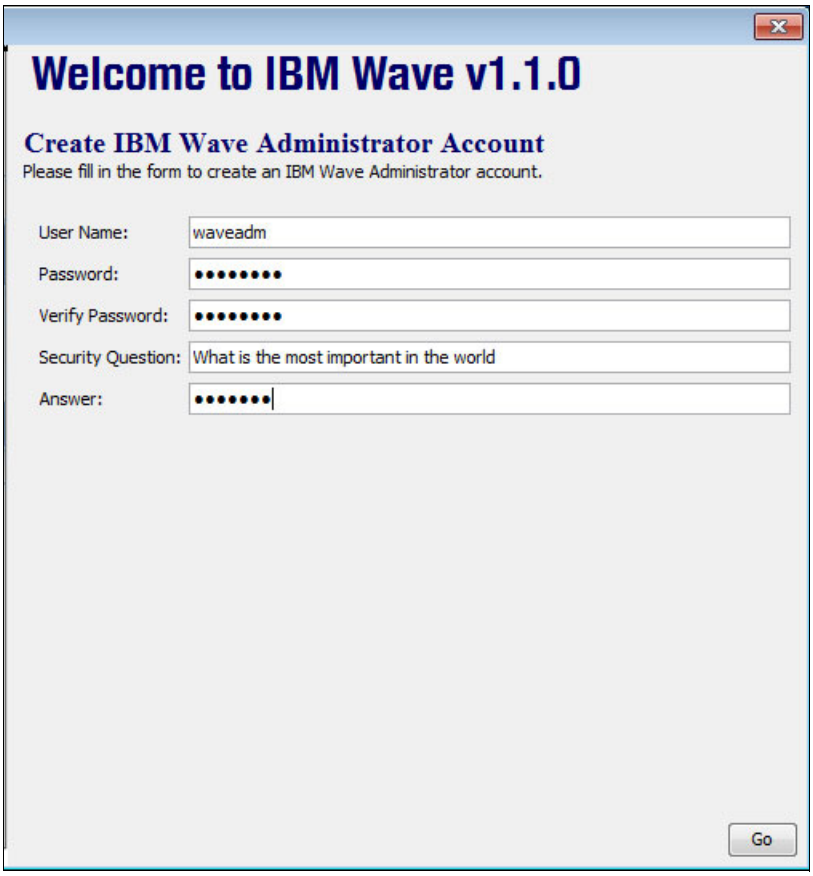


*Figure 4-52 IBM Wave first user ID creation*

Now, you can continue with the IBM Wave V1.2 installation.

## 4.2.2 IBM Wave V1.2 installation

After the installation of the Base Wave V1.1.0 RPM, or if you already installed IBM Wave V1.1 Fix Pack 10 in your environment, you must install IBM Wave V1.2 Fix Pack 0.5 to upgrade to IBM Wave V1.2.

In our environment, we transfer the compressed file in binary mode to the WAVESRV VM, and then run **tar** to extract the files from the tar archive, as shown in Figure 4-53.

```
tar -xvf IBM-Wave-1.2.0.00.5.tar
```

*Figure 4-53 Unpack IBM Wave V1.2*

After you extract the tar file (Figure 4-54 on page 55) into a temporary directory, IBM Wave V1.2.0.00.5 is ready to be installed.

```
wavesrv:/tmp # tar -xvf IBM-Wave-1.2.0.00.5.tar
IBM-Wave-1.2.0.00.5/
IBM-Wave-1.2.0.00.5/GUI/
IBM-Wave-1.2.0.00.5/GUI/appl/
IBM-Wave-1.2.0.00.5/GUI/appl/WAVE120.jar
IBM-Wave-1.2.0.00.5/GUI/Common/
IBM-Wave-1.2.0.00.5/GUI/Common/WAVECommon120.jar
IBM-Wave-1.2.0.00.5/GUI/index.html
IBM-Wave-1.2.0.00.5/GUI/style.css
IBM-Wave-1.2.0.00.5/GUI/TestAPI.jnlp
IBM-Wave-1.2.0.00.5/GUI/WAVE.jnlp
...
IBM-Wave-1.2.0.00.5/WAVECLI/
IBM-Wave-1.2.0.00.5/WAVECLI/WAVECLI120.jar
IBM-Wave-1.2.0.00.5/WAVECLI/networkUpdate.scr
IBM-Wave-1.2.0.00.5/WAVECLI/packWAVECLI
IBM-Wave-1.2.0.00.5/WAVECLI/wavecli.bat
IBM-Wave-1.2.0.00.5/WAVECLI/wavecli.scr
IBM-Wave-1.2.0.00.5/WAVECLI/wavecli.scr.cgi
IBM-Wave-1.2.0.00.5/doUpdate.sh
IBM-Wave-1.2.0.00.5/WAVE-SP-Upgrader.jar
```

*Figure 4-54 IBM Wave - unpacking the tar file*

The decompressed IBM Wave file creates the `IBM-Wave-1.2.0.00.5` directory in `/tmp`. Go to `IBM-Wave-1.2.0.00.5` and start the IBM Wave V 1.2 installation by running **`./doUpdate.sh`**. You see the results that are shown in Figure 4-55.

```
wavesrv:/tmp/IBM-Wave-1.2.0.00.5 # ./doUpdate.sh
Detected suse-11.4
Stopping IBM-Wave Background Services...
Shutting down WAVEBackgroundServices...                              done
Stopping mysql...
Shutting down service MySQL                                          done
stopping apache2...
Shutting down httpd2 (waiting for all children to terminate)         done
Backing up old version files...
Checking levels and updates...
Enabling migration from legacy GA1 to GA2
Removing legacy Jars
Updating files
...
No user passwords required updates.
Exiting...
Restarting Background Services...
Starting WAVEBackgroundServices...                                   done
Compressing /usr/wave/backup-2015-08-06-10-02-36...
-------------------------------------------
|                                         |
|       IBM Wave Successfully Updated!    |
|                                         |
-------------------------------------------
wavesrv:/tmp/IBM-Wave-1.2.0.00.5 #
```

*Figure 4-55 IBM Wave V1.2 installation*

All of the IBM Wave V1.2 components are now running and waiting for a connection through a browser to set up the environment.

### 4.2.3 Logging in to IBM Wave V1.2

The next step in the installation process is to start a browser and open a connection to the IP address of the WAVESRV system, as shown in Figure 4-56.

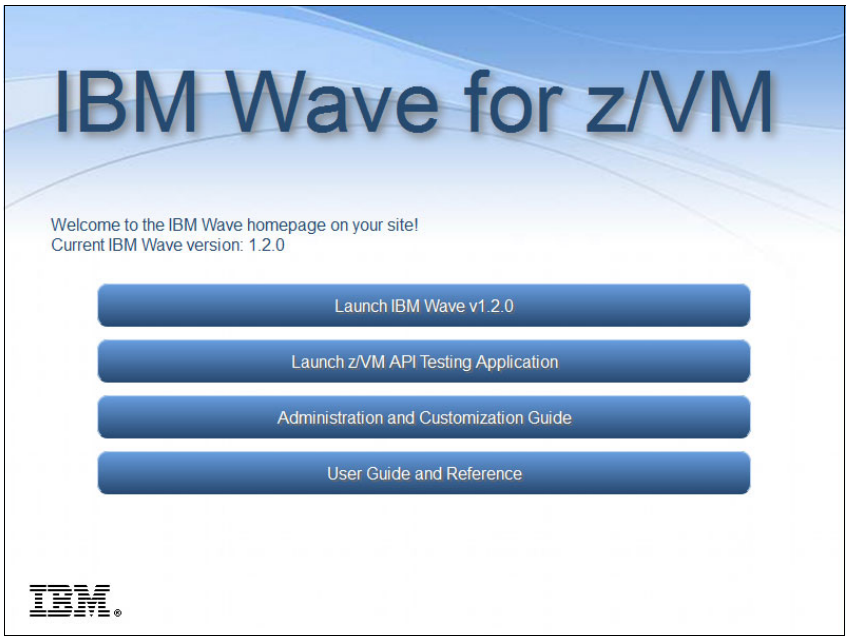


*Figure 4-56 IBM Wave Welcome window*

#### z/VM API testing

Before you start IBM Wave, test the z/VM APIs by clicking **Launch z/VM API Testing Application**.

The IP address that you enter is the same address as your z/VM system. The API user name and password corresponds to a user ID that is listed in the `VSMWORK1 AUTHLIST` file. When you click **Test**, you see a result similar to the example that is shown in Figure 4-57 on page 57.

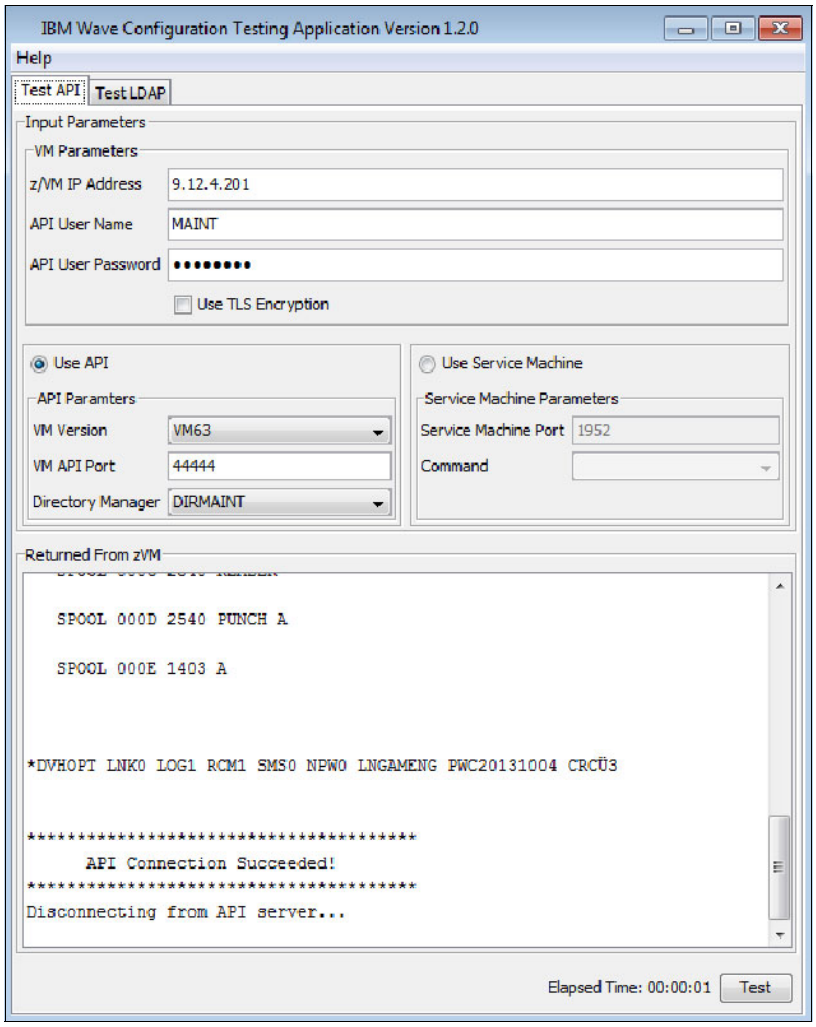


*Figure 4-57 Test API results*

After the API test is complete, you can log on to IBM Wave by using your IBM Wave administrative user ID. The main GUI opens. If you are performing an upgrade, you have IBM Wave V1.2 installed and ready to be use. If this is a new installation, you can continue by going to 4.2.4, “Adding a central processor complex” on page 58 and configuring your environment.

### 4.2.4 Adding a central processor complex

Now that IBM Wave is installed, your next task is to define a central processor complex (CPC) that hosts a z/VM system to be managed. Complete the following steps:

1. In the main window, right-click inside the **Hardware Viewer** pane. You see the Add New CPC option, as shown in Figure 4-58.

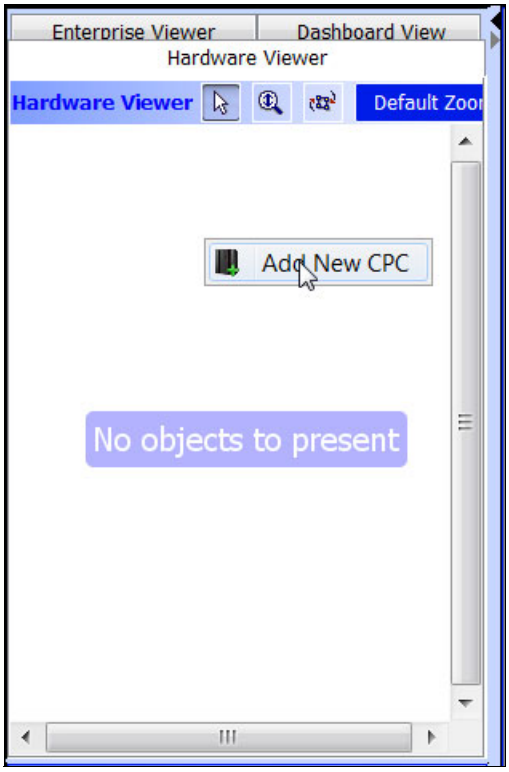


*Figure 4-58 Add New CPC*

2. Click **Add New CPC**. A window opens where you can enter information about the new CPC, as shown in Figure 4-59 on page 59.

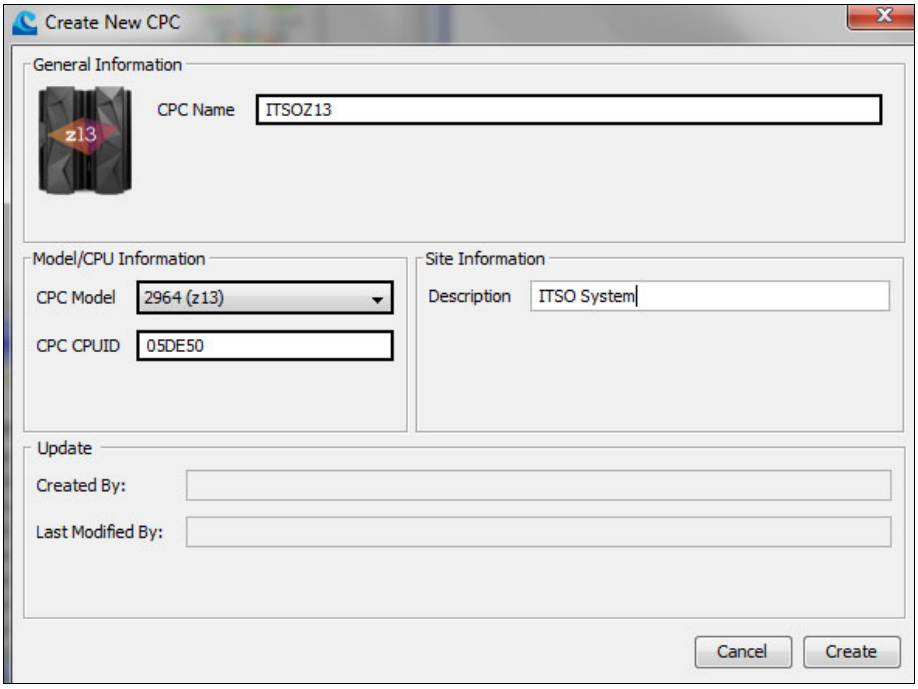


*Figure 4-59 Create a CPC*

The first field in the window is where you enter the name for this CPC. This name can be any name that you choose. Select the **CPC Model** from the CPC Model drop-down menu. Paste the CPU serial number for the CPC into the CPC CPUID field. You might also want to enter a description of the system in the Description field.

3. After all of the information is entered, click **Create**. If no errors are encountered, a success message is displayed.

4. Click **OK** and return to the main window. A graphic that represents the CPC that was added is shown in the **Hardware** Viewer section of the main window, as shown in Figure 4-60.


*Figure 4-60 Hardware Viewer CPC graphic*

## 4.2.5 Defining a system

To define a z/VM system on the CPC for IBM Wave, an external managed entity must be created for this system. Complete the following steps:

1. From the main window, click **Administration** at the top of the window, then select **Site Management**, and **External Entities Manager**, as shown in Figure 4-61 on page 61.

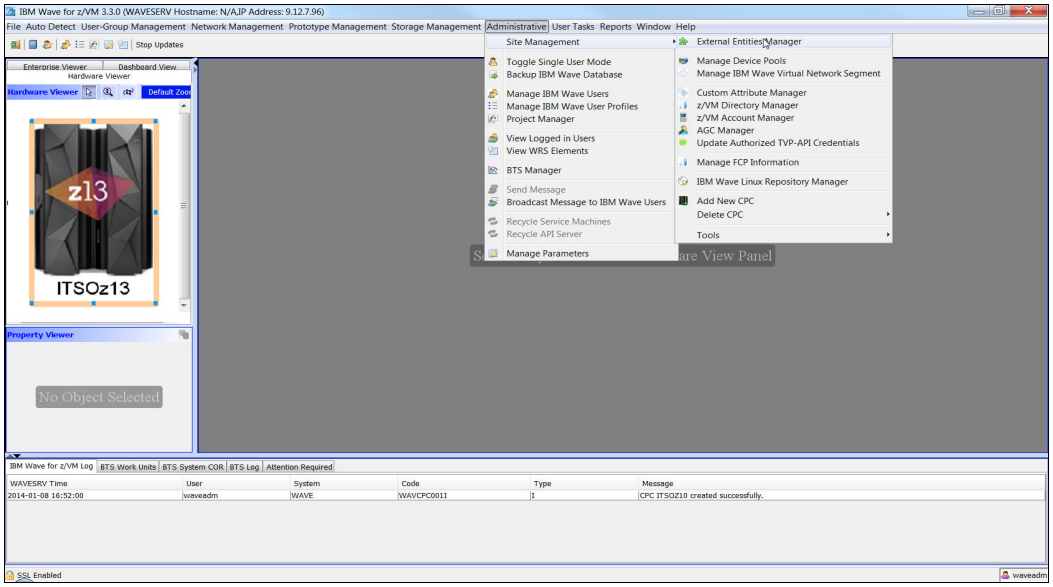


*Figure 4-61 Main window*

2. A window opens that lists any external entities that are defined, as shown in Figure 4-62. If you did not yet define any entities, the window is blank.

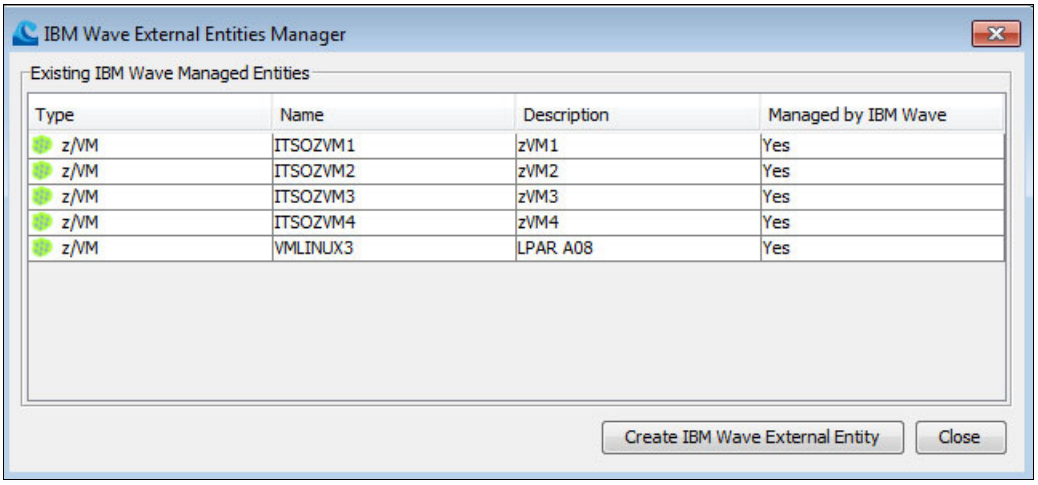


*Figure 4-62 External Entities Manager window*

3. Click **Create IBM Wave External Entity** to add an entity to IBM Wave. A window opens where you enter the name and type of entity that is added, as shown in Figure 4-63.

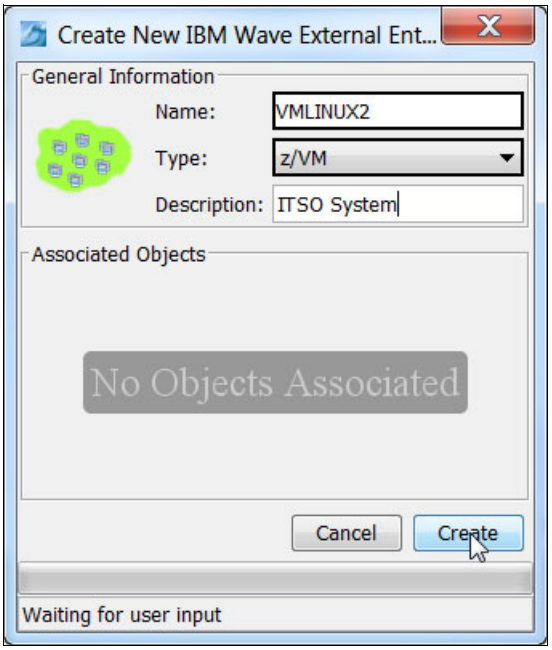


*Figure 4-63 External Entity VMLINUX2*

4. Enter the name of the LPAR that you want IBM Wave to manage. In this environment, we use VMLINUX2 as the name of the entity, and then select **z/VM** as the type from the drop-down menu. Add a description in the Description field (in this example, ITSO System).

   After all of the information is added, click **Create**. The system submits the add request to the background task scheduler (BTS) and opens a window with the information. Click **OK**, and when the add request is complete, a new external entity is shown in the External Entities Manager window, as shown in Figure 4-64.

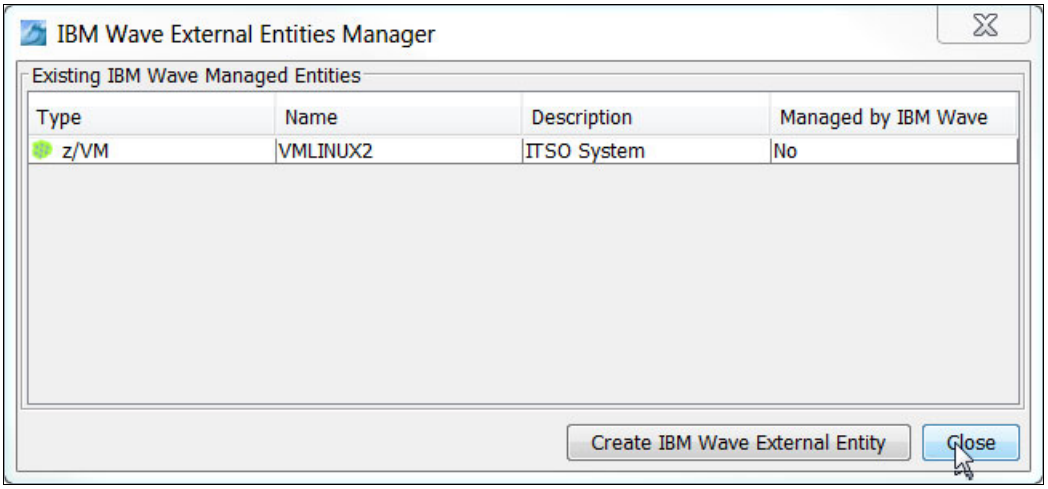


*Figure 4-64 External Entity Add complete*

5. Click **Close** to close the External Entities Manager window and return to the main window. Now that a type z/VM external entity is created, you can add a z/VM system for IBM Wave. Select **CPC** in the Hardware viewer and then right-click it. Click **More Actions** → **Add New System**, as shown in Figure 4-65.

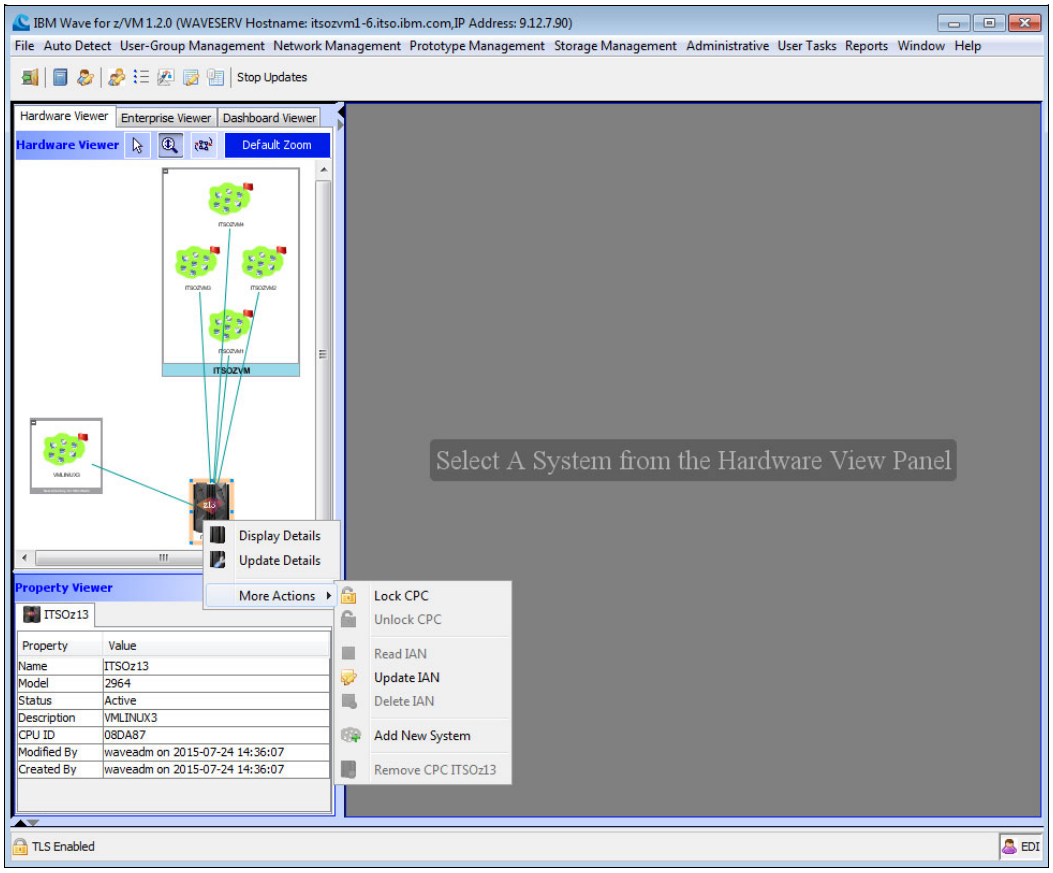


*Figure 4-65 Add New System window*

A window opens in which you enter information about the new z/VM system to be managed, as shown in Figure 4-66.

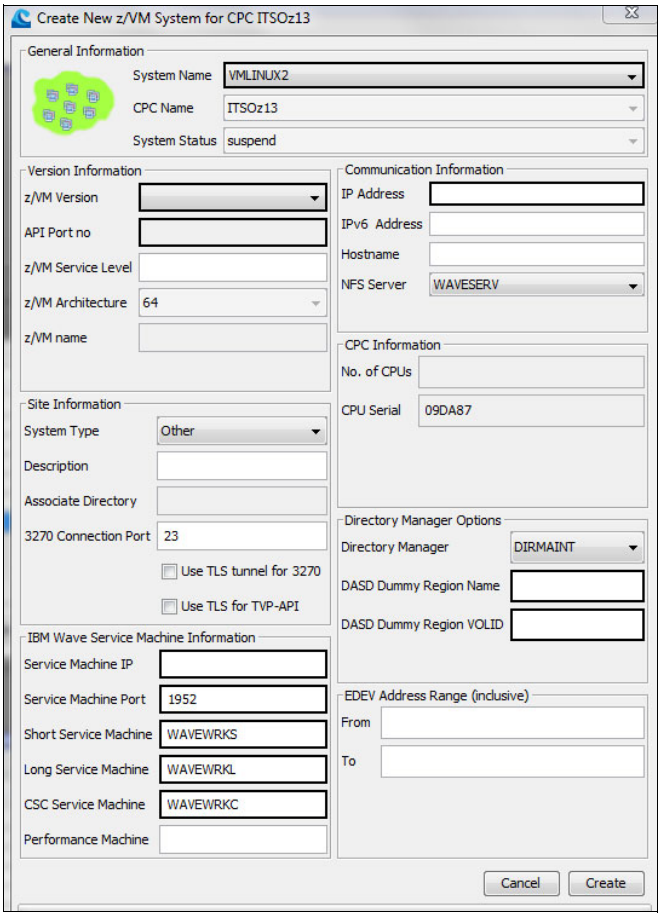


*Figure 4-66 New z/VM system information window*

The Add New System window opens with some of the information already entered. The System Name and CPC Name were taken from the External Entity definition and CPC Object. The IBM Wave Service Machine information contains typical defaults. Figure 4-67 on page 65 shows the window with the information that was added from our lab environment.

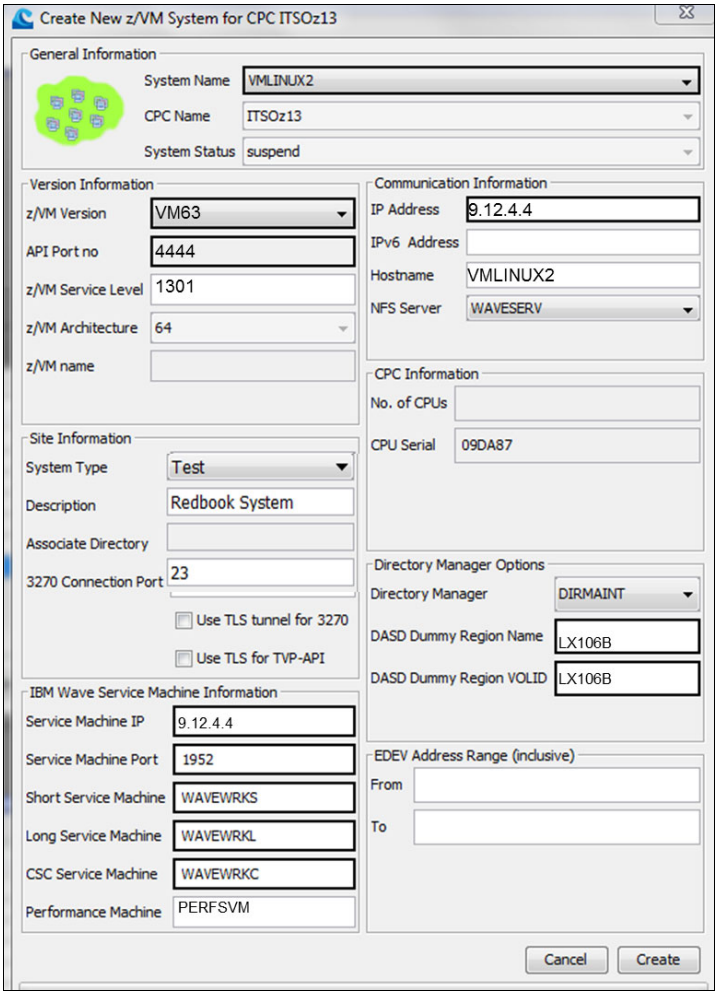


*Figure 4-67 Completed Create System window*

The IP address that is entered is the same that is associated with the z/VM TCP/IP stack because the IBM Wave worker machines listen on ports that are hosted by the z/VM TCP/IP stack. You enter the name of a DirMaint DASD Dummy Region and the VOLID for that device that is defined in the DirMaint `EXTENT CONTROL` file. All of this information is defined in the worksheets in Chapter 3, “Planning and design” on page 13 for easy reference.

If you want to use Performance Toolkit, you should enter PERFSVM as the Performance Machine.

6. After all of the information is entered, click **Create**. The system object is created and a window opens indicating the auto detect process is about to begin, as shown in Figure 4-68.

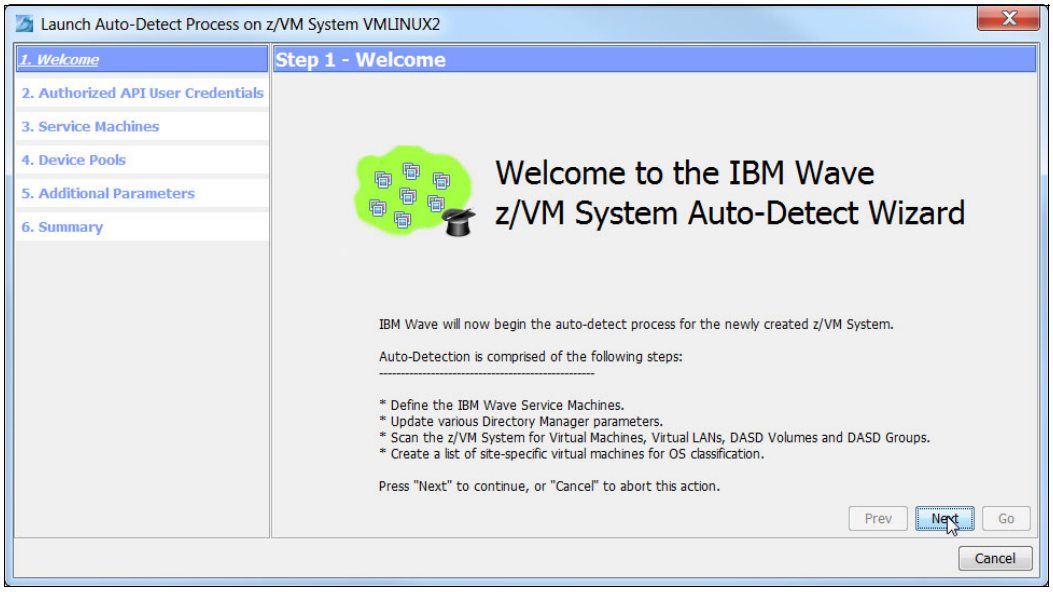


*Figure 4-68 Auto Detect window*

## 4.2.6 Auto detecting a system

The IBM Wave Auto Detect process creates the worker machines and then discovers as much information as possible from the z/VM system. This information discovery consists of all defined VMs, DASD information, networking information, and so on.

Complete the following steps:

1. To start the Auto Detect process, click **Next**. As shown in Figure 4-69 on page 67, a window opens where you enter the user name and password of the Authorized API User. (This is a user name that is defined in the VSMWORK1 `AUTHLIST` file).

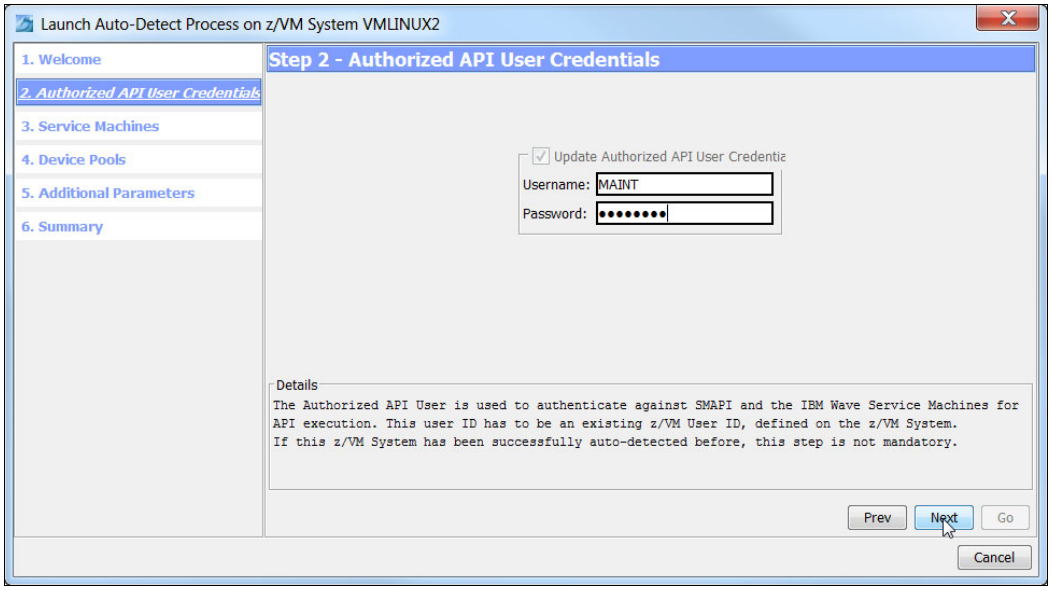


*Figure 4-69 Authorized API User Credentials window*

2. Click **Next**. A window opens with the information and actions that are related to the IBM Wave Service Machines, as shown in Figure 4-70.

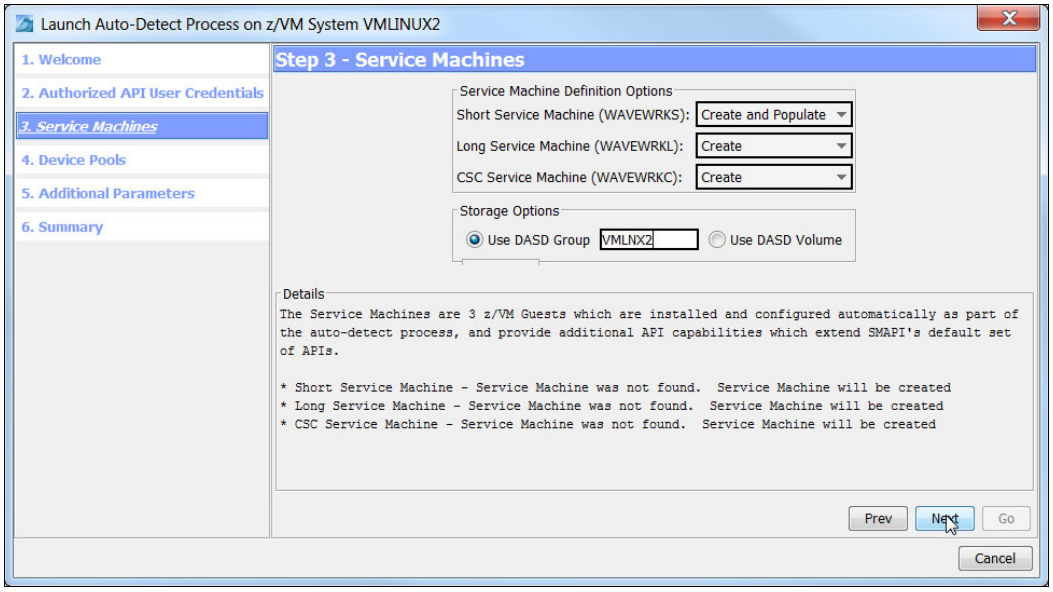


*Figure 4-70 Service Machines*

Because the IBM Wave Service Machines do not exist on the system, the default action is to Create and then Populate with code. The WAVEWRKS VM contains all the executable code. WAVEWRKL and WAVEWRKC link to the minidisk that is owned by WAVEWRKS to access the code. These service VMs increase the capability that is provided by the z/VM SMAPI. The WAVEWRKS VM is used for short duration requests and the WAVEWRKL VM is used for long duration (such as processing the `USER DIRECT` file). WAVEWRKC is used for cross-system cloning operations.

3. In the field that is next to Use DASD Group, enter the name of the group that is defined in `EXTENT CONTROL` to be used for the definition of the Service Machine minidisks. Click **Next**. A window opens with the device pool information, as shown in Figure 4-71.

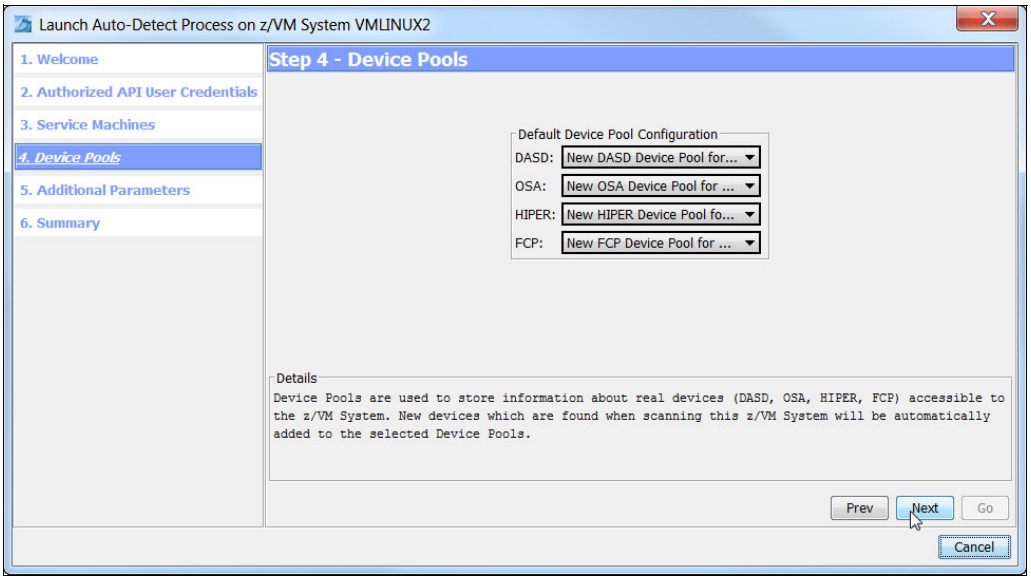


*Figure 4-71 Device Pools window*

4. Because information for this system is being acquired, the default action is to create device pools for all of the various pool types (DASD, OSA, HiperSockets, and FCP). You can accept these defaults and click **Next**. A window opens in which you enter other parameter information, as shown in Figure 4-72.

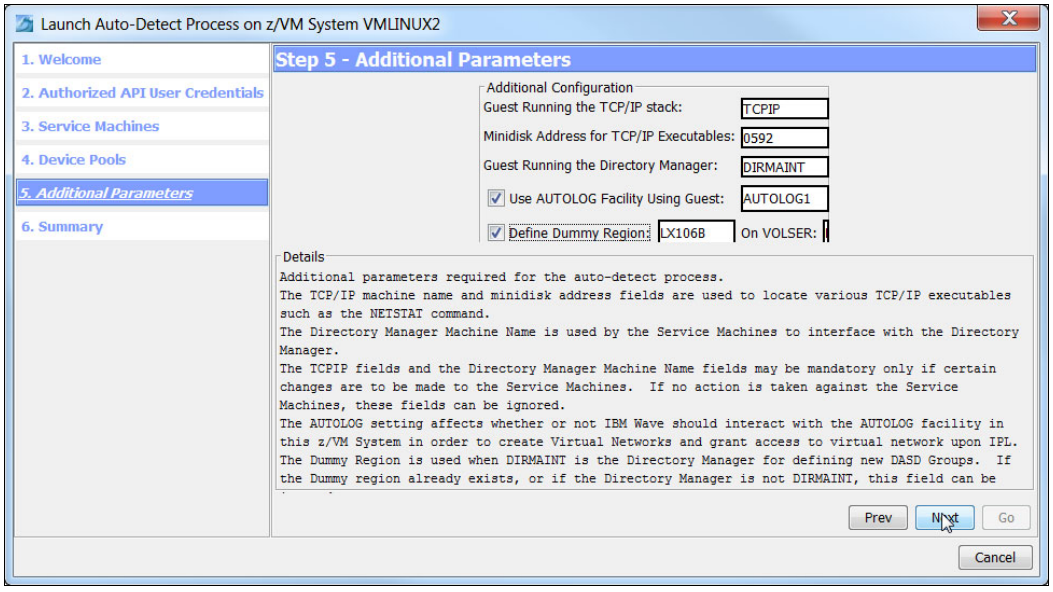


*Figure 4-72 Additional Parameters window*

The Additional Parameters window displays the typical information for an unmodified z/VM system. If you changed the name of the TCP/IP service machine from the IBM default of TCP/IP to something else, the name that you used in this window. Similarly, if you moved the TCP/IP client executable files from the default location of TCPMAINT 592, specify the new location in this window.

The only non-default piece of information is the specification of a Region name (from the `EXTENT CONTROL` file) that IBM Wave can use to create a Dummy Region. The Dummy Region is a one cylinder region that is never physically allocated. It exists to handle the limitation of a Group definition in `EXTENT CONTROL` and it cannot be empty. In our lab environment, we specified the name LX106B, which is the name of a Region in `EXTENT CONTROL` and the Volume Serial number of the device on which the region is defined.

5. Click **Next** and a summary window opens, as shown in Figure 4-73.

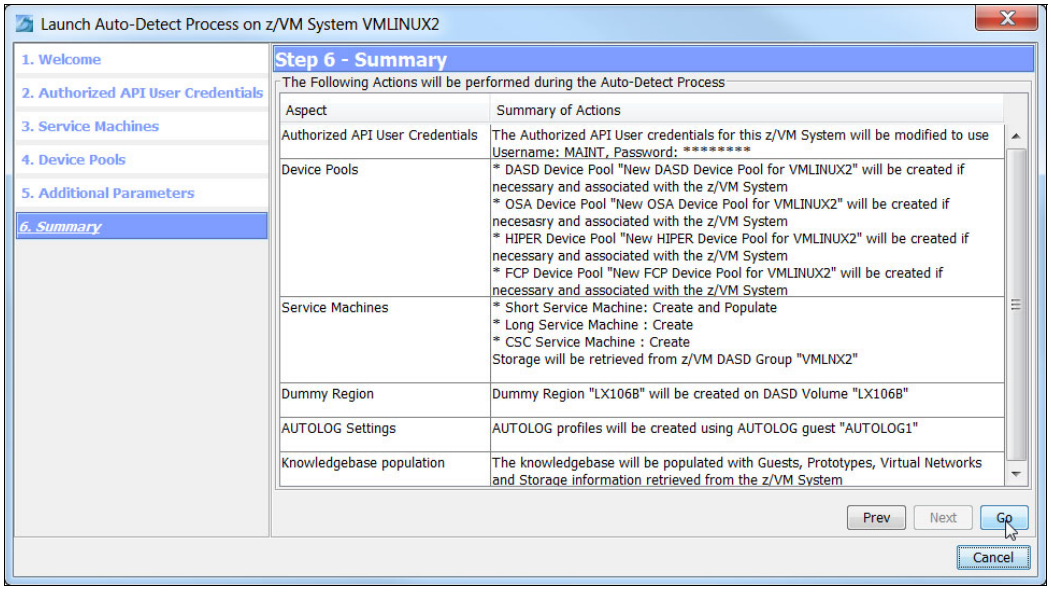


*Figure 4-73 Summary window*

6. If all of the information is correct, click **Go**. A confirmation window opens, as shown in Figure 4-74.

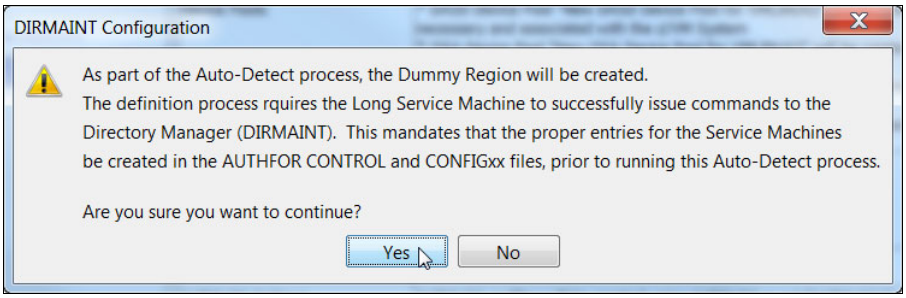


*Figure 4-74 DIRMAINT Warning window*

7. Because IBM Wave uses the long duration IBM Wave Service Machine (WAVEWRKL) to create the Dummy Region, this warning reminds you that the IBM Wave Service Machines must be authorized with DirMaint to run commands. If you completed the steps in 4.1.4, “DirMaint” on page 30 for updating and activating the changes to AUTHFOR CONTROL and the CONFIGxx DATADVH file, you completed the steps to which this warning is referring. Click **Yes** to continue with the auto detect process.

As the auto detect process runs, a BTS Work Unit shows a progress bar in the IBM Wave viewers section of the main window, as shown in Figure 4-75.

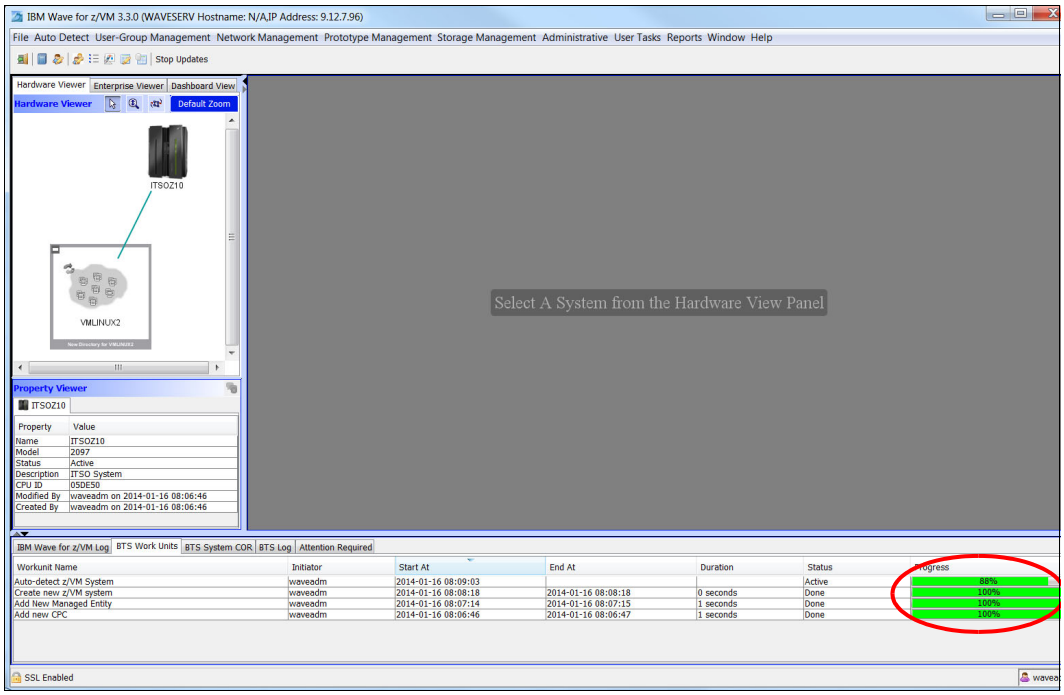


*Figure 4-75 Auto Detect Running window*

The details that are associated with this work unit can be viewed by double-clicking the row in the table under the **BTS Work Unit** tab. A Work Unit Details window opens, as shown in Figure 4-76 on page 71.

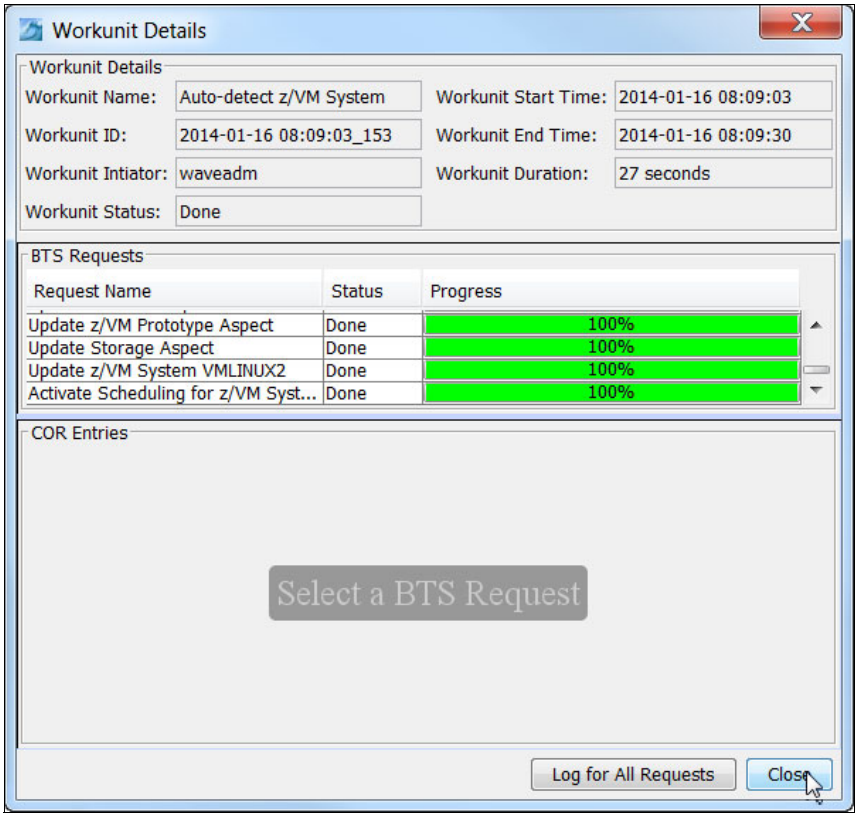


*Figure 4-76 Auto Detect Work Units window*

When the auto detect work unit completes, the system might still be processing the data that was gathered. A progress message is updated in the System Viewer portion of the main window, as shown in Figure 4-77.

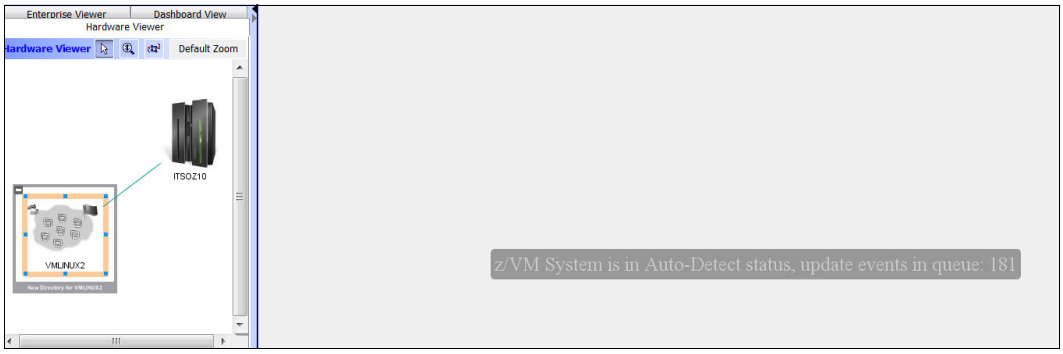


*Figure 4-77 System Auto Detect Status window*

When all of the queued update events are processed, the z/VM groups are displayed for the system (VMLINUX2), as shown in Figure 4-78.

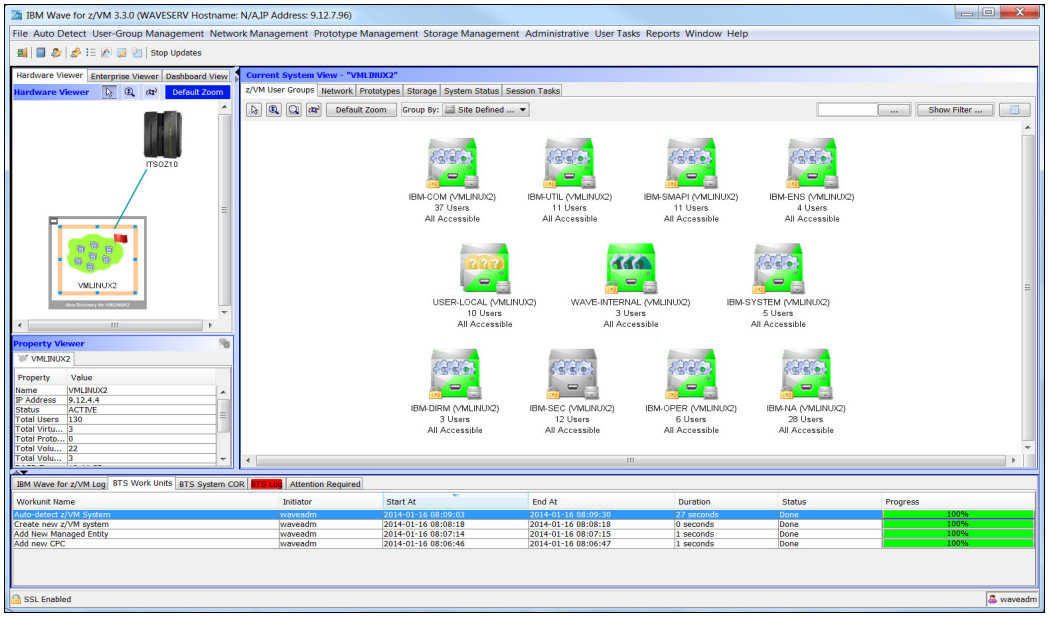


*Figure 4-78 Auto Detect Complete Site Defined Groups window*

## 4.2.7 How to correct alerts from Auto Detect

The System Status tab might show red in your system. In our system, it showed red for VMLINUX2, which indicates that there are some items that must be addressed. Clicking the **System Status** tab changes to that window in the system viewer, as shown in Figure 4-79.

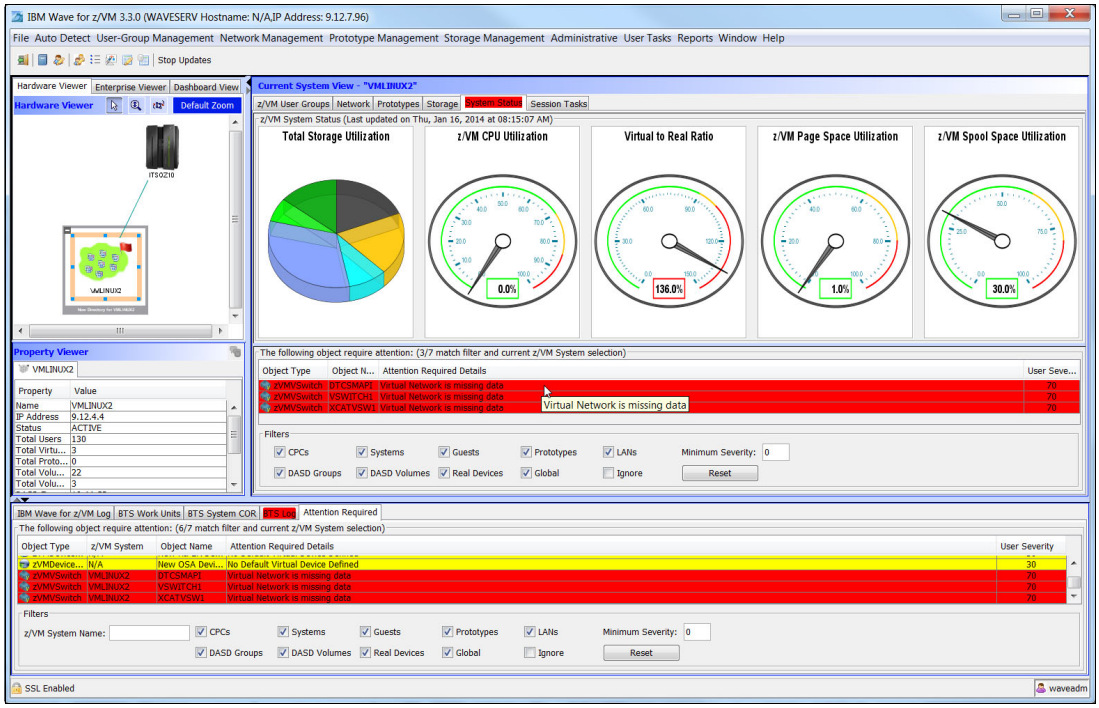


*Figure 4-79 System Status Tab window*

The red alert lines are displayed because IBM Wave does not know what default virtual NIC address should be used when it is connecting VMs to the three VSWITCH objects that were discovered on the system. This problem can be corrected from the Network tab in the system viewer.

## Fix alerts for VSWITCH

From the System Status tab, double-click the red alert for the first VSWITCH (DTCSMAPI), which opens the Network display that focuses on the DTCSMAPI VSWITCH icon, as shown in Figure 4-80.

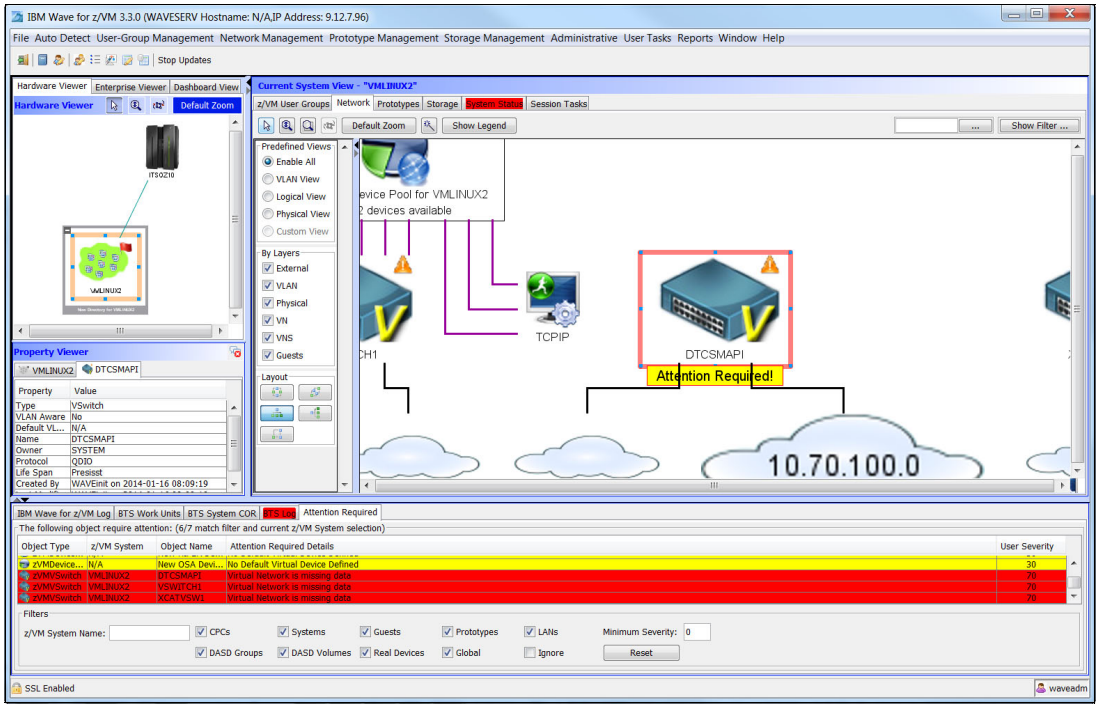


*Figure 4-80 Add Default NIC*

To update the information for this VSWITCH, right-click the icon and select **Update Information**, as shown in Figure 4-81.

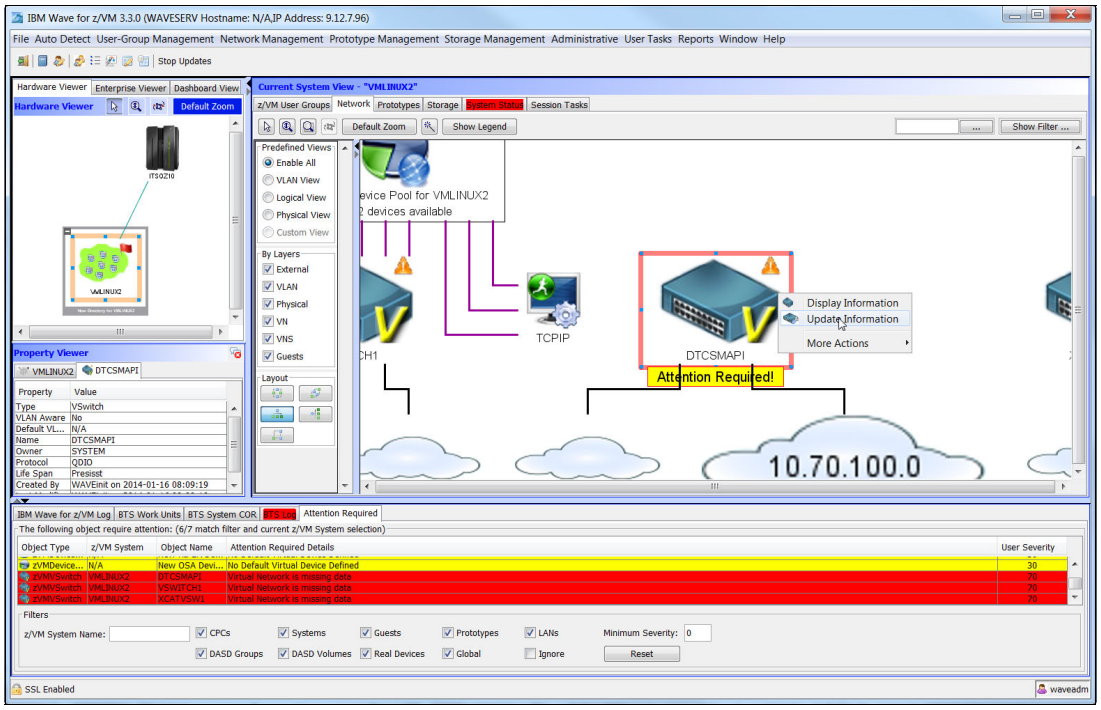


*Figure 4-81 Add Default NIC Navigation window*

Click **Update Information** and an information window opens, as shown in Figure 4-82.

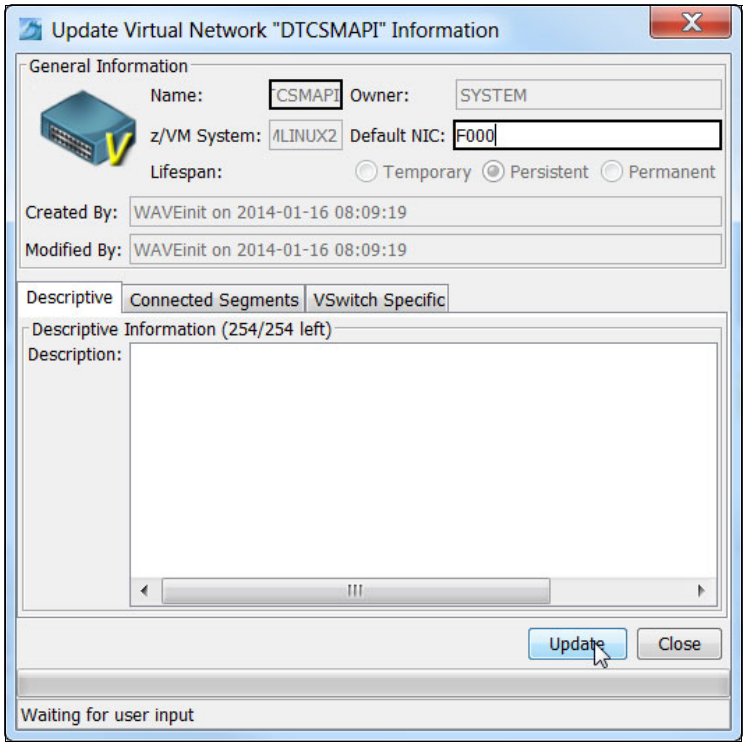


*Figure 4-82 Add Default NIC Data window*

To remove the alert, enter a virtual address in the Default NIC field. We entered F000 as the default NIC for this VSWITCH. If a guest is connected to a virtual network through this Virtual Network Segment (VNS) by using IBM Wave, the **NICDEF** statement in the guest's User Directory specifies F000 as the virtual NIC address, and all Linux configuration files refer to the F000 virtual address.

Click **Update** to complete the process. The properties for the VSWITCH are updated, and you are returned to the Network pane of the main window. We repeat this process on the remaining two VSWITCHes to remove the alert condition from them. When the updates are complete on all VSWITCHes, the icons in the network display no longer show the attention alert condition.

## Updating the virtual to real ratio threshold

Because memory is constrained in our lab environment, the virtual to real status gauge also shows red. Because we intend to run this system with a fairly high level of memory that is overcommitted, we can update the threshold so that IBM Wave does not display this item as an alert condition.

The threshold values are part of the global parameters that are maintained by IBM Wave. To modify these parameters, click **Administration** on the menu bar, then select **Manage Parameters**, as shown in Figure 4-83.

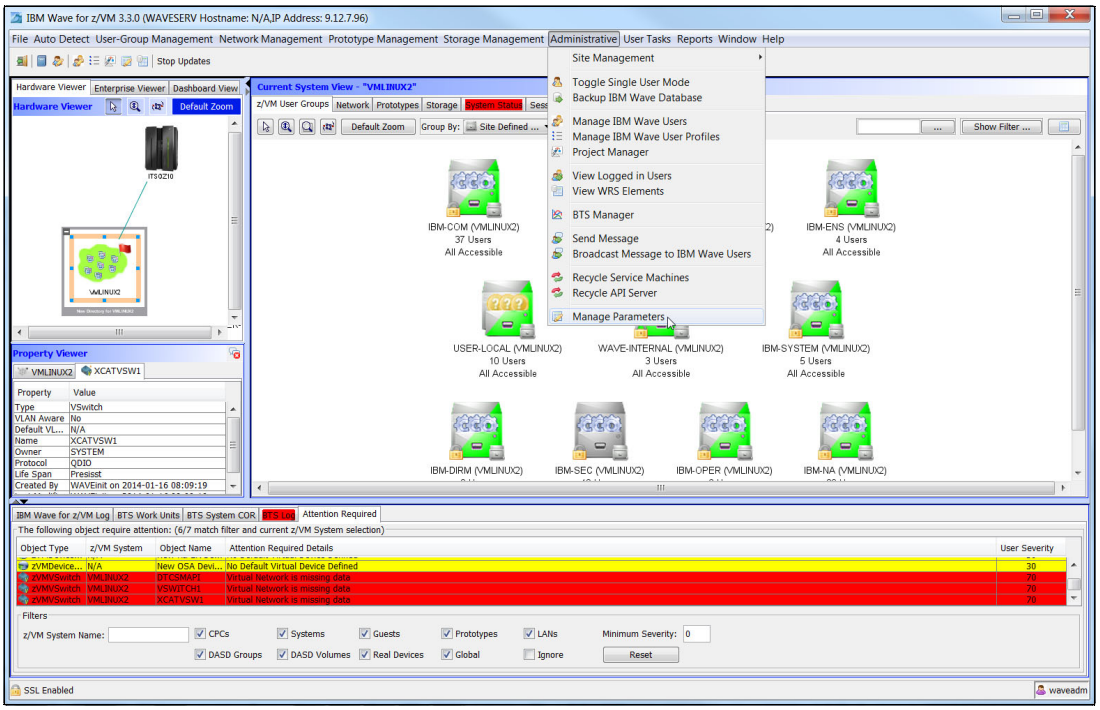


*Figure 4-83 Manage Parameters*

Clicking **Manage Parameters** opens the window that is shown in Figure 4-84.

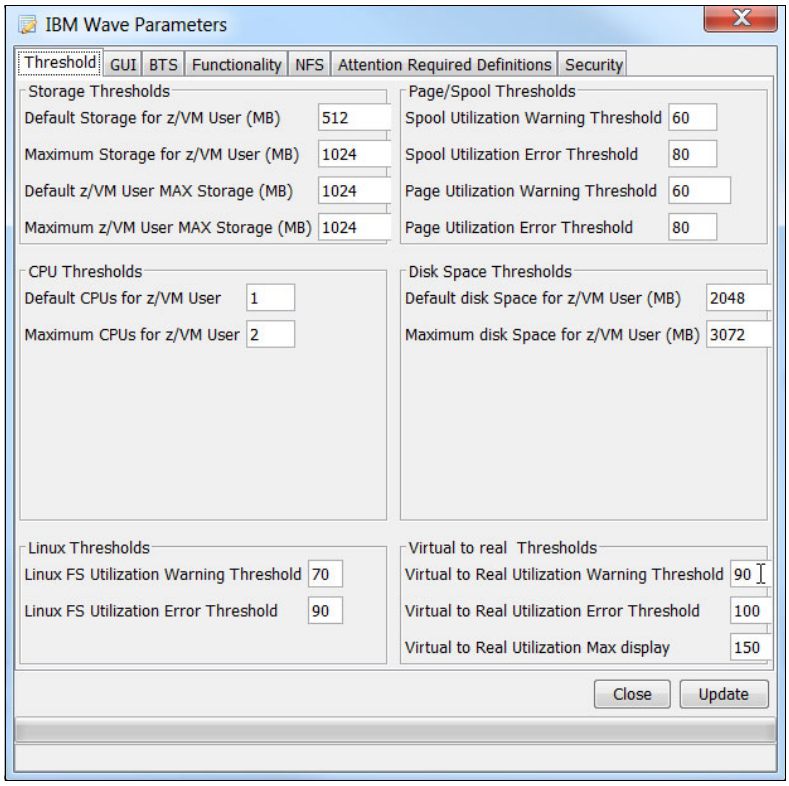


*Figure 4-84 Virtual To Real Threshold default settings*

We modified the defaults for our environment, as shown in Figure 4-85.

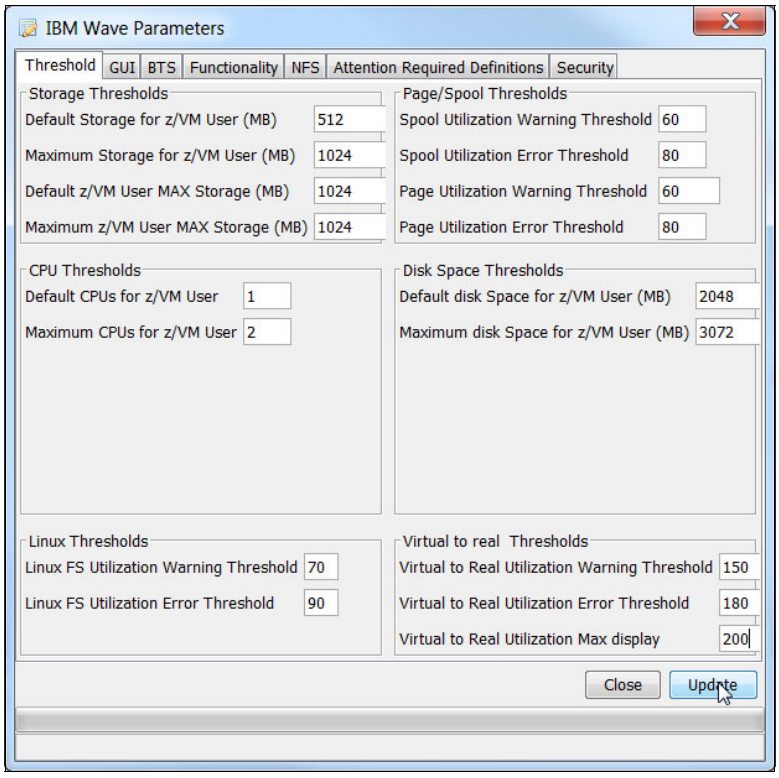


*Figure 4-85 Updated Virtual To Real Threshold settings*

After new values are entered for the thresholds, click **Update** to save and activate the new values.

The System Status tab also showed yellow alerts for the new device pools that were discovered during the auto detect process. The alerts indicate that the new device pools do not have default devices that are specified. We can eliminate these alerts by adding default devices to the device pool objects.

## Updating device pools

To update the device pool objects, click **Administration** in the menu bar, and then select **Manage Device Pools**, as shown in Figure 4-86.

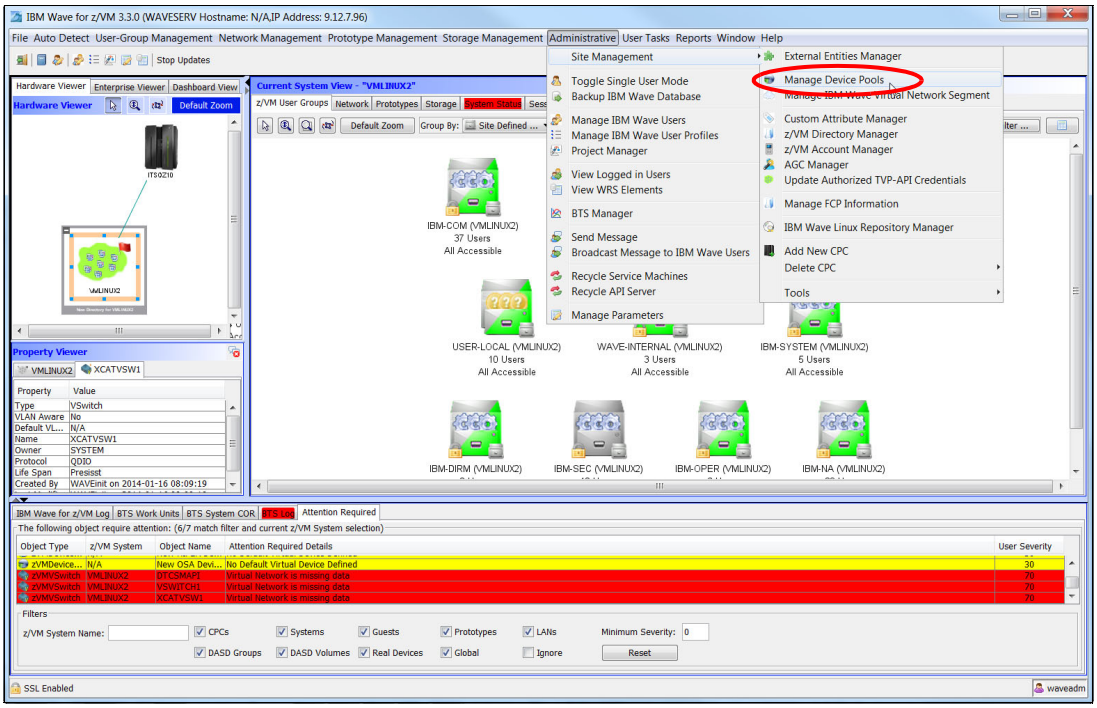


*Figure 4-86 Manage Device Pools*

When you click **Manage Device Pools**, the Device Pool Manager window opens. Each line in the window represents a new device pool that was discovered during the auto detect process. Click one of the lines to select that line, then right-click to open a menu and select **Update Details**, as shown in Figure 4-87.

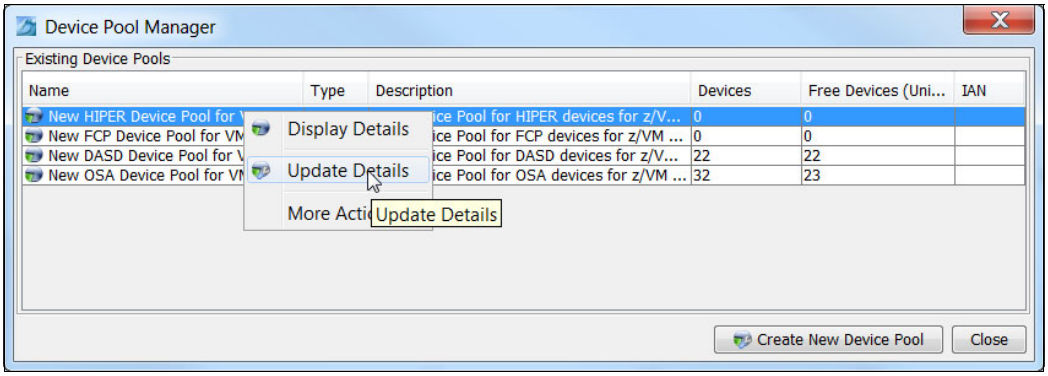


*Figure 4-87 Update Details*

When you click **Update Details**, the Update Device Pool window opens. Enter a value for Default Virtual Device and click **Update**, as shown in Figure 4-88 on page 79.

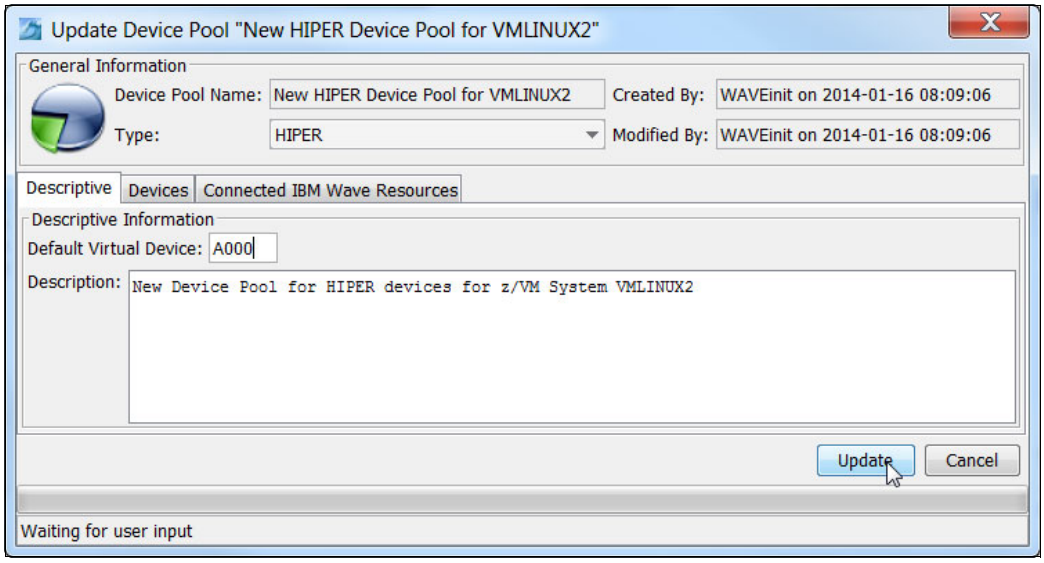


*Figure 4-88 Add Default Virtual Device*

When you click **Update**, the default virtual device is added to the particular device pool and you are returned to the Device Manager window. Continue entering default devices for all of the new device pools.

After all of these items are updated, the System Status tab and the Attention Required pane of the IBM Wave Viewer should no longer show attention items, as shown in Figure 4-89.

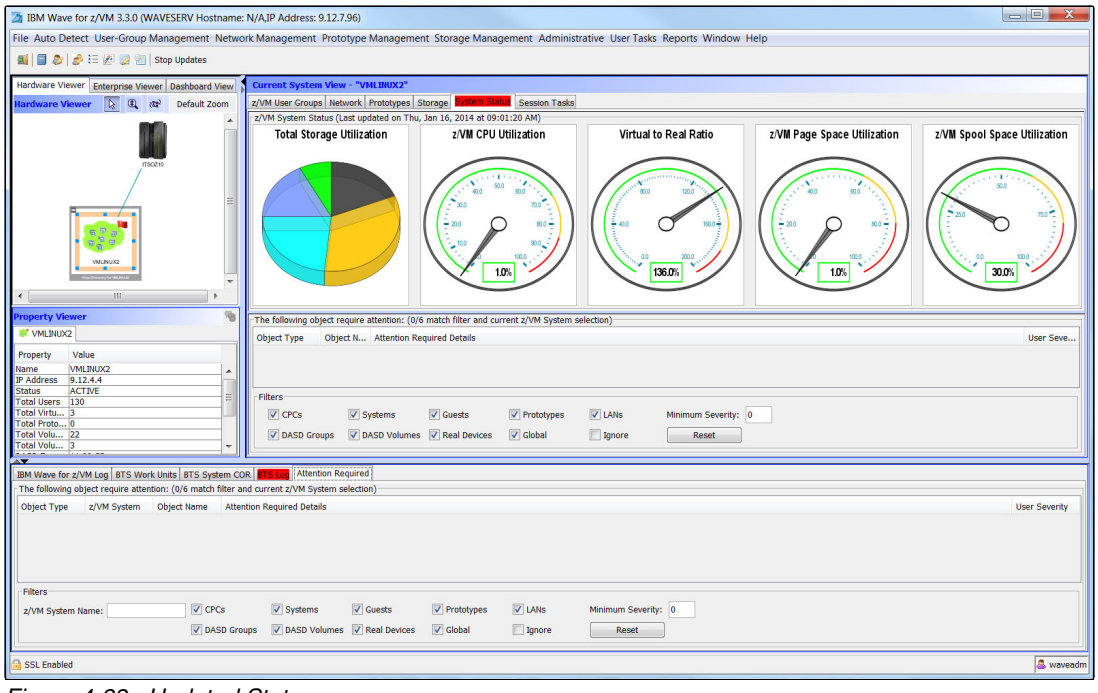


*Figure 4-89 Updated Status*

## Marking WAVESRV as an IBM Wave Service Machine

One last task to complete after the installation is to mark the WAVESRV VM (the Linux VM that is running the IBM Wave background services) as an IBM Wave Service Machine. This action prevents certain guest user actions against WAVESRV.

IBM Wave is aware which worker VMs are part of the product during the auto detect process. However, IBM Wave is not yet aware in which VM the IBM Wave background processes are running. To mark the WAVESRV VM in this manner, go to the z/VM User Groups view in the Current System view and expand the group USER-LOCAL. Right-click the icon for WAVESRV and select **More Actions** → **Mark As IBM Wave Service Machine** from the menu, as shown in Figure 4-90.

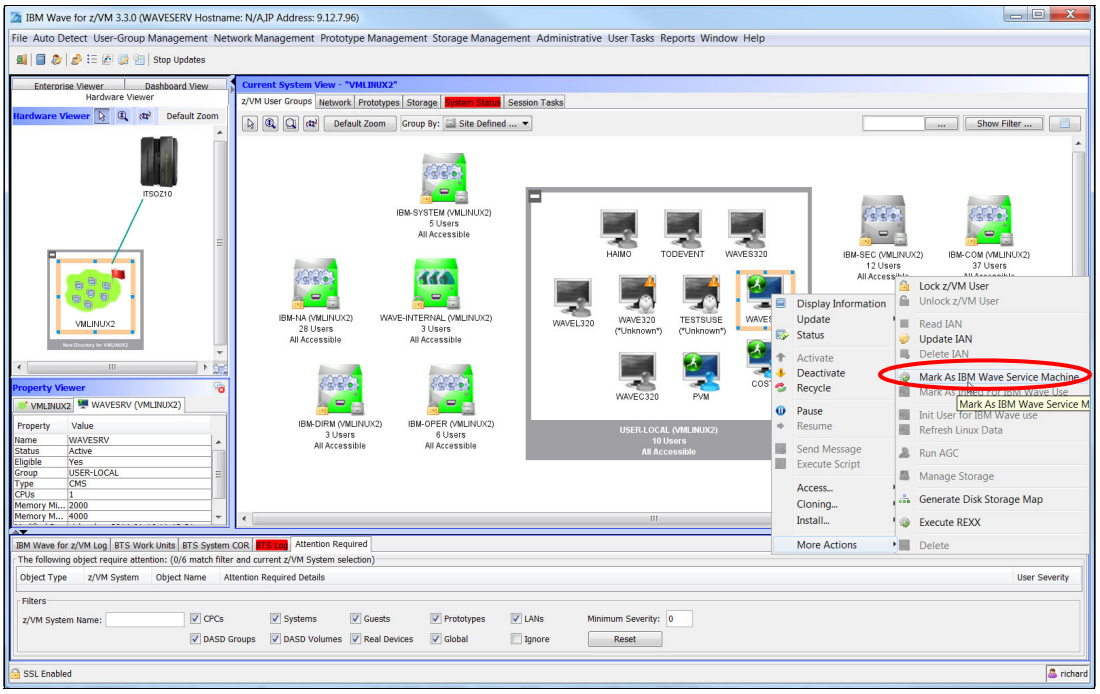


*Figure 4-90 Mark as IBM Wave Service Machine*

When you click **Mark As IBM Wave Service Machine**, a confirmation window opens. Click **Go**, as shown in Figure 4-91.

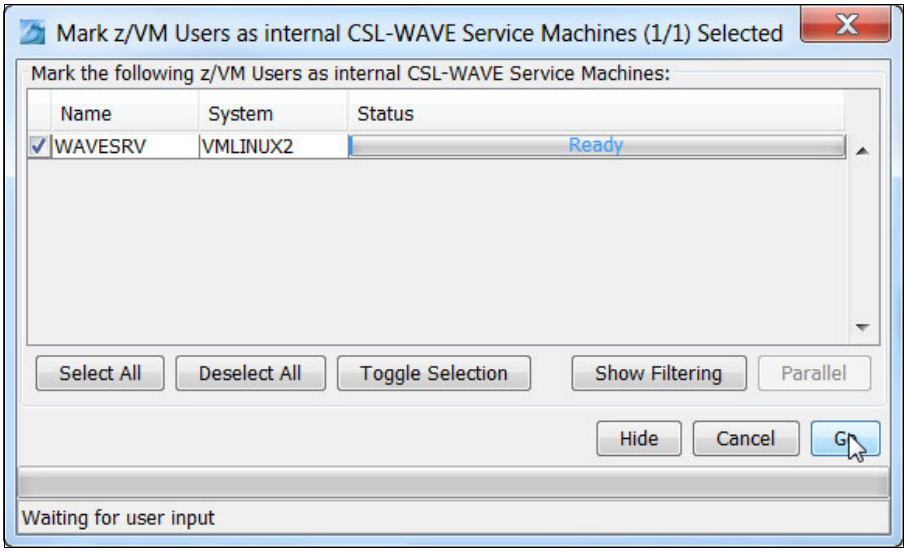


*Figure 4-91 Mark service machine action*

When you click **Go**, the system locks the VM from further changes and moves it to the group of VMs that are designated as WAVE-INTERNAL. The window indicates the outcome, as shown in Figure 4-92.

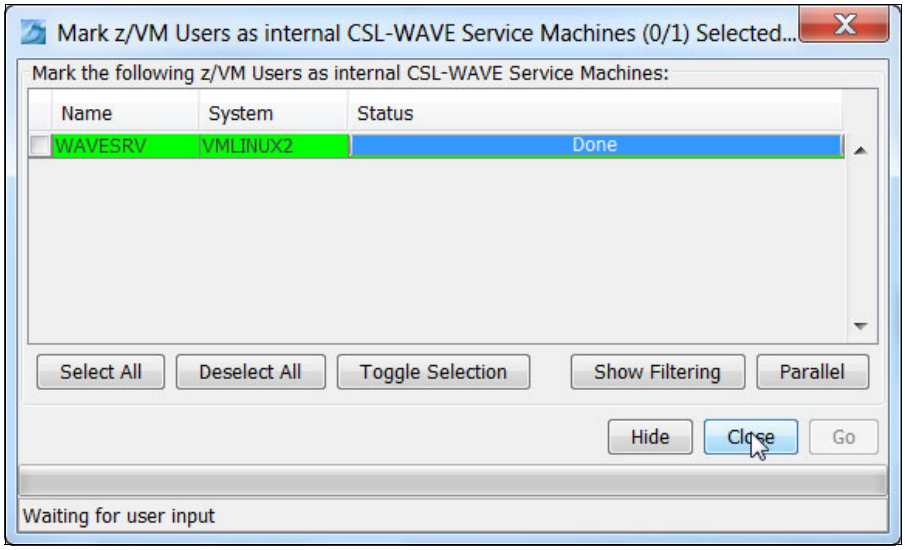


*Figure 4-92 Mark service machine complete*

Click **Close** to return to the main window. Expanding the WAVE-INTERNAL group icon shows that the WAVESRV VM is now part of that group, as shown in Figure 4-93.

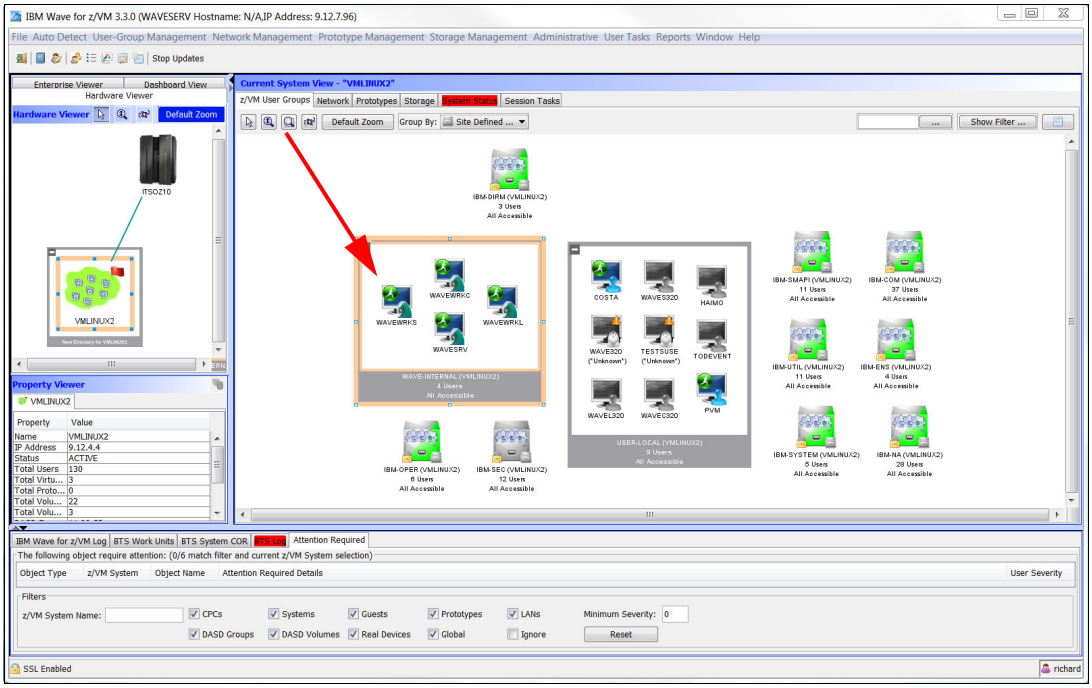


*Figure 4-93 Updated WAVE-INTERNAL Group.*

The alerts that result from the Auto Detect process are now addressed. IBM Wave is installed and a z/VM system is now ready to be managed.

## 4.3 IBM Wave Command-Line Interface installation

There is an IBM Wave Command-Line Interface (CLI) that can be used either on another Windows or Linux server or workstation. You must copy the appropriate file to the destination server or workstation. We did set up an anonymous FTP server in our WAVESRV server to access the files that are in `/usr/wave/WAVECLI`.

There is a `wavecli.zip` file for Windows and a `wavecli.tar` file for Linux. Extract the file into the destination folder.

### 4.3.1 IBM Wave CLI in Windows

In our environment, we extracted the Wave 1.2.05 `wavecli.zip` file into `/users/roland/WAVECLI` on a Windows workstation.

After extracting the Wave 1.2.05 `wavecli.zip` file, you have the following files

- `WAVECommon120.jar`
- `log4j-1.2.16.jar`
- `jdom.jar`
- `jcommon-1.0.12.jar`
- `jcommon-serializer-0.2.0.jar`
- `commons-collections-3.2.jar`
- `commons-dbcp-1.2.2.jar`
- `commons-pool-1.4.jar`
- `jsch-0.1.42.jar`
- `WAVECLI120.jar`
- `wavecli.bat`

For more information about this feature and its use, see 5.11, “Using the IBM Wave Command-Line Interface” on page 161.

# Part 2

# Using IBM Wave for z/VM

This part describes the use of IBM Wave for z/VM (IBM Wave) to manage your environment and describe the following tasks:

- Creating a z/VM guest
- Installing Linux in a virtual machine
- Creating a golden master
- Cloning virtual machines from a golden master

All of these tasks are done by using IBM Wave.

This part includes the following chapters and one appendix:

5

# Managing and operating the user interface environment

This chapter provides an introduction to the graphical user interface (GUI) for IBM Wave for z/VM (IBM Wave) and describes various tasks that use the GUI. It also describes the process of creating groups, projects, and custom attributes and how these objects help manage a z/VM system.

This chapter describes the process of creating IBM Wave users and assigning scopes, permissions, group associations, and project associations. Finally, this chapter describes the process of creating reports and backup functions that are available through the GUI.

The chapter includes the following topics:

- IBM Wave user interface
- Toolbar
- Viewers
- IBM Wave Enviroment Viewers
- The Status Bar
- Group management
- Project management
- Creating reports
- User management
- Backup option for IBM Wave
- Using the IBM Wave Command-Line Interface
- Communication-Less Connection technology

# 5.1 IBM Wave user interface

When you log in to IBM Wave, the first window that you see is the main IBM Wave window, as shown in Figure 5-1.

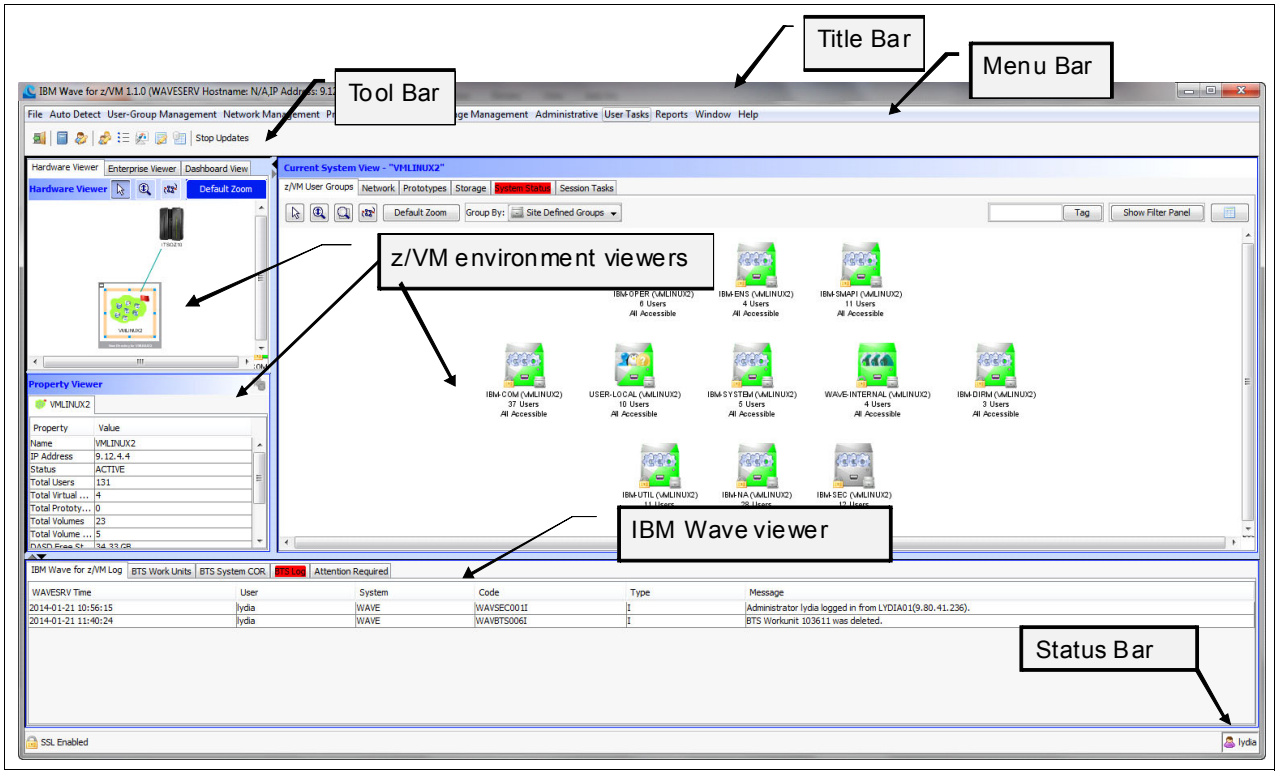


*Figure 5-1 IBM Wave Main window*

## 5.1.1 Title bar

Figure 5-2 shows the Title bar.

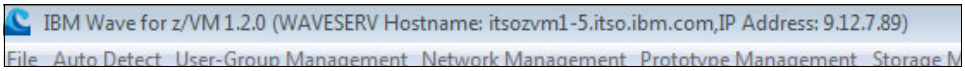


*Figure 5-2 IBM Wave Title bar*

The Title bar contains basic information about the IBM Wave server and buttons to manipulate the primary IBM Wave window. As shown in Figure 5-2 (from left to right), the Title bar contains the following components:

- IBM Wave icon
- Version of IBM Wave
- Host name of the server that is running IBM Wave
- IP address of the server that is running IBM Wave

## 5.1.2 Menu bar

The menu bar contains various submenus that provide various functions. Although many submenus can be intimidating, many of these menus contain only one function. Also, most of these functions can be accessed through other means, such as the toolbar or menus in various viewers.

### File menu

Figure 5-3 shows the File menu.

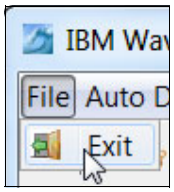


*Figure 5-3 File menu*

The File menu contains a single item that is selected to exit IBM Wave.

### Auto Detect menu

The Auto Detect menu contains all of the functions for auto detecting new z/VM systems and refreshing Viewers for IBM Wave systems. Because all of these functions operate at the level of a specific z/VM system or member, all selections are inaccessible unless a particular z/VM system is selected in the Hardware Viewer. Figure 5-4 shows the Auto Detect menu.

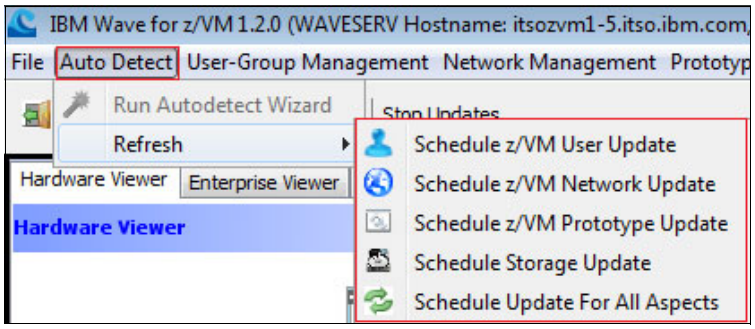


*Figure 5-4 Auto Detect menu*

The Run Auto Detect wizard item is used to perform the initial configuration that is needed to allow IBM Wave to manage a z/VM system. The process discovers all resources that are associated with a z/VM system and builds the IBM Wave metadata repository. This option is accessible only if a z/VM system is selected in the Hardware Viewer that has not yet run the Auto Detect wizard.

The Refresh Submenu provides the capability to refresh specific metadata for a selected z/VM system. Because IBM Wave updates its metadata repository on a periodic basis, these selections provide a manual method of ensuring that the metadata repository is fresh.

## User-Group, Network Management, Prototype, and Storage Management menus

Figure 5-5 shows the User-Group, Network Management, Prototype, and Storage Management menu selections.

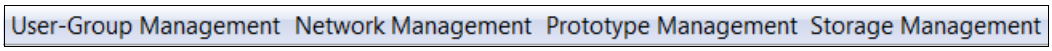


*Figure 5-5 User-Group, Network, Prototype, and Storage Management menu links*

Using these menus, you can select which Viewer to display in the Current System View. All of the options in these menus are accessible only when a specific z/VM system is selected in the z/VM Hardware Viewer.

## Administrative menu

The Administrative menu contains various items that are related to the IBM Wave Administrator role. This menu is available only to IBM Wave users who have the IBM Wave administrator role. Figure 5-6 shows the Administrative menu.

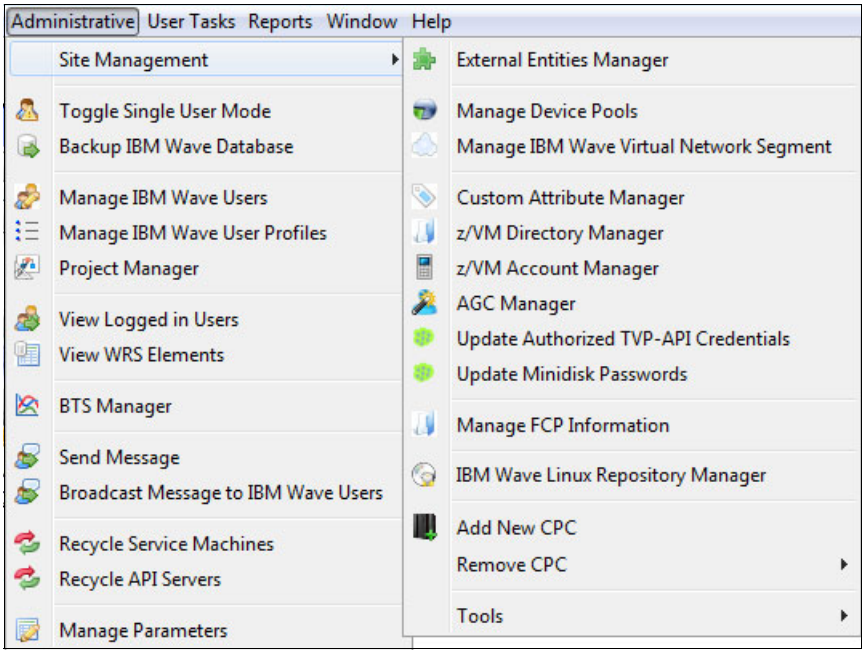


*Figure 5-6 Administrative menu*

The Administrative menu features the following selections:

- The Site Management menu contains various items that are related to managing physical resources, such as mainframe central processor complexes (CPCs), z/VM systems, routers, z/OS system, Linux repositories, device pools, FCP devices, and network segments. It also contains the following selections that are used for configuring custom attributes for managed guests and managing z/VM account codes:
  - The External Entities Manager is used to define metadata for major entities that exist external to IBM Wave, including z/OS or z/VM systems, storage controllers, and routers.
  - Manage Device Pools is used to create and manage pools of DASD, OSA, HiperSockets, and FCP devices.

  - Manage IBM Wave Virtual Network Segment is used to display the virtual network topology of the z/VM environment to IBM Wave.
  - The Custom Attribute Manager is used to define possible custom attributes that IBM Wave users can assign to guest virtual machines (VMs).
  - The z/VM Directory Manager is used to view and update the z/VM directories that are managed by IBM Wave.
  - The z/VM Account Manager is used to create, update, and delete z/VM accounts that are associated with VMs.
  - The AGC Manager is used to configure the Automatic Guest Configuration Engine of IBM Wave. This engine is used to create rules that allow auto-discovery to assign automatically IBM Wave metadata or allow changes to IBM Wave metadata to make changes to the z/VM Directory.
  - Manage FCP Information is used to assign WWPNs to FCP storage controllers to allow IBM Wave to manage FCP disks.
  - The IBM Wave Linux Repository Manager is used to create, modify, and delete Linux Repositories for use by the IBM Wave Bare Metal Installation function.
  - Add New CPC is used to add a mainframe CPC to IBM Wave.
  - Delete CPC is used to fully delete an existing CPC from IBM Wave. Each of the CPCs that are managed by IBM Wave are listed in this menu.
  - The Tools menu contains the Import Guest Information option, which allows an IBM Wave user to import guest metadata from a `.csv` file that is formatted according to the criteria in the *IBM Wave User Guide*.
- Toggle Single User Mode puts IBM Wave in single user mode, which is a mode where the IBM Wave administrator is the only IBM Wave user that is logged in to IBM Wave. This mode is required for certain IBM Wave actions, such as shutting down a z/VM system.
- Backup IBM Wave Database is used to export the IBM Wave metadata repository to unencrypted flat files.
- Manage IBM Wave Users is used to create or delete IBM Wave users and assign users roles, privileges, and scope.
- Manage IBM Wave User Profiles is used to create User Profiles to associate with LDAP groups when LDAP is used for access control.
- Project Manager is used to create or delete IBM Wave projects.
- View Logged In Users is used to see the users that are logged in to IBM Wave and optionally force the user to log off if the user is hung up because the user did not gracefully log out of IBM Wave.
- View WRS Elements views all of the active IBM Wave Resource Serialization (WRS) Elements, which are the locks on various IBM Wave elements. WRS elements are used to ensure that certain windows are locked when an IBM Wave user has it open to prevent two users from submitting conflicting requests to the IBM Wave server at the same time. However, a WRS element can get hung up, which prevents any user from opening a particular window. This option allows the IBM Wave administrator to end forcibly these WRS elements if they get hung up.
- BTS Manager allows an administrator to display and manipulate various information about the Background Task Scheduler (BTS) of IBM Wave, including connected users, the status of each IBM Wave worker thread, the schedule that determines how often IBM Wave periodic updates occur, historical statistics for all work units that are submitted to the BTS, and a log of all BTS activity.

- Send Message and Broadcast Messages allows an IBM Wave administrator to message directly a specific IBM Wave user or broadcast a message to all IBM Wave users. This ability is useful when you are notifying users about systems that are coming down or when you are requesting that users log off IBM Wave to allow the administrator to enter single-user mode for system-wide actions.
- Recycle Service Machines and Recycle API Server are used to recycle the WAVEWRK* service machines and the z/VM Service Management API.
- Manage Parameters is used to configure various system-wide IBM Wave settings.

## User Tasks menu

As shown in Figure 5-7, the User Tasks menu includes tasks that regular IBM Wave users can perform.

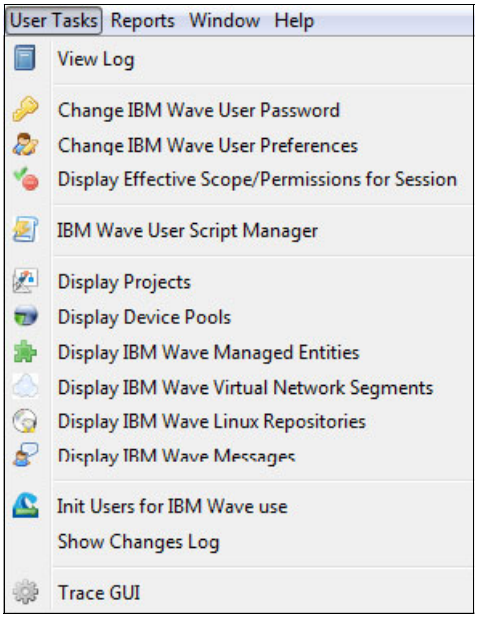


*Figure 5-7 User Tasks menu*

The menu has the following selections:

- View Log allows an IBM Wave user to display all of the logs that are produced by the IBM Wave server.
- Change IBM Wave User Password allows a user to change their password.
- Change IBM Wave User Preferences configures user settings for the current IBM Wave user, including display settings, configuration settings for an external SSH terminal, and SSH and 3270 connection settings.
- Display Effective Scope/Permissions for Session displays the scope and permission within IBM Wave, which determines the tasks that the user can perform.
- The IBM Wave User Scripts Manager allows a user to create or delete Linux shell scripts that the IBM Wave server can run across the managed z/VM environment.
- Display IBM Wave Managed Entities, Display IBM Wave Virtual Network Segments, Display IBM Wave Linux Repositories, and Display IBM Wave Messages gives a non-administrator IBM Wave user read-only access to a subset of the elements that the IBM Wave administrator configured by using the actions in the Administrative menu.

- Init Users for IBM Wave User configures a z/VM Linux guest to be eligible for management by the IBM Wave interface. This task must be completed to manage the guest by using IBM Wave interfaces.
- Show Changes Log shows the changes that the IBM Wave user performed during the session.
- Trace GUI toggles whether the client-side IBM Wave GUI is outputting status and error messages to the local Java console.
- Gather Debug Information is used to generate a large report that summarizes the current state of the installed IBM Wave environment to allow IBM to troubleshoot IBM Wave errors.

A user can turn on the Trace GUI option in the User Tasks menu, as shown in Figure 5-8.

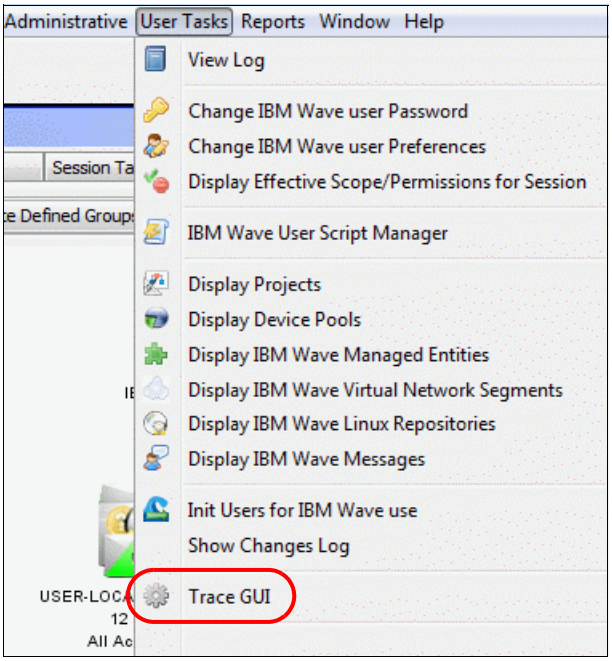


*Figure 5-8 Trace GUI*

If the Trace GUI option is selected, the user can set the GUI Debug Level by clicking **Administrative** → **IBM Wave Parameters**. A window opens that is similar to the window that is shown in Figure 5-9. Select the **GUI** tab, choose the trace level that you want, and set the GUI Entry Limit, which is the maximum number of log entries that appear in the Log Viewer in the General Status Viewer.

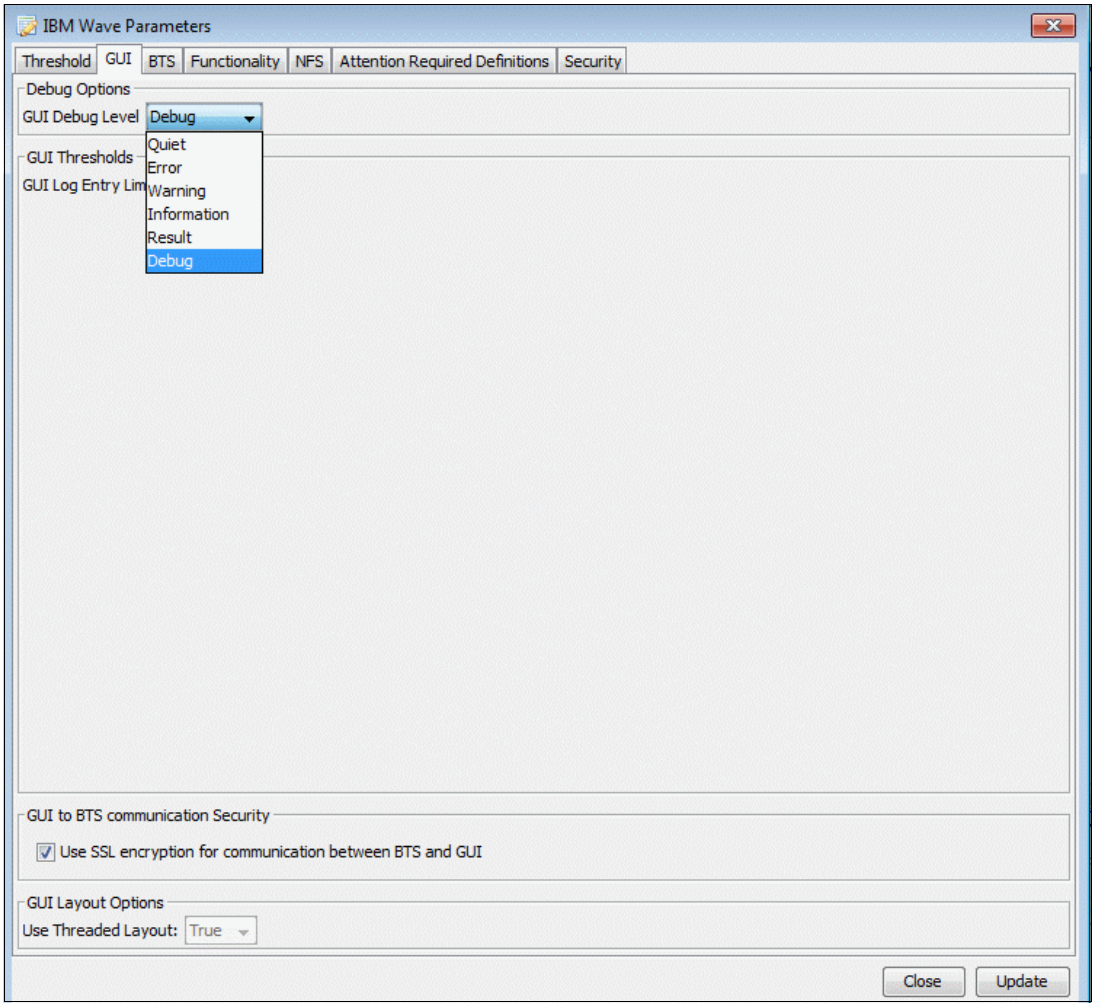


*Figure 5-9 Set the Debug Level for Trace GUI*

The GUI Debug Log is on the client workstation in a temporary directory. You can find it in `C:\Users\<<username>>\AppData\Local\Temp`. The format for the name of the log is `WAVEGuiLog-<<dd-mm-yy>>.txt`.

## Reports menu

The Reports menu (Figure 5-10) contains the Report Manager selection, which allows an IBM Wave user to generate printable reports that show the status of selected aspects of the z/VM environment.

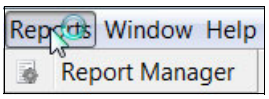


*Figure 5-10 Reports menu*

For more information about reports, see 5.8 “Creating reports” on page 129.

## Window menu

The Window menu contains options for controlling how viewers are displayed, as shown in Figure 5-11.

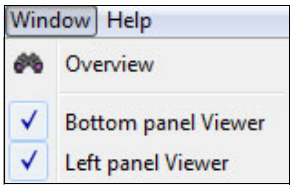


*Figure 5-11 Window menu*

The Window menu includes the following options:

- Overview provides a tool to more easily control the zoom level and positioning of the Current System View.
- Bottom Panel Viewer toggles whether the IBM Wave Viewers are displayed or minimized.
- Left Panel Viewer toggles whether the tabbed element that contains the Hardware Viewer, Enterprise Viewer, and Dashboard Viewer is displayed or minimized.

## Help menu

The Help menu contains information about the versions of IBM Wave and your Java virtual machine (JVM) and a few links to IBM Wave documentation. Figure 5-12 shows the Help menu and the About Wave window.


*Figure 5-12 Help menu*

The Help menu includes the following choices:

- The About menu shows the version information about IBM Wave.
- Show Java Information shows the version of the JVM that is running on your workstation.
- View User Guide opens the *IBM Wave User Guide* in PDF form.
- View Administration Guide opens the *IBM Wave Administration Guide* in PDF form.

## 5.2 Toolbar

As shown in Figure 5-13, the IBM Wave toolbar is composed of small icons that start various functions. Some of the functions are visible only to IBM Wave Administrators.

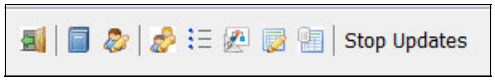


*Figure 5-13 Toolbar*

From left to right in Figure 5-13, the toolbar features the following functions:

- Clicking **Exit** (icon of a door) logs off the IBM Wave user and exits the IBM Wave program. This function is also accessible by clicking **File** → **Exit** in the menu bar or by clicking the **Close** button on the Title Bar.
- Clicking **View Log** (icon of a book) displays all of the logs that are produced by the Wave server. This function is also accessible by clicking **User Tasks** → **View Log** in the menu bar or by double-clicking the **IBM Wave Log** tab in the IBM Wave Environment Viewer.
- Clicking **Change IBM Wave User Preferences** (icon of a person with pencil) configures the user's settings, including display settings, configuration settings for an external SSH terminal, and SSH and 3270 connection settings. This function is also accessible by clicking **User Tasks** → **Change IBM Wave User Preferences** in the menu bar.
- Clicking **Manage IBM Wave Users** (icon of two people with a pencil) creates or deletes IBM Wave users and assigns users roles, privileges, and scope. This option is available only to Administrative users. This function is also accessible by clicking **Administrative** → **Manage IBM Wave Users** in the menu bar.
- Clicking **Manage IBM Wave User Profiles** (icon of bullet points and lines) creates User Profiles to associate with LDAP groups when LDAP is used for access control. This option is available only to Administrative users. This function is also accessible by clicking **Administrative** → **Manage IBM Wave User Profiles** in the menu bar.
- Clicking **Project Manager** (icon of a flow chart and a triangle) creates or deletes IBM Wave projects. This action is available only to Administrative users. This function is also accessible by clicking **Administrative** → **Manage Projects** in the menu bar or the menu when the z/VM User Groups tab of the Current System Viewer is grouped by project.
- Clicking **Manage Parameters** (icon of a pencil and a spreadsheet) configures various system-wide IBM Wave settings. This action is available only to Administrative users. This function is also accessible by clicking **Administrative** → **Manage Parameters** in the menu bar.
- Clicking **View WRS Elements** (icon of a cylinder and a spreadsheet) views and clears the locks on various IBM Wave elements. This option is available only to Administrative users. This function is also accessible by clicking **Administrative** → **View WRS Elements** in the menu bar.
- Clicking **Stop Updates** halts and queues all of the periodic update processes of IBM Wave to ensure that the IBM Wave metadata repository reflects the current state of the z/VM environment. After it is stopped, this option is relabeled as Process Updates and runs the queue of IBM Wave updates that were halted. This option is available only to Administrative users, and is the only function on the toolbar that cannot be run through other means in the IBM Wave GUI.

## 5.3 Viewers

The System Viewers (Figure 5-14) are an important piece of the IBM Wave interface because they are where systems management tasks are run.

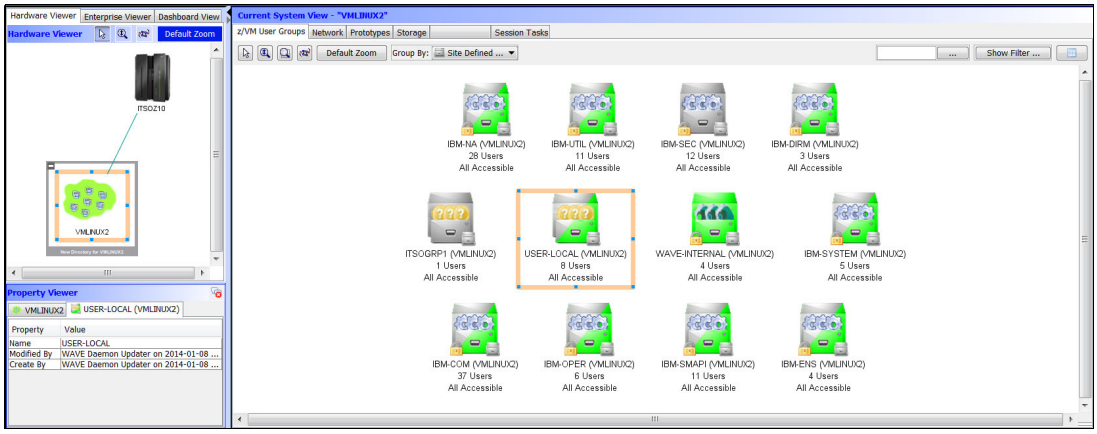


*Figure 5-14 Viewers window*

System Viewers include the following components:

- The Viewer Navigator is on the left side of the GUI. It is divided into the Hardware Viewer, Enterprise Viewer, and Dashboard Viewer tabs. These viewers are the high-level organization for this portion of the GUI and are used to select and filter the information that is displayed by IBM Wave in the View area, which is on the right side of the GUI.
- The View Area displays the view that is associated with the selection that is shown in the Viewer Navigator. The Current System View corresponds to the z/VM system that is selected in the Hardware Viewer. The Enterprise View corresponds to the filter criteria that is set in the Enterprise Viewer. The Dashboard View corresponds to the filter criteria that is set in the Dashboard Viewer.
- The Property Viewer displays tabs of information that are related to items that are selected in the Viewer Navigator and the View Area.

### 5.3.1 Hardware Viewer

The Hardware Viewer visualizes mainframe CPCs, instances of the z/VM hypervisor, and single-system image (SSI) clusters. Figure 5-15 shows the menus that are available from the Hardware Viewer.

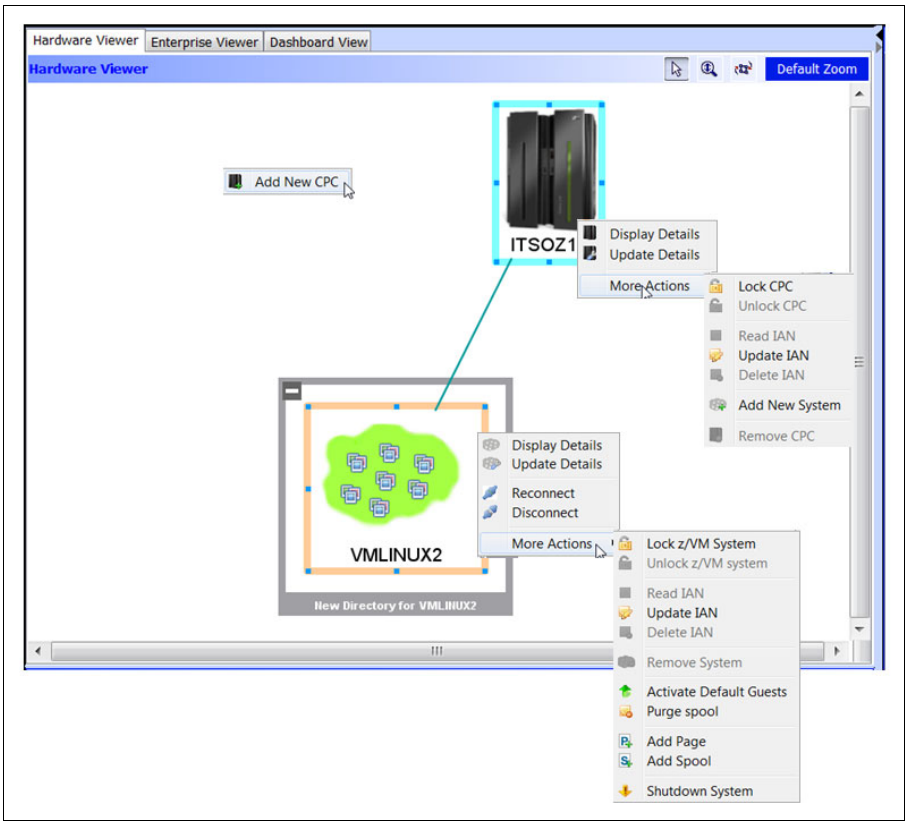


*Figure 5-15 Hardware Viewer*

Right-clicking the white space in this area opens a menu with which an IBM Wave administrator can add a mainframe server through the Add New CPC task.

Right-clicking a mainframe CPC opens a menu with which an IBM Wave administrator can delete the CPC, update CPC information (such as model and CPUID), and add a z/VM system.

Right-clicking a z/VM system opens a menu with which an IBM Wave administrator can delete the system from the IBM Wave metadata repository, update the z/VM system's metadata, purge the z/VM system's spool, add page or spool space to the system, and shut down (and optionally restart or reload) the entire z/VM system.

Clicking any element in this view opens a property tab for that element in the Property Viewer. Additionally, clicking z/VM systems opens a Current System View of that z/VM system in the View Area.

### 5.3.2 Current System View

The Current System View shows the system that is selected in the Hardware Viewer.

## z/VM User Groups tab

The z/VM User Groups tab of the Current System View displays the guest VMs that are running on the selected z/VM system. Through the Group By menu, the guest VMs can be organized by Site Defined Groups, Projects, Operating System Distribution, or z/VM Account. Figure 5-16 shows the z/VM Groups Tab and the menus that are available at the group level.

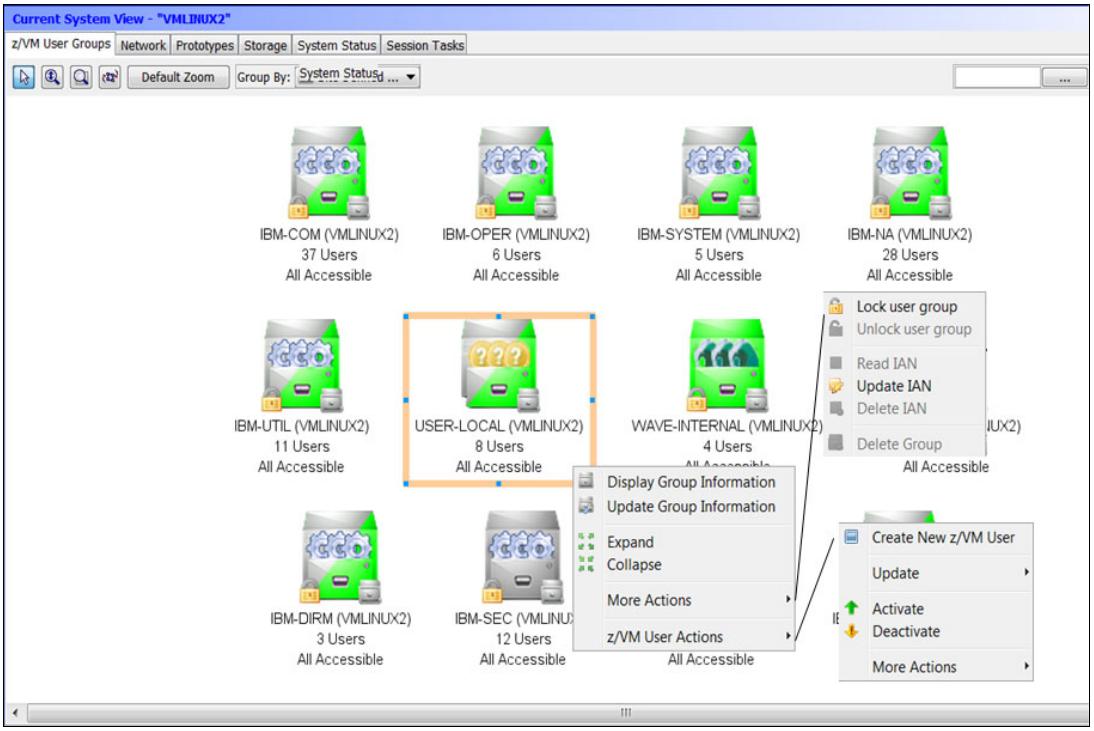


*Figure 5-16 z/VM Groups tab*

Right-clicking the white space in this view opens a menu with which the IBM Wave user can use to arrange items to clean up the display, collapse or expand all groups, or add an instance of the group that is selected in the Group By menu.

Right-clicking a group opens a menu with which the IBM Wave user can display information about the group or collapse and expand the group. Collapsing and expanding the group also can be done by double-clicking the group.

Right-clicking a VM within a group opens a menu that shows many tasks that can be done with the VM, as shown in Figure 5-17.

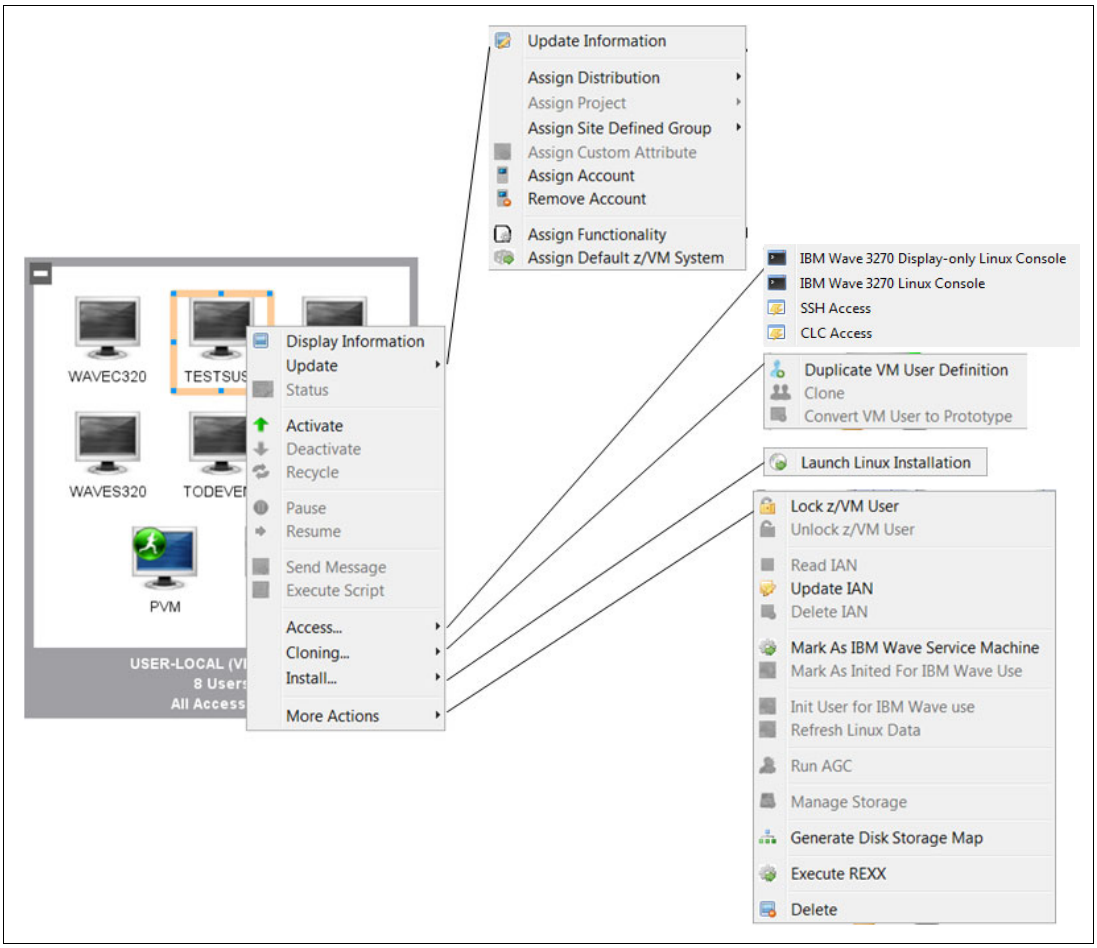


*Figure 5-17 Virtual machine menus*

Right-clicking a guest VM opens a menu with which a user can run the following IBM Wave functions on this guest:

- Display Information on all aspects of the guest from the hypervisor to guest operating system to IBM Wave metadata.
- Update Information that is related to the guest's z/VM directory entry (virtual memory, processors, and guest password).
- Assign and Remove Metadata, including OS Distribution, Projects, Site Defined Groups, Custom Attributes, accounts, and functions.
- Define the Default z/VM system that you want this guest to load (start) when they are starting a z/VM system.
- Access the guest through Display-only or full-featured 3270 terminals, access a Linux shell through SSH or through a feature that is called Communication-Less Connection (CLC) that screen-scrapes the guest's z/VM console (which is intended for emergency access when Internet Protocol network connectivity in unavailable). Details about CLC can be found in 5.12 “Communication-Less Connection technology” on page 166.

The following figures show some examples of guest access.

– IBM Wave 3270 Linux Console

  If you click in the IBM Wave 3270 Linux Console, you must enter the z/VM password of that guest, as shown in Figure 5-18.

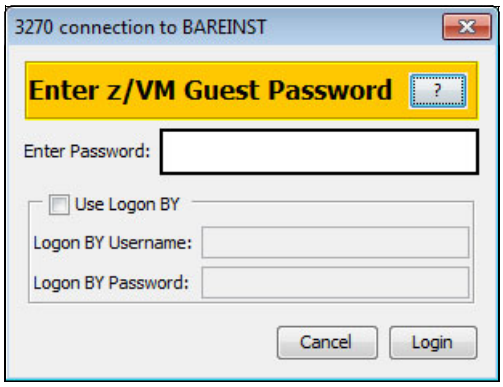


*Figure 5-18 3270 Linux Console Password Prompt*

Then, a 3270 Console, including a keypad, opens, as shown in Figure 5-19.

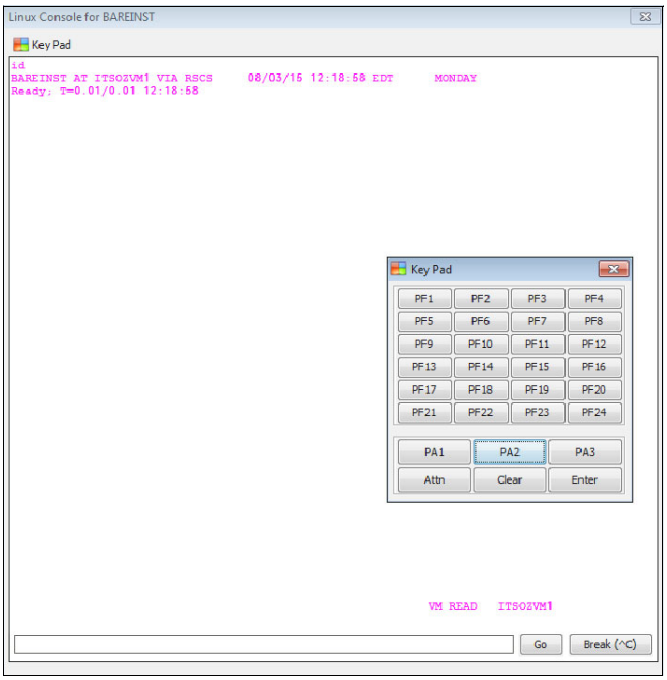


*Figure 5-19 3270 Console with keypad*

– SSH Access

  You must define which external SSH application you want to use, as described in 7.1.3 “Configuring an external SSH program to start sessions with X services” on page 221.

– CLC Access

If you are using the CLC access to a Linux z/VM guest, then you must provide the Linux user ID and password and the z/VM password, as shown in Figure 5-20.

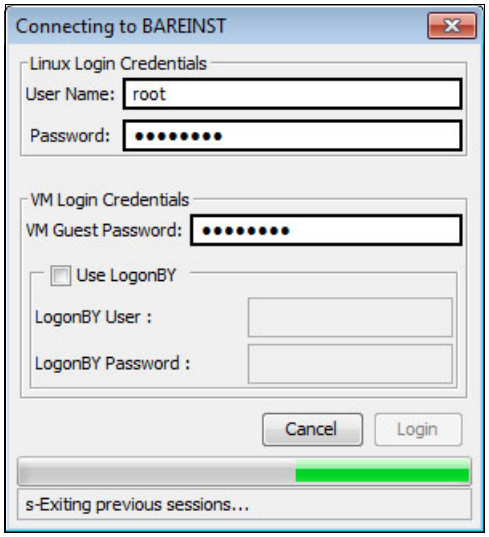


*Figure 5-20 CLC Access*

Then, the CLC window opens, as shown in Figure 5-21.

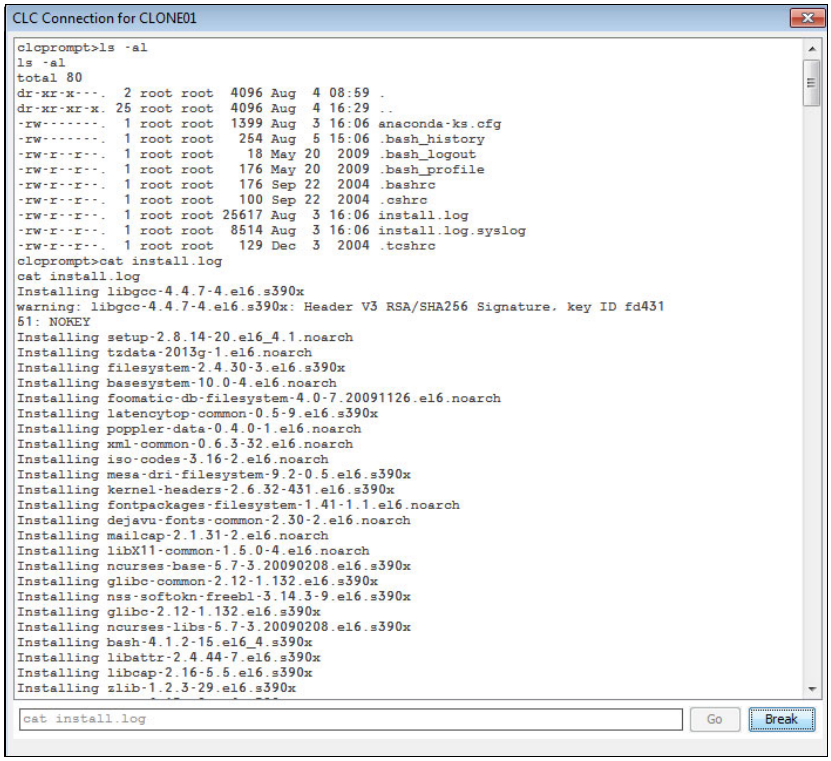


*Figure 5-21 CLC window*

You may use CLCEDIT to edit files, as shown in Figure 5-22 on page 101. The vi editor is not available in CLC.

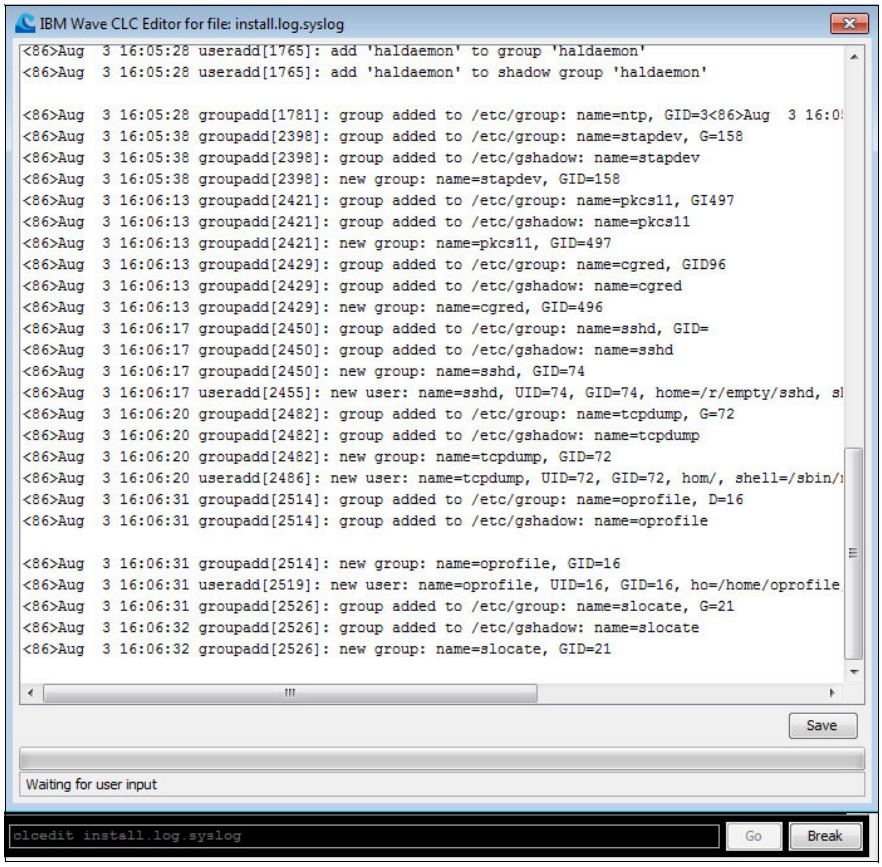


*Figure 5-22 CLCEDIT*

- Create a guest by cloning a guest's z/VM directory entry and disk volumes.
- Convert a guest into an Associated Prototype for use as a golden master.
- Install Linux on this guest through the Bare Metal Installation feature.
- Lock or unlock the guest to prevent or allow changes through IBM Wave.
- Add, update, or delete a sticky note to the guest to communicate information to all IBM Wave users.
- Update metadata.
- Create or expand LVMs in the guest's Linux file system.
- Run a REXX script on the guest.
- Delete the guest.

## Network tab

As shown in Figure 5-23, the Network tab of the Current System View displays the end-to-end virtual network topology of the selected z/VM system, including virtual LANs, network device pools, virtual switches, virtual network segments, and guests with TCP/IP stacks and IP addresses. This view can be customized by using the pane on the left side in which the IBM Wave user can select predefined views, create a view, and pick a particular layout.

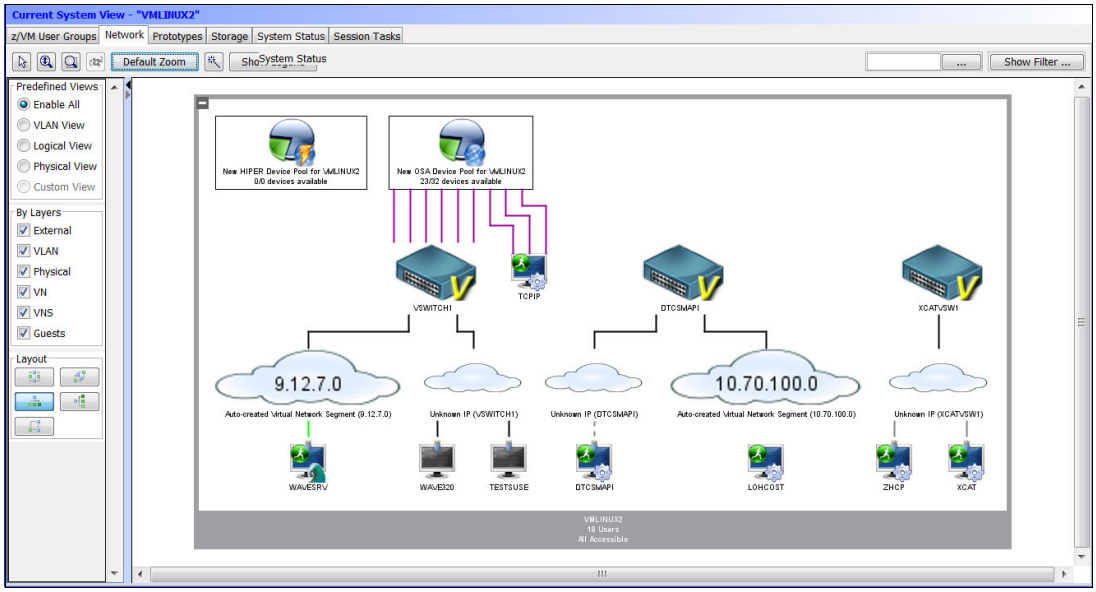


*Figure 5-23 Network tab*

Double-clicking any of the elements opens a detailed view of that element. By right-clicking any of the elements, an IBM Wave user can lock or unlock the selected element. The behavior of guests in this view is identical to their behavior in the z/VM User Groups window.

Right-clicking the white space in the window opens a detailed menu with which an IBM Wave user can define a guest LAN or virtual switch and configure the display of the network view, as shown in Figure 5-24 on page 103.

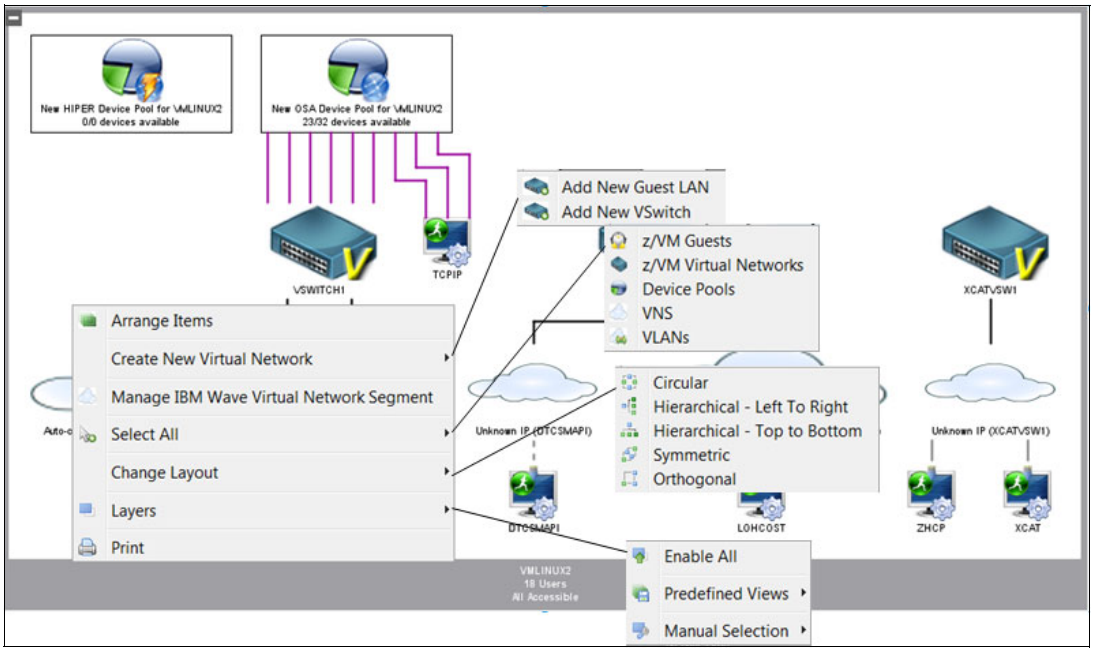


*Figure 5-24 Network tab menus*

## Prototypes tab

The Prototypes tab of the Current System View shows the prototypes that are defined in your Directory Manager and the associated prototypes that are defined in IBM Wave. When DirMaint is installed, two prototype definitions are supplied by default, one for Linux and one for Conversational Monitoring System (CMS). These definitions are displayed in the Prototypes tab, as shown in Figure 5-25.

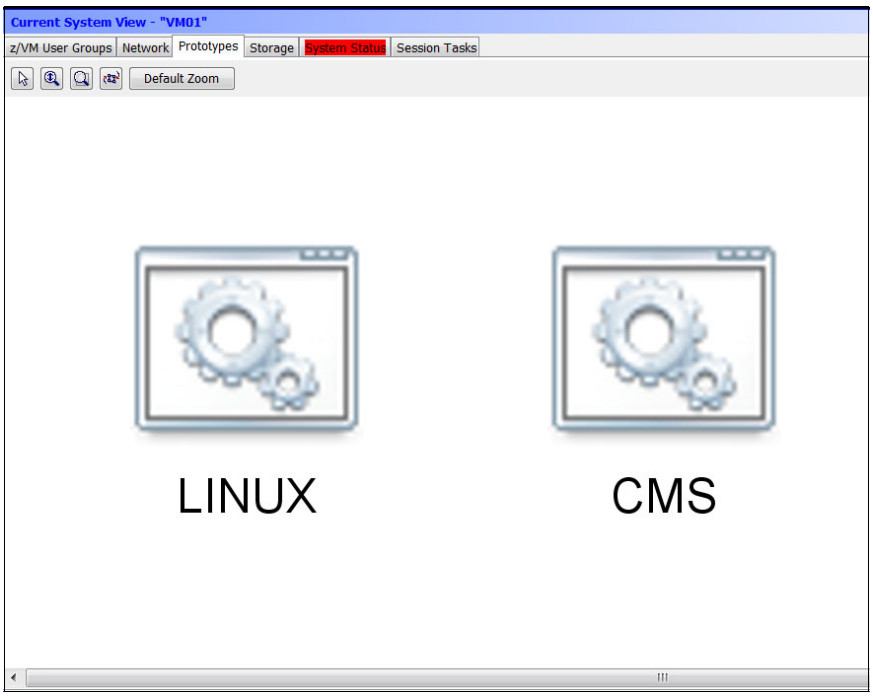


*Figure 5-25 Prototypes tab*

Right-clicking a prototype opens a menu with which an IBM Wave user can use to create guests by cloning this prototype. It also shows options for converting the prototype into a normal z/VM guest so that the master image can be updated and options for deleting the prototype, as shown in Figure 5-26.

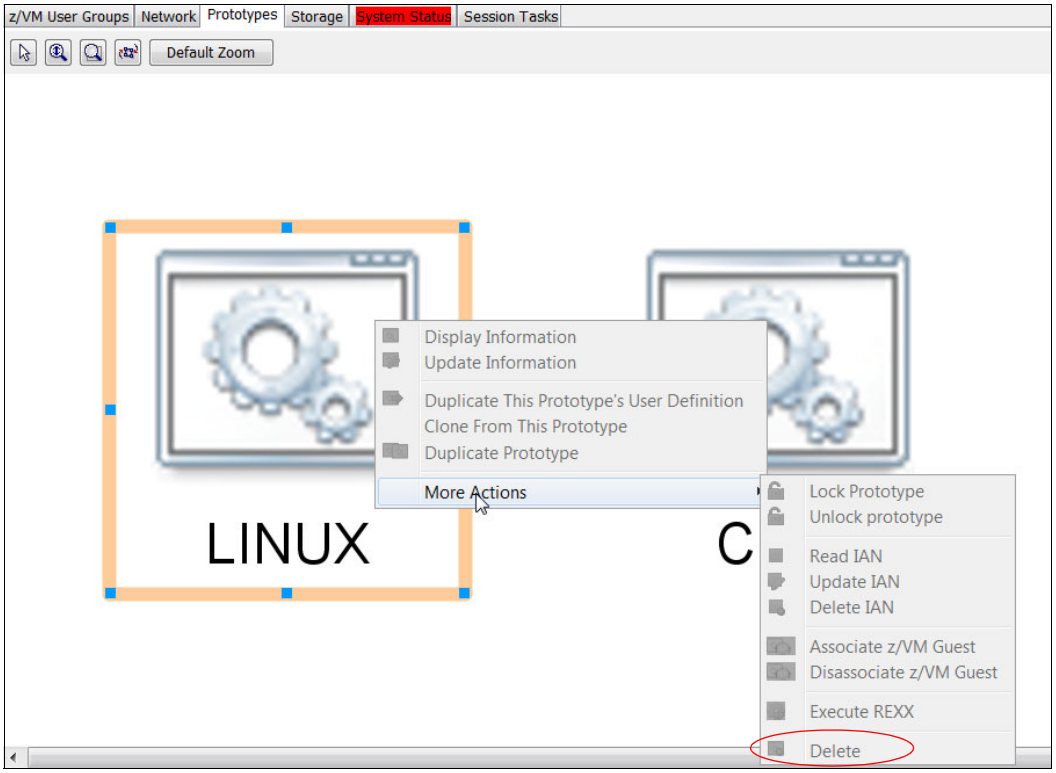


*Figure 5-26 Prototype menus*

## Storage tab

The Storage tab of the Current System View contains three different views of the storage system at various levels of granularity. These views can be selected by selecting the tabs on the left side of the window.

The first storage view is Distribution, which presents a pie chart that shows the portion of the storage distribution in each possible status. Each pie slice can be double-clicked to drill down into a view of the DASD volumes with that status. Figure 5-27 on page 105 shows the distribution view in the Storage tab.

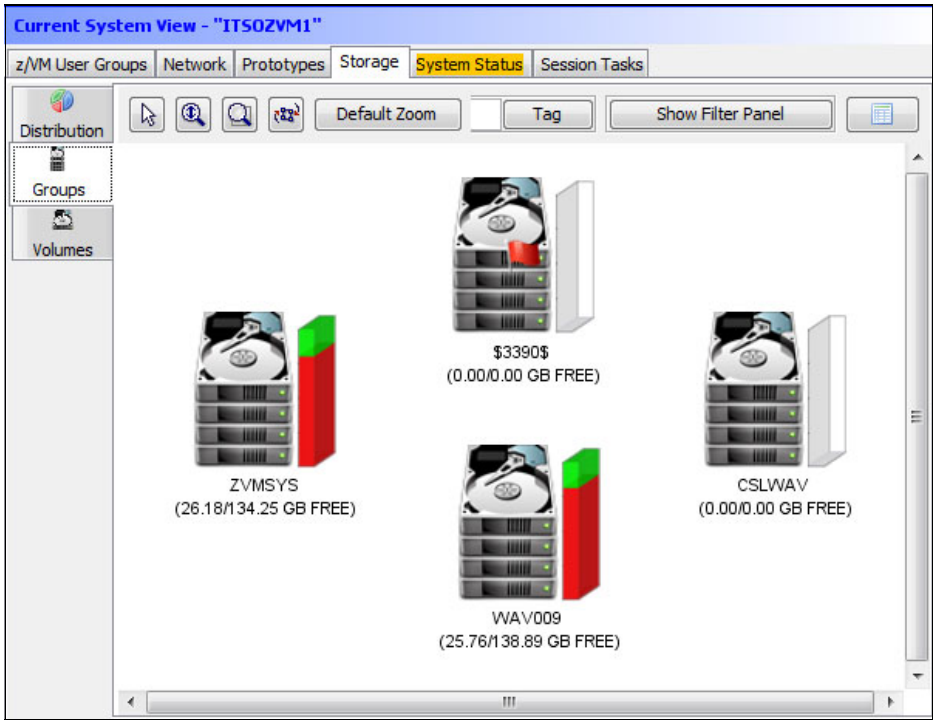


*Figure 5-27 Storage Distribution*

The second storage view is Groups, which shows each storage group with graphical and textual depictions of the amount of storage space that is free for allocation as minidisks to guest VMs.

Right-clicking the white space opens a menu with which DASD groups can be created. Right-clicking an existing DASD group opens the options to Update Information, such as renaming the group, as shown in Figure 5-28.

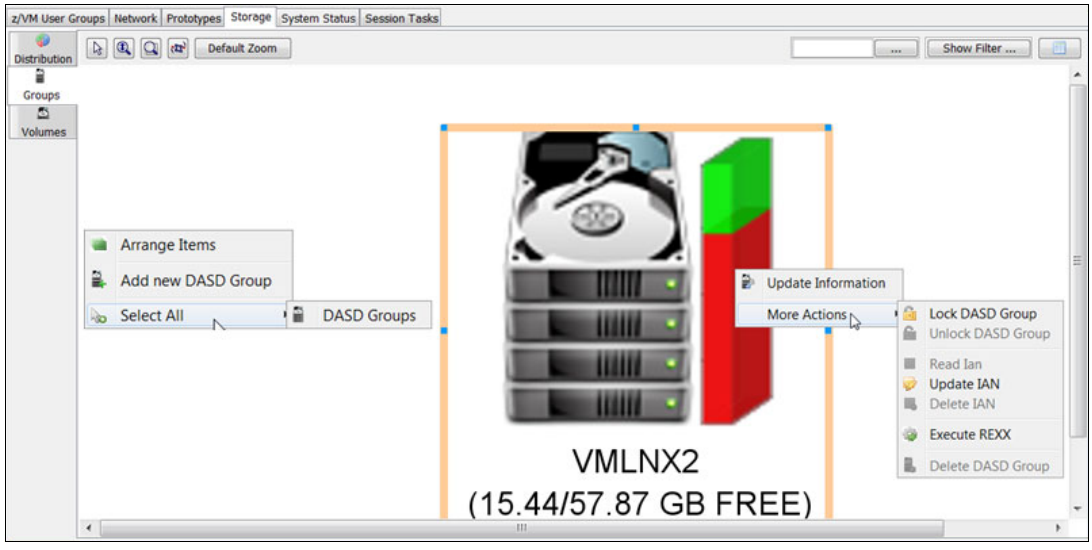


*Figure 5-28 Storage tab groups*

The third storage view is Volumes, which shows actual DASD volumes. Because a typical z/VM environment features too many DASD volumes to be effectively visualized on a single view, this view uses a filter that is at the bottom of the view, as shown in Figure 5-29.

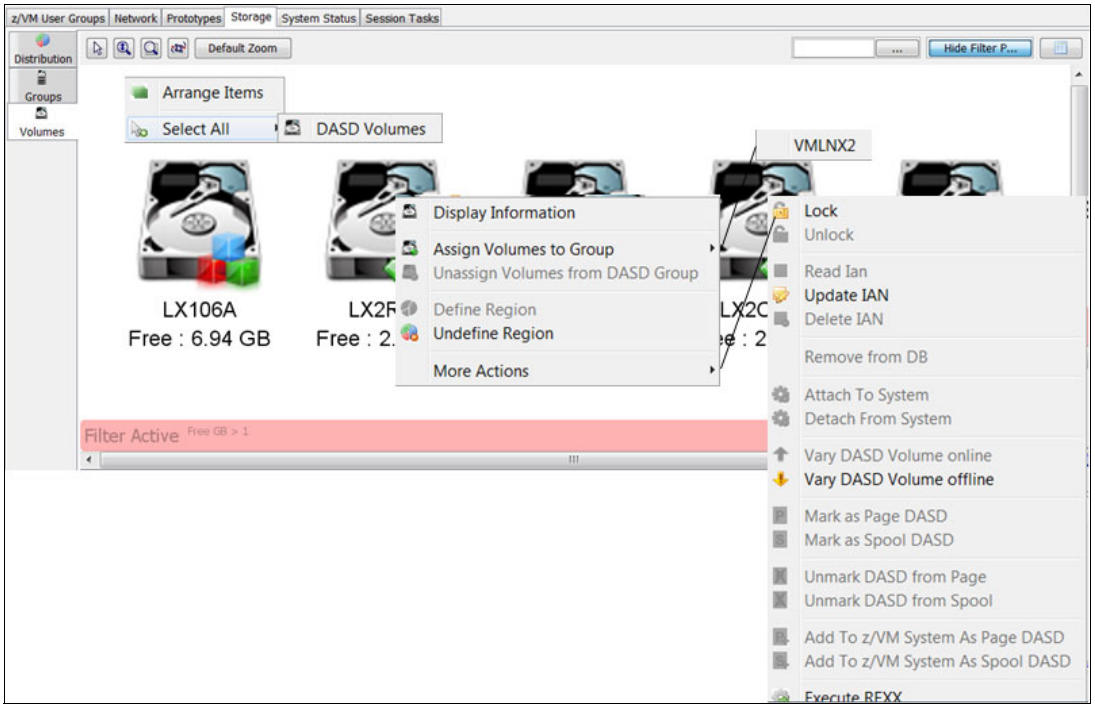


*Figure 5-29 Storage tab volumes*

By right-clicking a specific DASD volume, a menu opens with which an IBM Wave user can run the following tasks:

- Assign or unassign the volume to or from a DASD group.
- Define or undefine a region on the volume.
- Vary the DASD volume on or offline.
- Mark or unmark the DASD as page or SPOOL space for z/VM.
- Add the DASD volume to z/VM as a page or SPOOL volume.

## System Status tab

As shown in Figure 5-30, the System Status tab of the Current System View displays performance and storage information that is related to the guest and any administrative actions that IBM Wave believes that an IBM Wave administrator must complete. All of the panes can be double-clicked to open a window that features more detailed information. The objects that require attention can be right-clicked to open a menu to clear or ignore the item.

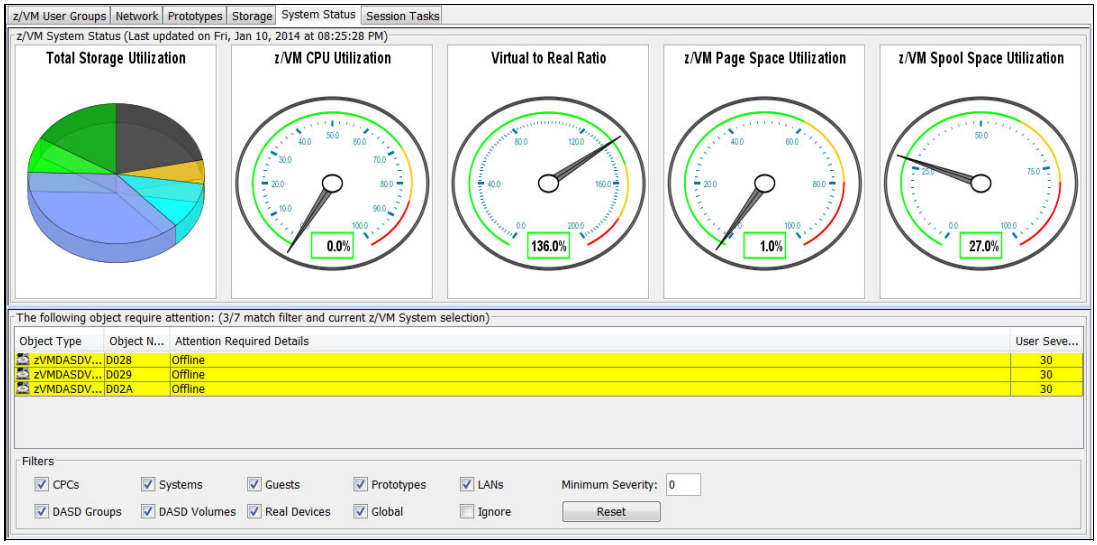


*Figure 5-30 System Status tab*

## Session Tasks tab

The Session Tasks tab displays activities that are running on multiple z/VM systems in parallel. When the environment consists of a single z/VM system, this tab always displays the No Session Tasks Exist string, as shown in Figure 5-31.

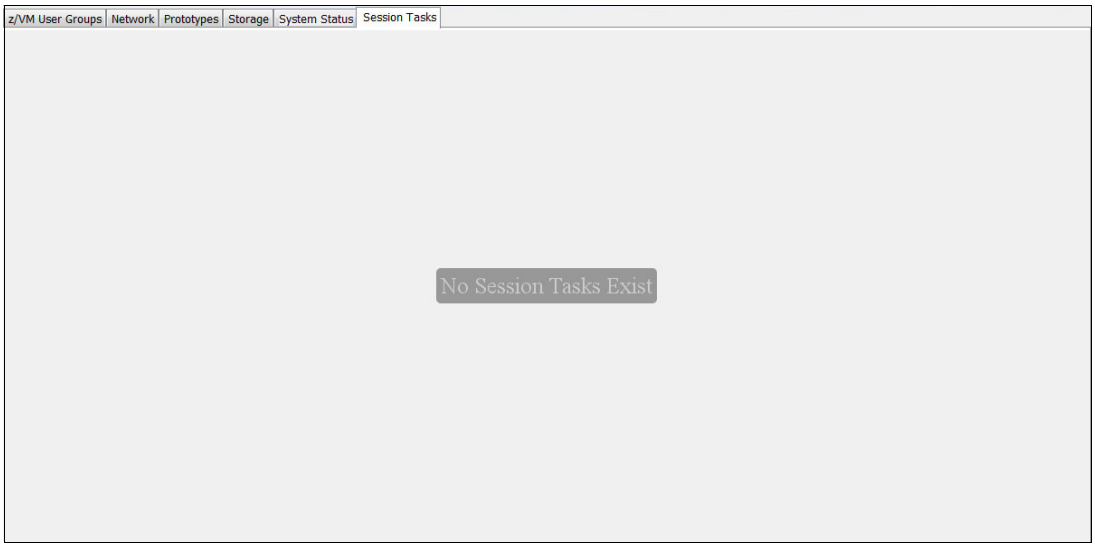


*Figure 5-31 Session Tasks tab*

## 5.3.3 Enterprise Viewer

The Enterprise Viewer and View differs from the Hardware Viewer and Current System View because it displays information across all instances of the z/VM hypervisor. You can use this function to see an enterprise-wide view of z/VM LPARs, as shown in Figure 5-32.

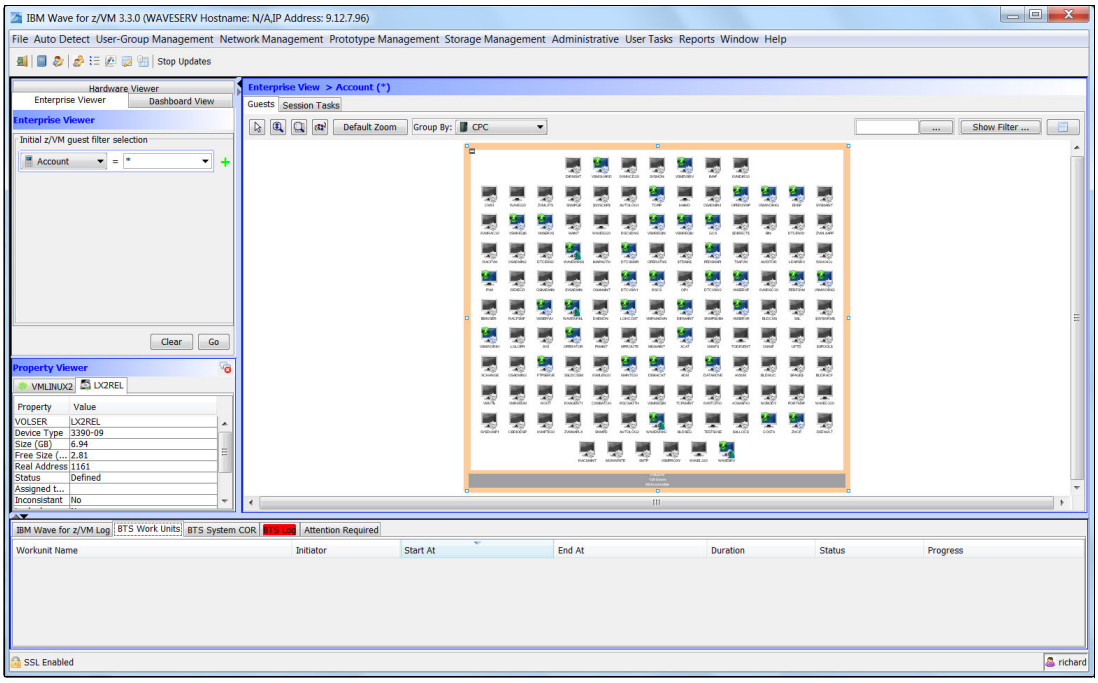


*Figure 5-32 Enterprise Viewer*

The left side of the Enterprise Viewer area consists of an area in which one or more filter criteria can be applied to the view in the Guests tab on the right. This method is an ideal way to visualize virtual resources across z/VM systems by project, account, or other attributes.

The functions in the Guests tab of the Enterprise View are identical to the ones in the User Group View of the Hardware View. The only exception is that the Enterprise View is the only location that can be used for the IBM Wave dragging implementation of LGR between z/VM systems.

## 5.3.4 Dashboard Viewer

The Dashboard Viewer and View displays a single window that shows performance data for all of the z/VM systems that are managed by IBM Wave, as shown in Figure 5-33 on page 109.

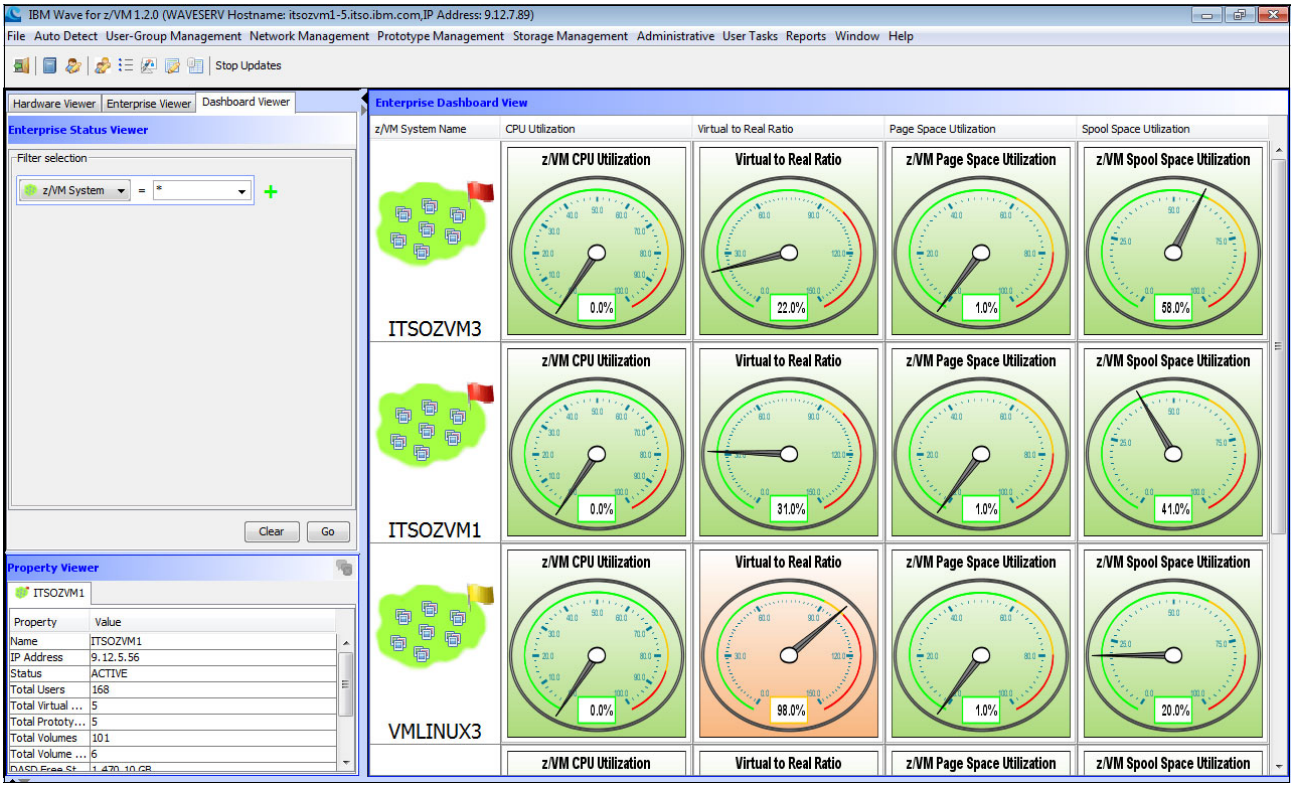


*Figure 5-33 Dashboard Viewer and View*

Double-clicking the z/VM system Name column or selecting **Expand All** or **Collapse All** alternates the view between the analog dials that are shown in Figure 5-33 and the compressed view that is shown in Figure 5-34, which can show more data on a single window.

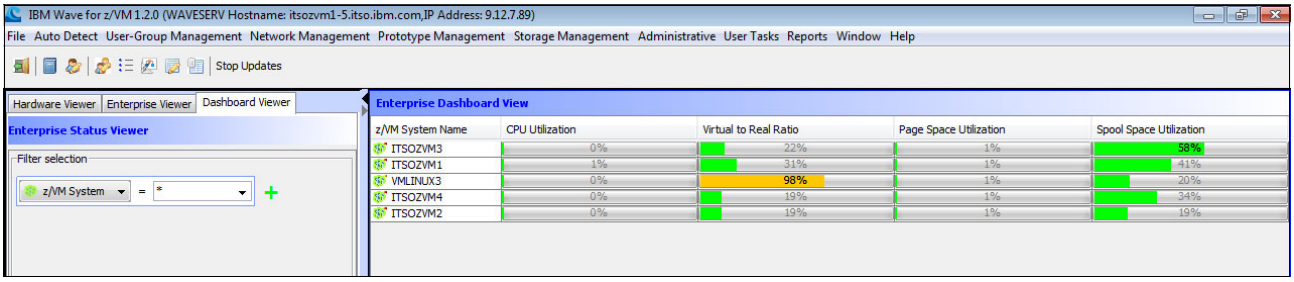


*Figure 5-34 Dashboard View Compressed*

Double-clicking the dials or % fields shows more detailed performance information and a textual view of the z/VM Performance Toolkit data that is used to build these views.

When the Dashboard View tab is open, the Viewer area on the left becomes a filter that can be used to limit the report to particular z/VM system or mainframe hardware.

# 5.4 IBM Wave Enviroment Viewers

The IBM Wave Environment Viewers are focused on IBM Wave concepts, such as messages, BTS data that is related to work units that are submitted by IBM Wave actions, and lists of IBM Wave objects that require manual intervention by an IBM Wave administrator.

The IBM Wave Log tab (Figure 5-35) displays messages that are received from the IBVM Wave server or users with the IBM Wave administrator role. The full message can be displayed by double-clicking the row of the wanted message. By double-clicking the IBM Wave Log tab, an IBM Wave user runs the View Log action with which users can search for messages that were received during previous sessions.

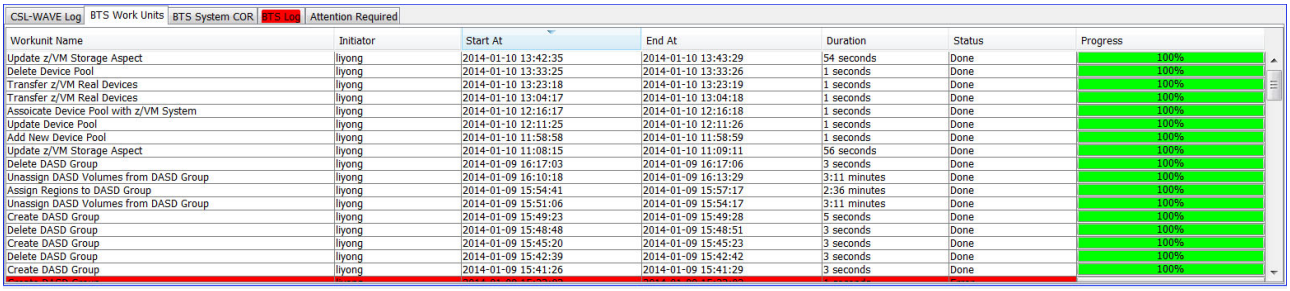
CSL-WAVE Log | BTS Work Units | BTS System COR | BTS Log | Attention Required

| Workunit Name | Initiator | Start At | End At | Duration | Status | Progress |
|---|---|---|---|---|---|---|
| Update z/VM Storage Aspect | liyong | 2014-01-10 13:42:35 | 2014-01-10 13:43:29 | 54 seconds | Done | 100% |
| Delete Device Pool | liyong | 2014-01-10 13:33:25 | 2014-01-10 13:33:26 | 1 seconds | Done | 100% |
| Transfer z/VM Real Devices | liyong | 2014-01-10 13:23:18 | 2014-01-10 13:23:19 | 1 seconds | Done | 100% |
| Transfer z/VM Real Devices | liyong | 2014-01-10 13:04:17 | 2014-01-10 13:04:18 | 1 seconds | Done | 100% |
| Assoicate Device Pool with z/VM System | liyong | 2014-01-10 12:16:17 | 2014-01-10 12:16:18 | 1 seconds | Done | 100% |
| Update Device Pool | liyong | 2014-01-10 12:11:25 | 2014-01-10 12:11:26 | 1 seconds | Done | 100% |
| Add New Device Pool | liyong | 2014-01-10 11:58:58 | 2014-01-10 11:58:59 | 1 seconds | Done | 100% |
| Update z/VM Storage Aspect | liyong | 2014-01-10 11:08:15 | 2014-01-10 11:09:11 | 56 seconds | Done | 100% |
| Delete DASD Group | liyong | 2014-01-09 16:17:03 | 2014-01-09 16:17:06 | 3 seconds | Done | 100% |
| Unassign DASD Volumes from DASD Group | liyong | 2014-01-09 16:10:18 | 2014-01-09 16:13:29 | 3:11 minutes | Done | 100% |
| Assign Regions to DASD Group | liyong | 2014-01-09 15:54:41 | 2014-01-09 15:57:17 | 2:36 minutes | Done | 100% |
| Unassign DASD Volumes from DASD Group | liyong | 2014-01-09 15:51:06 | 2014-01-09 15:54:17 | 3:11 minutes | Done | 100% |
| Create DASD Group | liyong | 2014-01-09 15:49:23 | 2014-01-09 15:49:28 | 5 seconds | Done | 100% |
| Delete DASD Group | liyong | 2014-01-09 15:48:48 | 2014-01-09 15:48:51 | 3 seconds | Done | 100% |
| Create DASD Group | liyong | 2014-01-09 15:45:20 | 2014-01-09 15:45:23 | 3 seconds | Done | 100% |
| Delete DASD Group | liyong | 2014-01-09 15:42:39 | 2014-01-09 15:42:42 | 3 seconds | Done | 100% |
| Create DASD Group | liyong | 2014-01-09 15:41:26 | 2014-01-09 15:41:29 | 3 seconds | Done | 100% |

*Figure 5-35 BTS Log tab*

The BTS Work Units tab (Figure 5-36) displays a list of the work units that is submitted to the BTS by all IBM Wave users. Work units can be double-clicked for more information, including a step-by-step list of what the work unit did to the system. A work unit can be deleted from the repository by right-clicking a work unit and selecting the Delete BTS Workunit action from the menu.

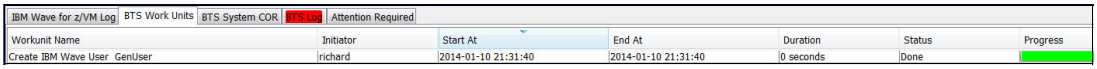
IBM Wave for z/VM Log | BTS Work Units | BTS System COR | BTS Log | Attention Required

| Workunit Name | Initiator | Start At | End At | Duration | Status | Progress |
|---|---|---|---|---|---|---|
| Create IBM Wave User GenUser | richard | 2014-01-10 21:31:40 | 2014-01-10 21:31:40 | 0 seconds | Done | |

*Figure 5-36 BTS Work Units*

The BTS System Common Output Repository (COR) tab displays a list of output messages from the BTS that are logged to the current COR log. The rows in this tab cannot be selected. Double-clicking the **BTS System COR** tab opens a window with which an IBM Wave user can browse the contents of older COR log files. Figure 5-37 shows the BTS Common Output Repository tab.

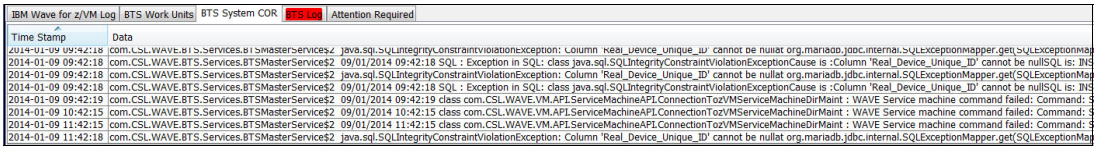
IBM Wave for z/VM Log | BTS Work Units | BTS System COR | BTS Log | Attention Required

| Time Stamp | Data |
|---|---|
| 2014-01-09 09:42:18 | com.CSL.WAVE.BTS.Services.BTSMasterService$2 java.sql.SQLIntegrityConstraintViolationException: Column 'Real_Device_Unique_ID' cannot be nullat org.mariadb.jdbc.internal.SQLExceptionMapper.get(SQLExceptionMap |
| 2014-01-09 09:42:18 | com.CSL.WAVE.BTS.Services.BTSMasterService$2 09/01/2014 09:42:18 SQL : Exception in SQL: class java.sql.SQLIntegrityConstraintViolationExceptionCause is :Column 'Real_Device_Unique_ID' cannot be nullSQL is: INS |
| 2014-01-09 09:42:18 | com.CSL.WAVE.BTS.Services.BTSMasterService$2 java.sql.SQLIntegrityConstraintViolationException: Column 'Real_Device_Unique_ID' cannot be nullat org.mariadb.jdbc.internal.SQLExceptionMapper.get(SQLExceptionMap |
| 2014-01-09 09:42:18 | com.CSL.WAVE.BTS.Services.BTSMasterService$2 09/01/2014 09:42:18 SQL : Exception in SQL: class java.sql.SQLIntegrityConstraintViolationExceptionCause is :Column 'Real_Device_Unique_ID' cannot be nullSQL is: INS |
| 2014-01-09 09:42:19 | com.CSL.WAVE.BTS.Services.BTSMasterService$2 09/01/2014 09:42:19 class com.CSL.WAVE.VM.API.ServiceMachineAPI.ConnectionTozVMServiceMachineDirMaint : WAVE Service machine command failed: Command: S |
| 2014-01-09 10:42:15 | com.CSL.WAVE.BTS.Services.BTSMasterService$2 09/01/2014 10:42:15 class com.CSL.WAVE.VM.API.ServiceMachineAPI.ConnectionTozVMServiceMachineDirMaint : WAVE Service machine command failed: Command: S |
| 2014-01-09 11:42:15 | com.CSL.WAVE.BTS.Services.BTSMasterService$2 09/01/2014 11:42:15 class com.CSL.WAVE.VM.API.ServiceMachineAPI.ConnectionTozVMServiceMachineDirMaint : WAVE Service machine command failed: Command: S |
| 2014-01-09 11:42:18 | com.CSL.WAVE.BTS.Services.BTSMasterService$2 java.sql.SQLIntegrityConstraintViolationException: Column 'Real_Device_Unique_ID' cannot be nullat org.mariadb.jdbc.internal.SQLExceptionMapper.get(SQLExceptionMap |

*Figure 5-37 BTS COR*

The BTS Log tab contains a user-readable log of the processing that is done by the BTS. By right-clicking the text and selecting **Scroll to Bottom when Messages Arrive**, this log functions as a console and automatically scrolls as messages arrive. Double-clicking the **BTS Log** tab opens the BTS Manager, which configures the settings for the BTS and stores all BTS logs. Figure 5-38 on page 111 shows the BTS Log Tab.

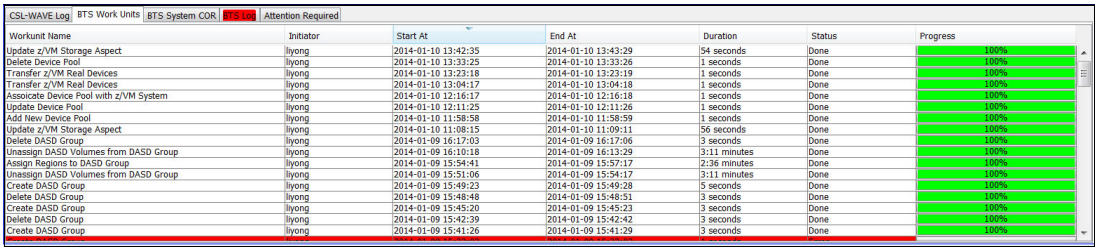


*Figure 5-38 BTS Log tab*

The Attention Required tab (Figure 5-39) displays a list of all the actions that IBM Wave believes require manual IBM Wave administrator intervention. Right-clicking a row opens a menu with which an IBM Wave administrator can ignore the issue or change the numeric severity level.

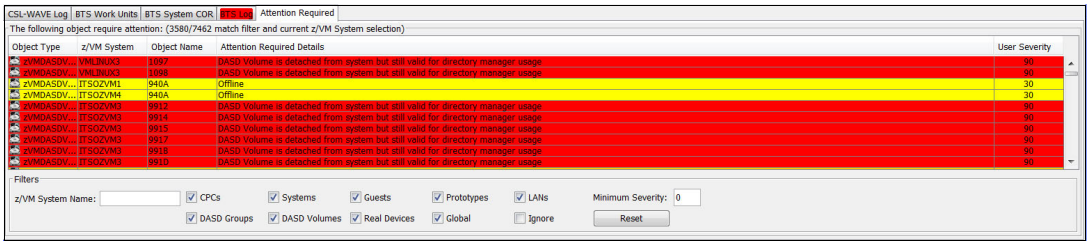


*Figure 5-39 Attention Required tab*

## 5.5 The Status Bar

The Status Bar displays the items that are shown in Figure 5-40.

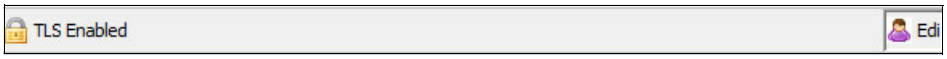


*Figure 5-40* Status Bar

From left to right, the Status Bar includes the following features:

- A Message Area that displays messages from the IBM Wave client. In this case, it indicates that SSL was enabled.
- A cell that shows the current IBM Wave user.

## 5.6 Group management

IBM Wave is designed to help system programmers and system administrators manage complex z/VM environments. A large z/VM system might consist of hundreds or thousands of VMs. Tracking all of these VMs and managing them requires some mechanism to summarize what is in the directory. IBM Wave accomplishes this task by assigning VMs to groups during the auto detect process. It then provides the capability to create site-defined groups for better granularity.

The grouping capability exists only within the context of IBM Wave and does not involve any modifications to a managed z/VM system. The default groups for a small z/VM V6.3 system are shown in Figure 5-41.

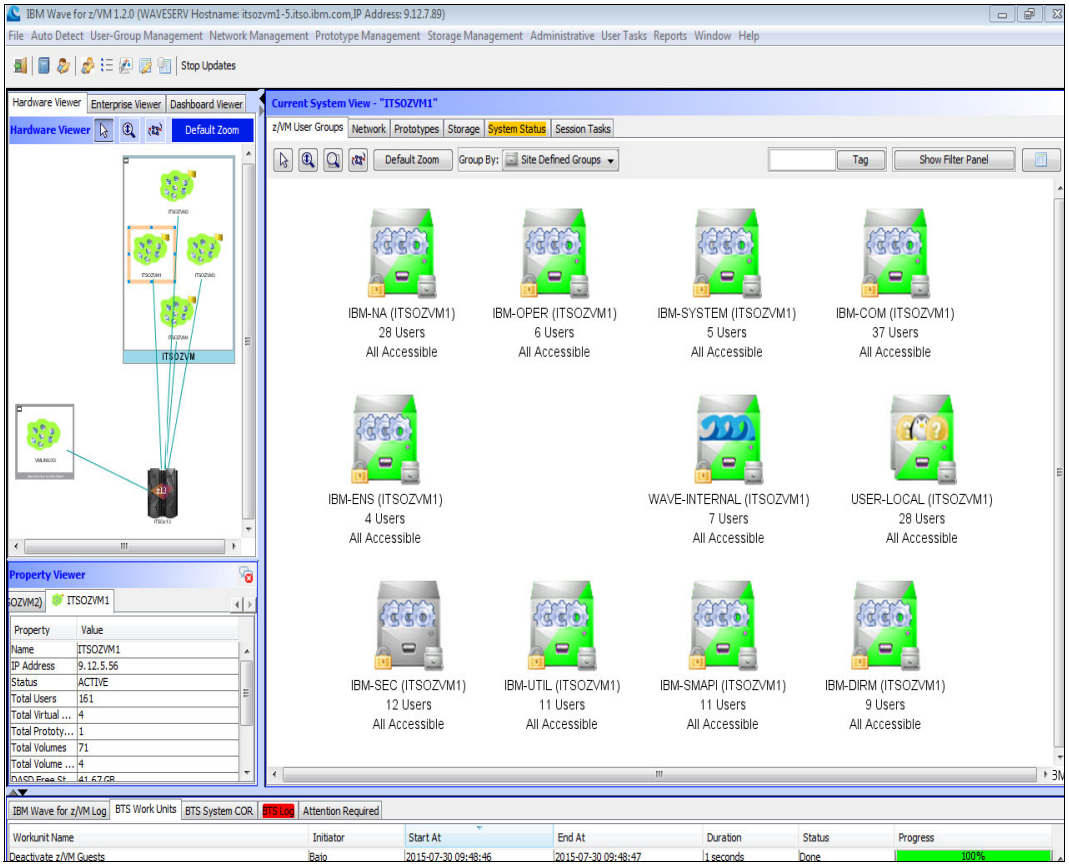


*Figure 5-41 Default site defined groups*

IBM Wave places VMs that are part of the z/VM product installation into particular groups so that related VMs are kept together. For example, the IBM SMAPI group contains all of the z/VM system Management API requester and worker VMs. The USER-LOCAL group contains all of the VMs that IBM Wave could not place in a specific group during the auto detect process. On a large z/VM system, this group might contain many VMs. Often, it is the content of this group that you move to groups that you create to better categorize the VMs that are running on a managed z/VM system.

To create a group, right-click any portion of the white background of the z/VM User Groups tab. A menu opens that includes a selection to Add New Site Defined Group, as shown in Figure 5-42 on page 113.

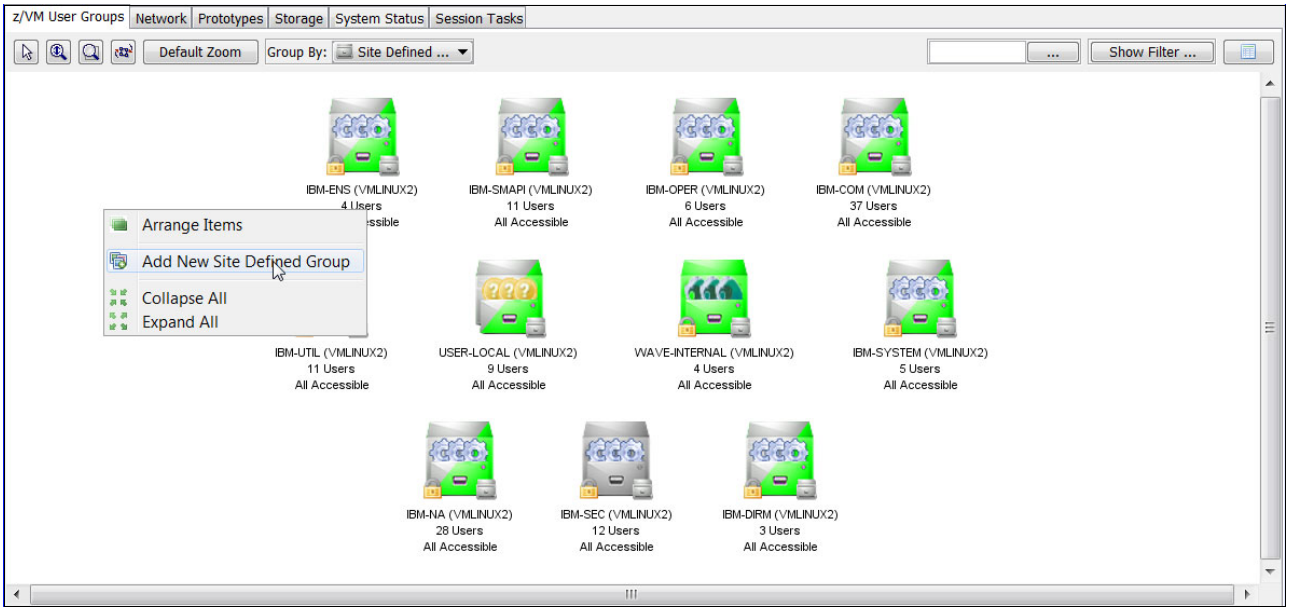


*Figure 5-42 Add New Site Defined Group option*

Clicking the **Add New Site Defined Group** option opens a window in which the user enters information about the new group, as shown in Figure 5-43.

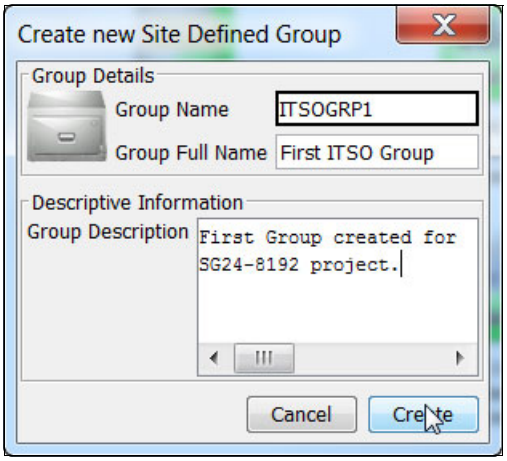


*Figure 5-43 Create new Site Defined Group window*

In the window, you must enter a name for the new group (completing the Group Full Name or Description fields is optional). When you click **Create**, a group is created and displayed in the z/VM User Groups tab, as shown in Figure 5-44.

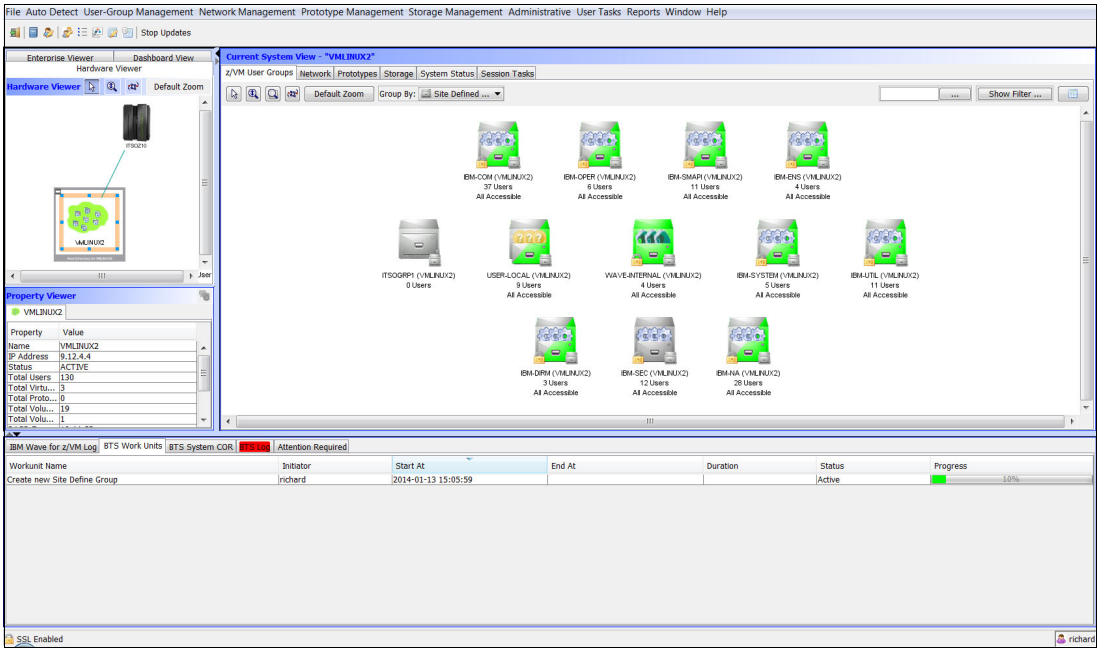

*Figure 5-44 Group is created*

After a group is created, VMs can be created in it and existing VMs can be moved from an existing group to the new group. To move an existing VM to the new group, right-click the VM and select **Update** → **Assign Site Defined Group**, and select the group from the list, as shown in Figure 5-45.

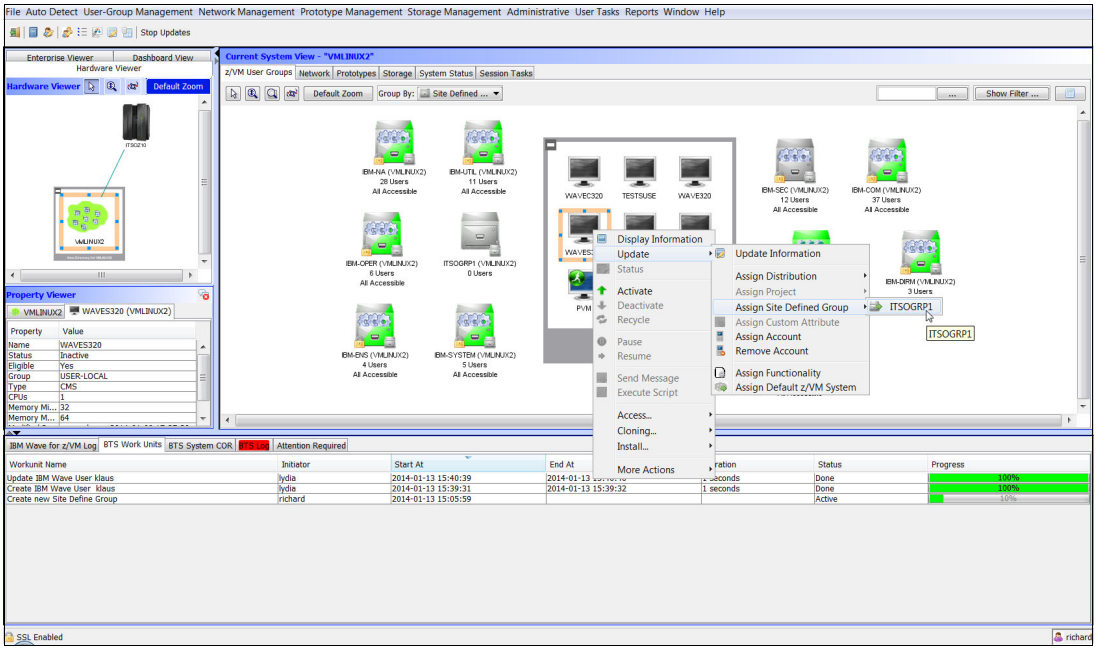

*Figure 5-45 Assign Site Defined Group option*

You also can expand the existing and new groups, select the VM from the existing group, and hold down the Ctrl key and drag the VM icon to the new group, as shown in Figure 5-46.

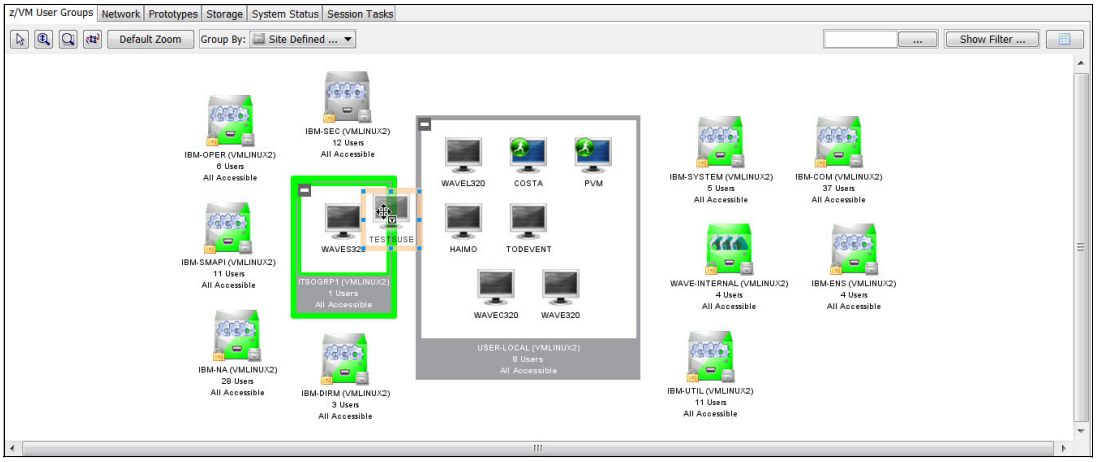


*Figure 5-46 Drag to a new group*

**Note:** A VM can belong to only one group.

IBM Wave groups are local to a particular z/VM system. In an SSI cluster, you must define those groups in every system participating in the SSI. The next section describes the organization of VMs across z/VM systems and CPCs.

## 5.7 Project management

Because site-defined groups exist only within the context of a single z/VM system, such groups are of limited use for organizing VMs in clustered environments or environments that span multiple LPARs or CPCs. For these types of complex environments, IBM Wave provides the concept of projects, with which users can categorize and logically group VMs and other virtual resources across multiple z/VM systems and CPCs.

Projects are also the primary means by which IBM Wave limits the scope of IBM Wave users. Projects are created by using the Project Manager option from the Administrative menu in the menu bar, as shown in Figure 5-47.

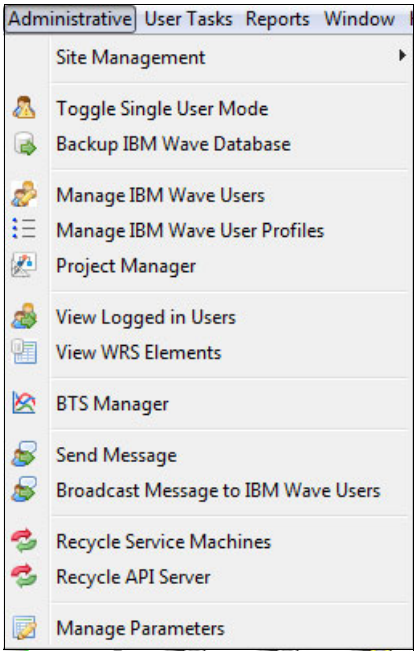


*Figure 5-47 Administrative menu options*

Complete the following steps:

1. Click **Project Manager** to open the Project Manager window in which all of the projects that are defined are listed. Click **Add New Project** to create a project, as shown in Figure 5-48.

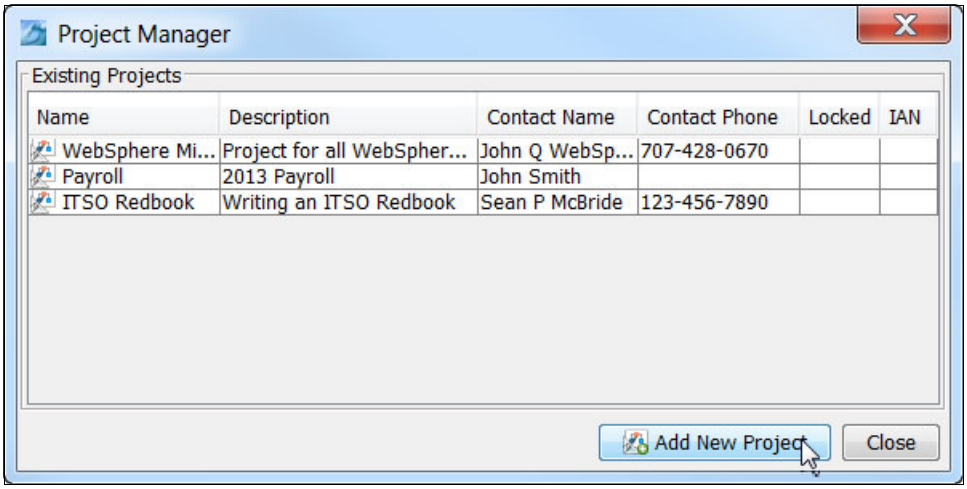


*Figure 5-48 Project Manager window*

2. In the Add New Project window, enter a name, description, contact name, phone, and email for the project. Enter also the Project Team and Project Department. Click **Add** and the system creates the project and returns you to the Project Manager window, as shown in Figure 5-49 on page 117.

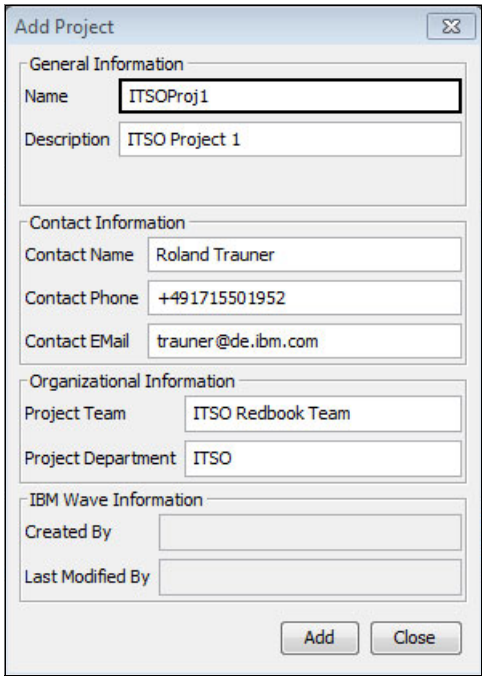


*Figure 5-49 Add New Project window*

The new project is now listed in the Project Manager window, as shown in Figure 5-50.

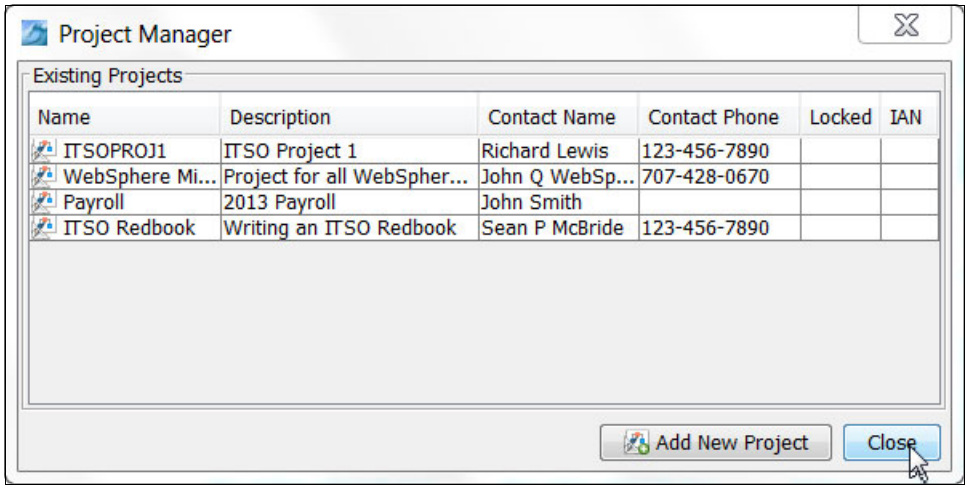


*Figure 5-50 Project added in the Project Manager window*

3. After a project is created, you can assign VMs to the project and give IBM Wave users permission to access that project. To assign a VM to a project, right-click the icon for a particular VM, select **Update** → **Assign Project**, and then select a project from the list that is displayed, as shown in Figure 5-51.

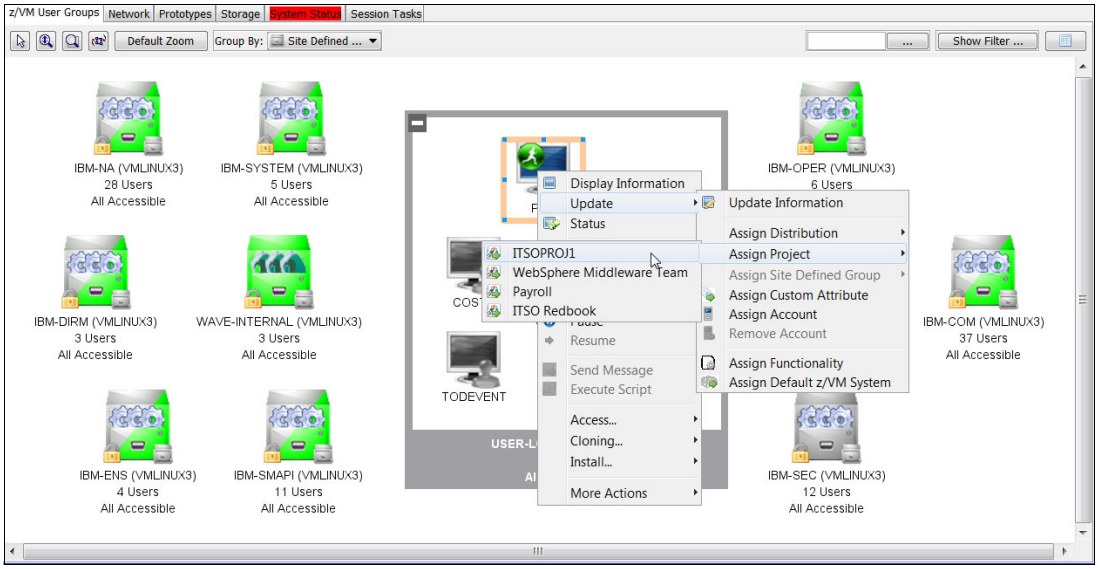


*Figure 5-51 Assign a virtual machine to a project*

When a project is selected, a BTS transaction is started to associate the VM with the selected project. The benefit of assigning VMs to projects is seen when you are using the Enterprise viewer with the Group by filter set to Project, as shown in Figure 5-52.

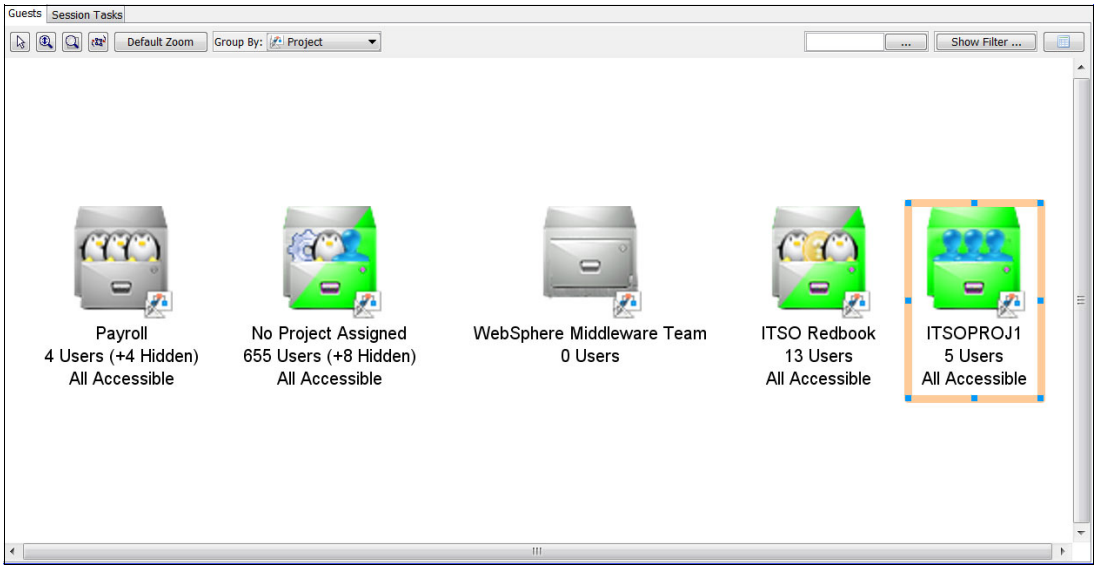


*Figure 5-52 Enterprise Viewer projects*

4. The project that is named No Project Assigned contains all of the VMs throughout the enterprise that are not yet assigned to projects. To see the VMs that are assigned to the project that was created (which is named ITSOPROJ1), expand the project icon, as shown in Figure 5-53 on page 119.

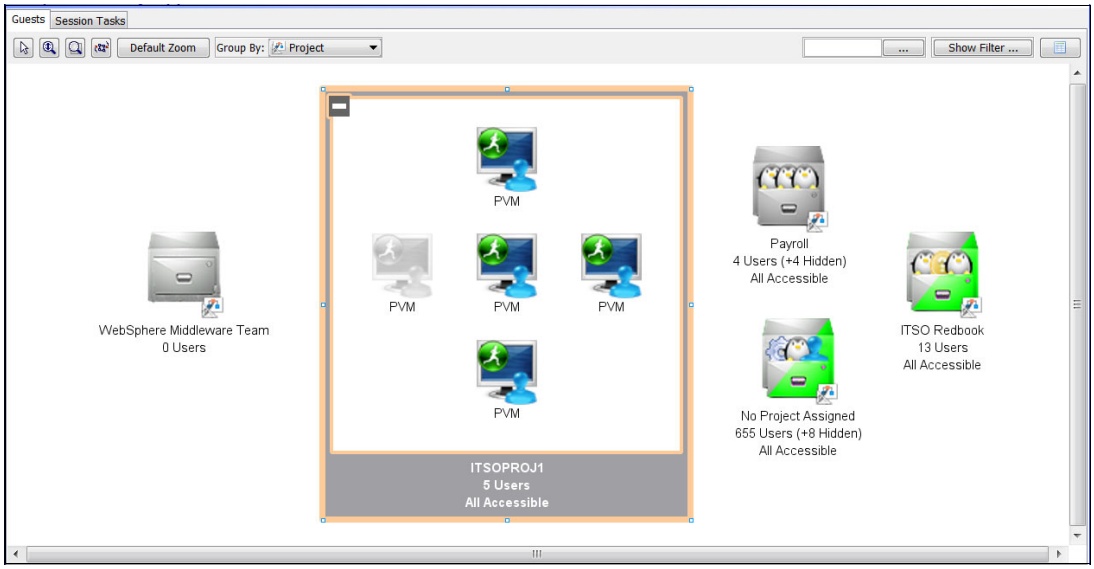


*Figure 5-53 Expand the project icon*

Each of the VMs that are shown in Figure 5-53 are from different z/VM LPARs.

## 5.7.1 Automatic guest classification

When you implement projects to help organize and manage VMs across your enterprise, one of the challenges becomes maintaining the association between VMs and projects. The Automatic Guest Classification (AGC) capability of IBM Wave can assist you with this task. With AGC, you can set up an association between z/VM account codes that are assigned to VMs with the **ACCOUNT** directory statement and projects that are defined within IBM Wave by using the Project Manager.

AGC then automatically changes the project to which a VM is associated if a change occurs to the z/VM account code that is associated with the VM. If a change to the project that is associated with a particular account code changes, IBM Wave ensures that the directory information for the VMs that are associated with that project is updated to reflect the new value on the **ACCOUNT** directory statement.

To use the AGC function, complete the following steps:

1. When IBM Wave is installed, the AGC function is not enabled by default. Before you can use this capability, you must enable the AGC function by clicking **Administration** from the menu bar, and then selecting **Manage Parameters**. Select the **Functionality** tab at the top of the Manage Parameters tab, and then select the **Automatic Guest Classification** option that is near the bottom of the window, as shown in Figure 5-54.

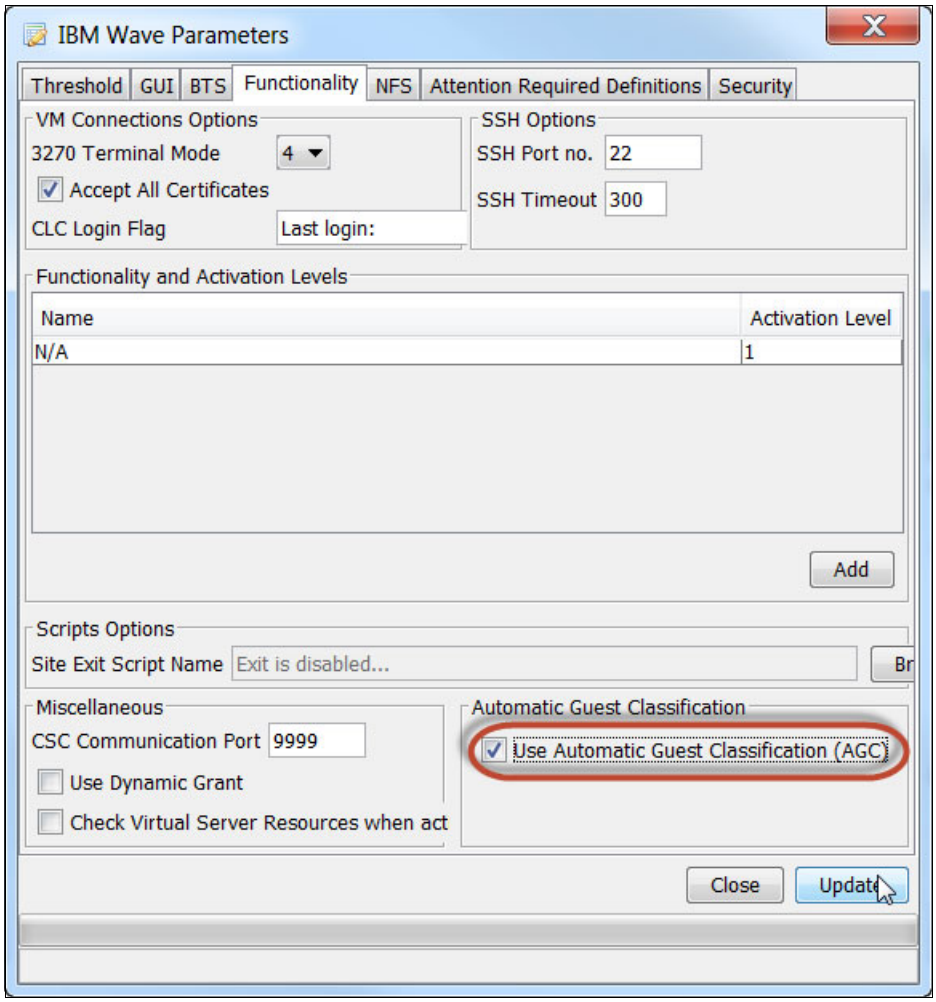


*Figure 5-54 Manage parameters by using AGC*

2. Click **Update** to save the changes and activate the AGC capability.

   Before you define AGC entries, you must ensure that the projects that are to be associated with z/VM **ACCOUNT** record values are defined. Section 5.7 “Project management” on page 115 describes creating projects. In our example, we now use the AGC Manager to create the association between an **ACCOUNT** value and a particular project.

3. Start the AGC Manager by clicking **Administration** → **Site Management** → **AGC Manager** from the menu bar, as shown in Figure 5-55 on page 121.

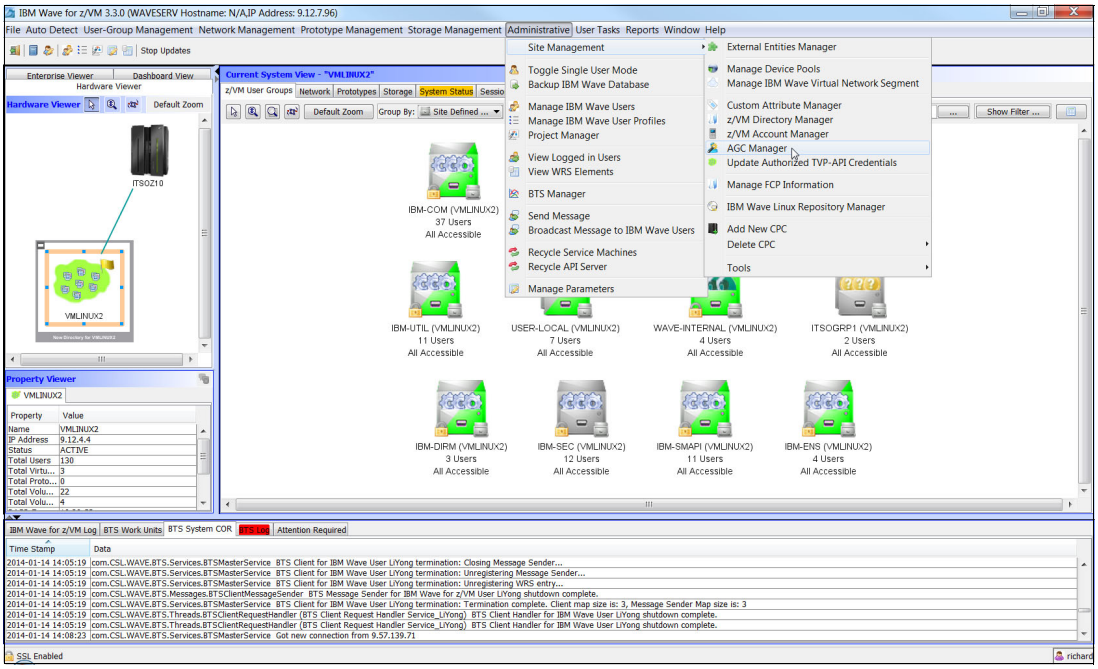


*Figure 5-55 Select AGC Manager*

4. When you click **AGC Manager**, a new window opens. Click **Create New AGC Entry** to define a new criteria for automatic guest classification, as shown in Figure 5-56.

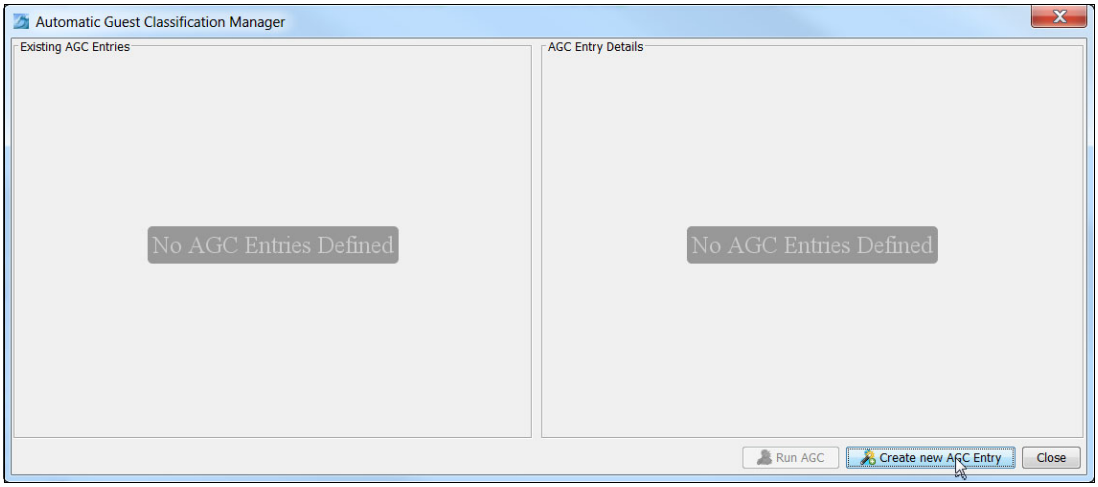


*Figure 5-56 AGC Manager window*

5. A window opens in which you enter the AGC information. Enter a name for the entry, and then select an AGC Value from the selection box. The values are all of the various values that were discovered on **ACCOUNT** statements when the USER Directory was processed during the auto detection of the particular z/VM system.

6. In our example, we want to associate the **ACCOUNT** value of SERVICE with a project that is named Service. Entering a description of the AGC entry is optional. Right-click in the area at the bottom of the window to open the Create New AGC Property association option. Clicking this option opens the Create New AGC Entry window, in which you can select the project to associate with the AGC value that is selected, as shown in Figure 5-57.

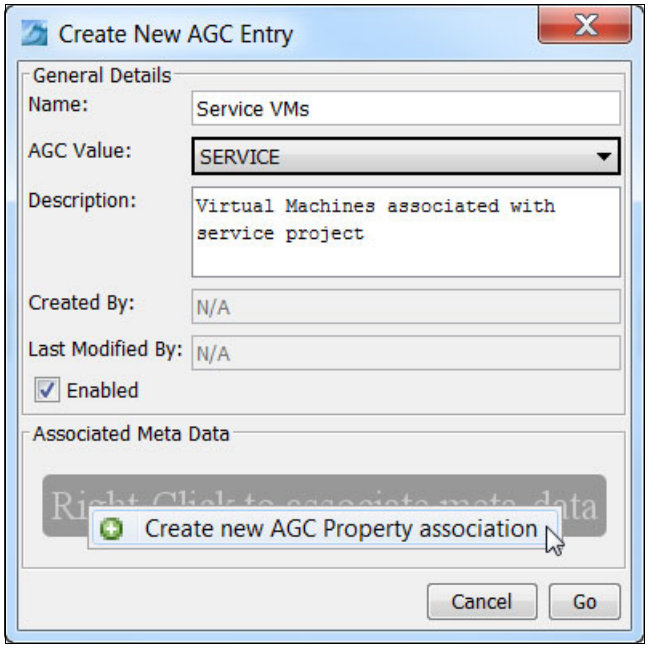


*Figure 5-57 Create New AGC Entry window*

7. In the AGC Entry data properties window, you can choose from Project, OS Distribution, Default z/VM system, Lock, Site Defined Group, or Functionality. Functionality is an attribute that is assigned to VMs, which is defined in the global IBM Wave parameters. In our example, we choose the Project Value of Service to associate with the AGC entry. Click **Add** to return to the Create New AGC window, as shown in Figure 5-58.

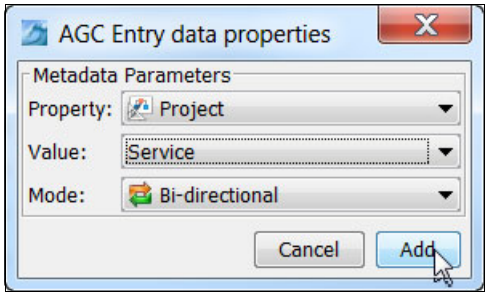


*Figure 5-58 AGC Entry Data properties window*

In the Create New AGC window, the information that you entered is now displayed, as shown in Figure 5-59 on page 123.

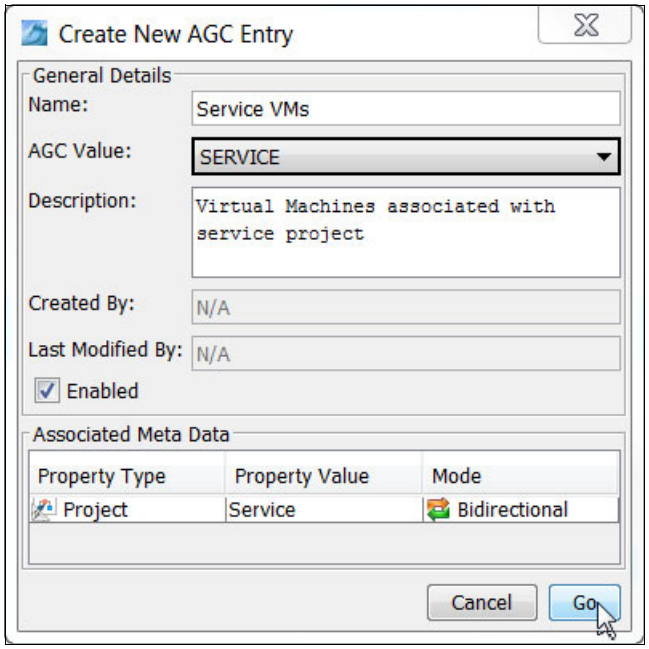


*Figure 5-59 Create AGC Entry*

8. When you click **Go**, the new AGC entry is created and displayed on the left side of the AGC Manager window. Click the entry to see the details that are displayed on the right side of the AGC Manager window.
9. After the AGC entry is created, you can click **Run AGC** (Figure 5-60) to submit a BTS entry to process all VMs and apply the defined automatic classification criteria.

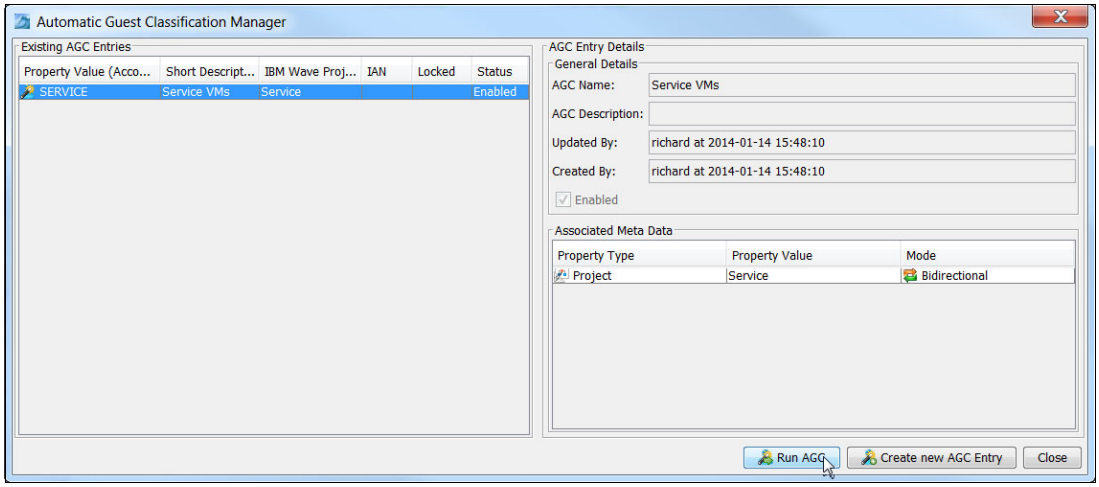


*Figure 5-60 Run AGC*

A confirmation window opens in which you are prompted to confirm that you want to run AGC against all existing z/VM guests, as shown in Figure 5-61.

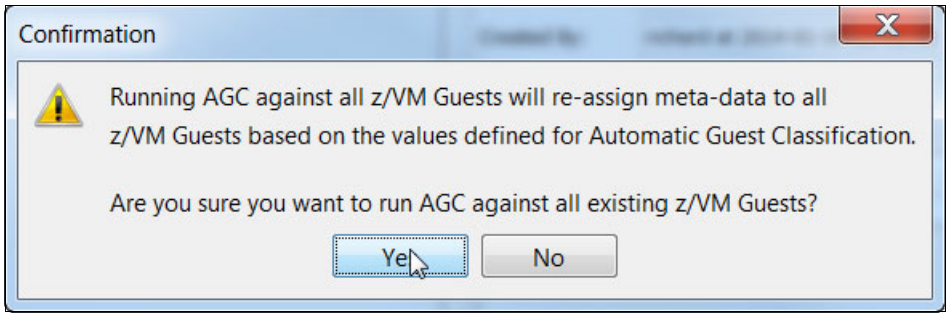


*Figure 5-61 Run AGC Confirmation*

10. Click **Yes** to proceed with running AGC against all z/VM VMs. You also can run AGC against one or a small group of VMs by selecting them from the Enterprise Viewer or from a group in the Site Defined Groups window. Right-click the guests and select **More Actions** → **Run AGC**.

To see the results of running AGC, it is easiest to switch the Group By filter on the Current System display from Site **Defined Groups** to **Project**. By making this switch, you can see the VMs that were added to the project that are associated with the AGC Entry. In our example, when we display the Service Project, we see a single VM that is named TESTSUSE was added. When we display the information for that VM, we see that it is now associated with this project, as shown in Figure 5-62.

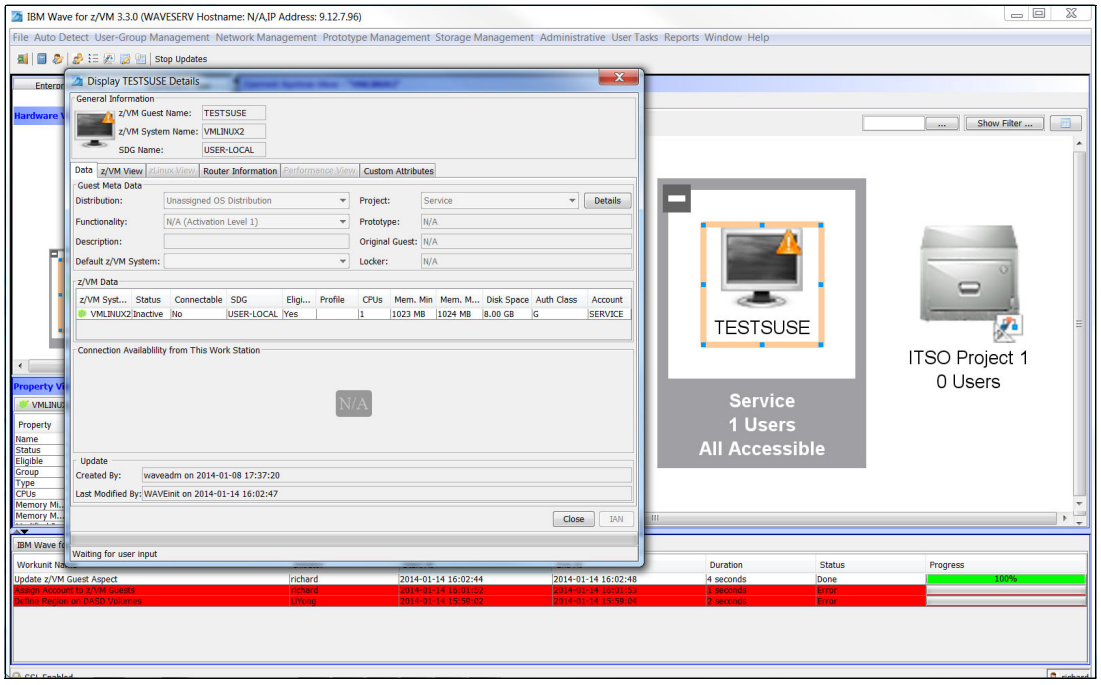


*Figure 5-62 View Project Service*

Because we set up our AGC Entry with bidirectional synchronization, if we update the account information of another VM to `SERVICE`, it also automatically appears in this Project folder. Figure 5-63 on page 125 shows the result of updating the WAVE320 VM **`ACCOUNT`** to `SERVICE`.

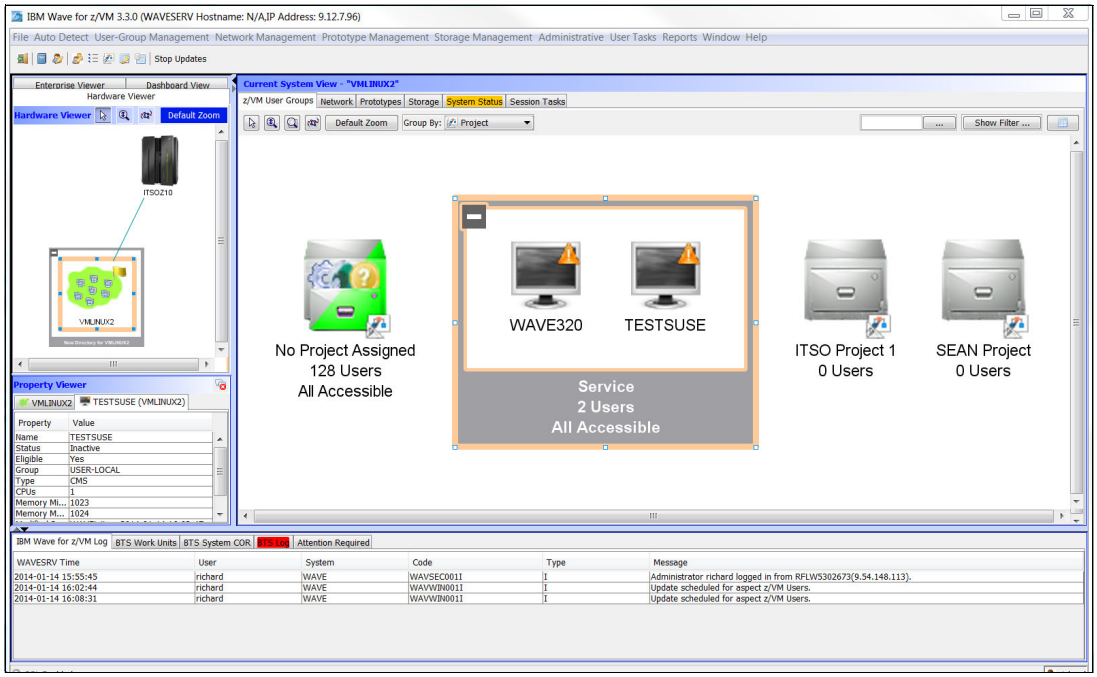


*Figure 5-63 Change virtual machine ACCOUNT*

With this AGC entry, when the primary **ACCOUNT** value for a VM is changed to the value SERVICE, it automatically is associated with Project Service. If we change the AGC Value in the AGC Entry to a different **ACCOUNT** value later, IBM Wave changes the primary **ACCOUNT** value in the z/VM Directory to the new value for all of the VMs in Project Service. If this behavior is not what is wanted, you can change the AGC Entry to use a synchronization method that is different from the default of bidirectional.

## 5.7.2 Defining custom attributes

Custom attributes apply to all virtual servers. If a defined attribute is not assigned to a virtual server, the value is the default value Not Assigned.

To specify Custom Attributes and their values, complete the following steps:

1. Click **Custom Attribute Manager**, as shown in Figure 5-64.

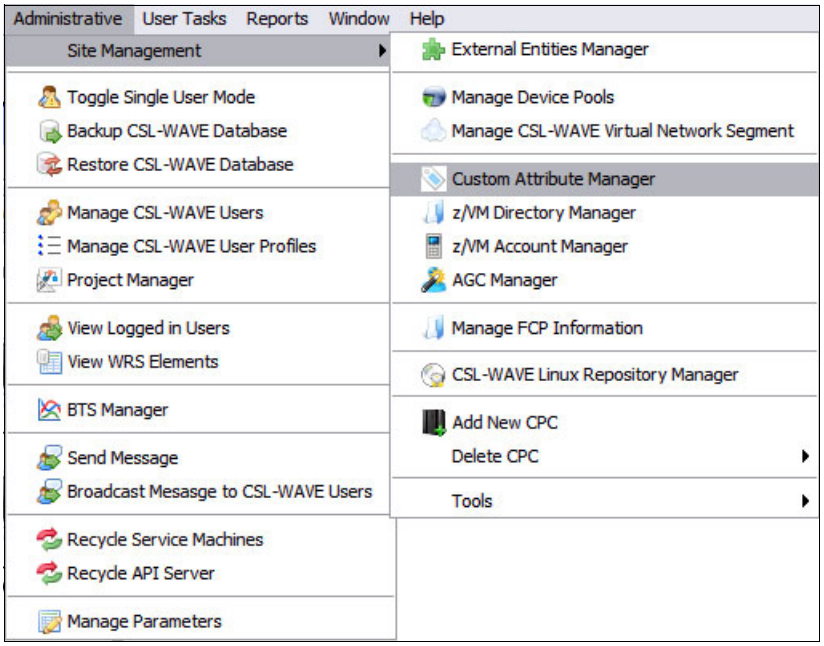


*Figure 5-64 Custom Attribute Manager*

2. Click **Add Custom Attribute** and define an attribute; for example, Company, and its possible values, as shown in Figure 5-65.

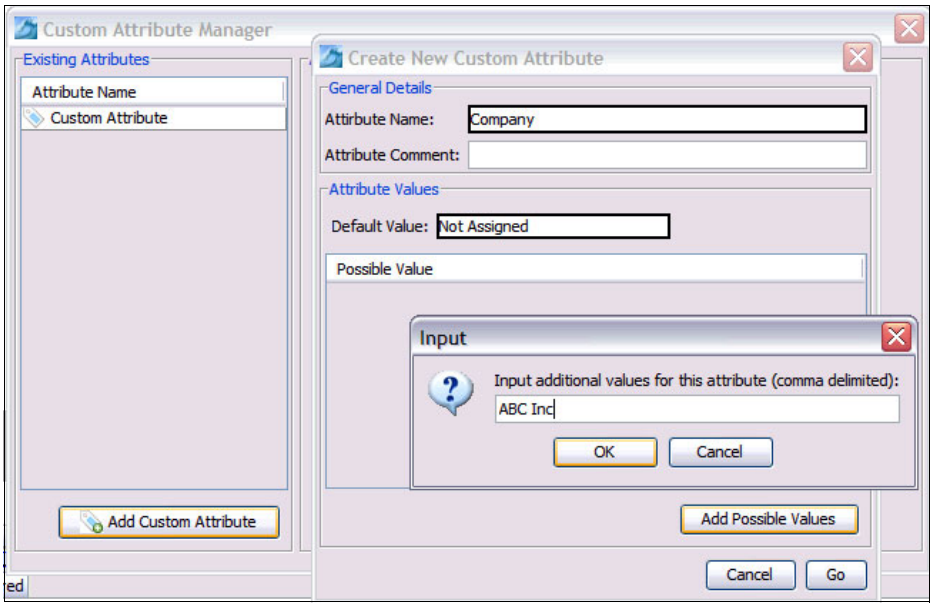


*Figure 5-65 Set values for a custom attribute*

3. To change a custom attribute or its values, from the Custom Attribute Manager, right-click the custom attribute to be changed, delete the attribute value by right-clicking the value (Figure 5-66), and add the correct value.

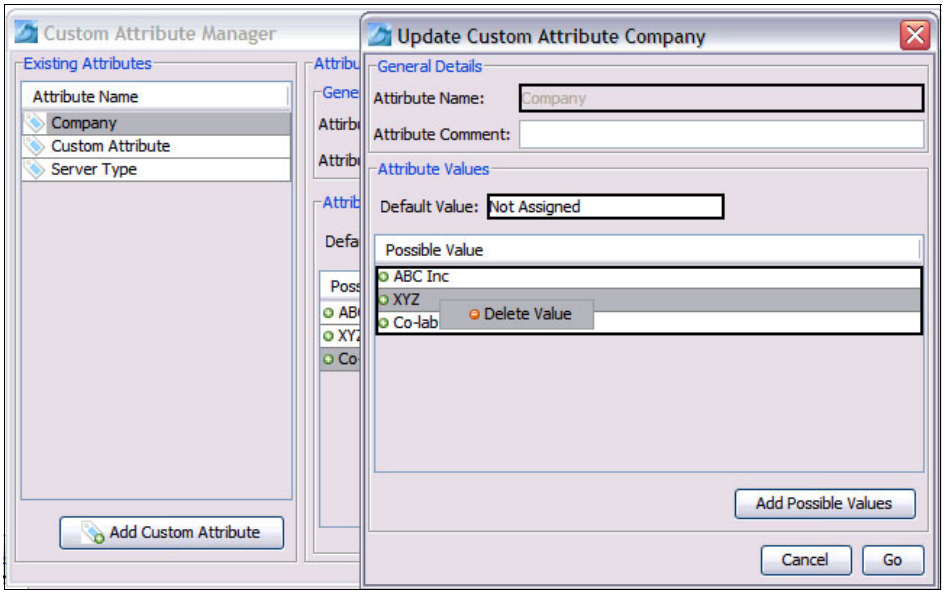


*Figure 5-66 Update custom attributes*

4. To assign a custom attribute to a virtual server, right-click the virtual server in the z/VM guest view and select **Update** → **Assign Custom Attribute**, as shown in Figure 5-67.

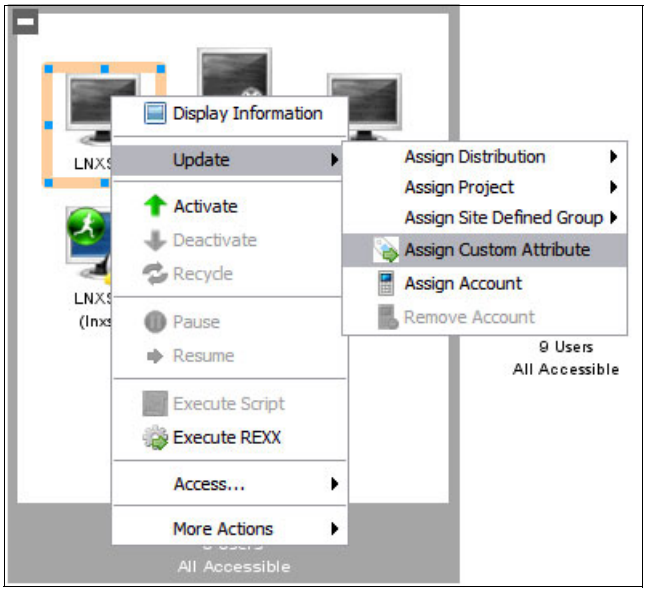


*Figure 5-67 Assign custom attribute*

5. In the Assign Custom Attributes to z/VM guests window, select the wanted attribute and value and click **Go**, as shown in Figure 5-68. Each attribute assignment sends a request to the BTS server.

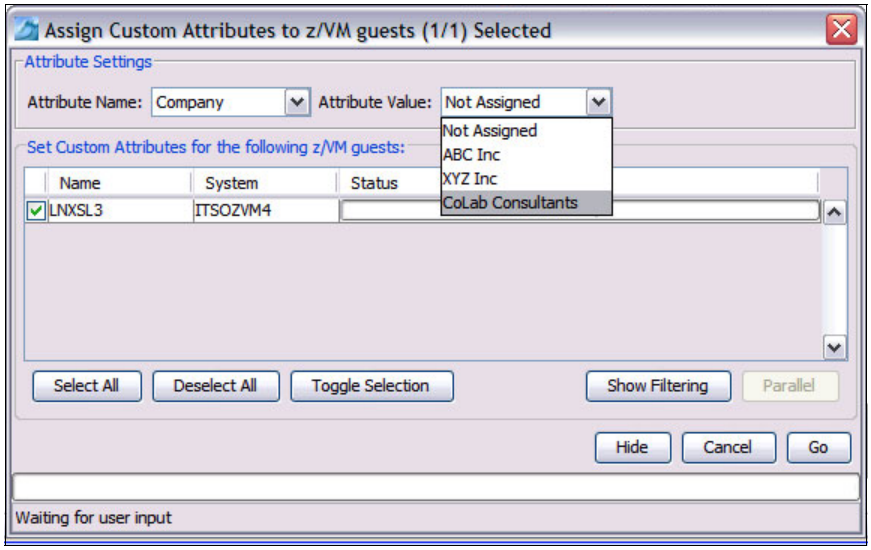


*Figure 5-68 Assign custom attributes to z/VM guests*

6. To display with filtering by custom attributes, use the Enterprise Viewer to display guests across all LPARs that are based on the attributes that are selected. In Figure 5-69, the Enterprise Viewer is used to display guests that are associated with Company = Colab Consultants, with Server Type = Production Server type. In the Guests tab, the Grouped by z/VM system option is selected and a filter of Show Empty Groups = No was selected so that LPARs with no entries are not displayed.

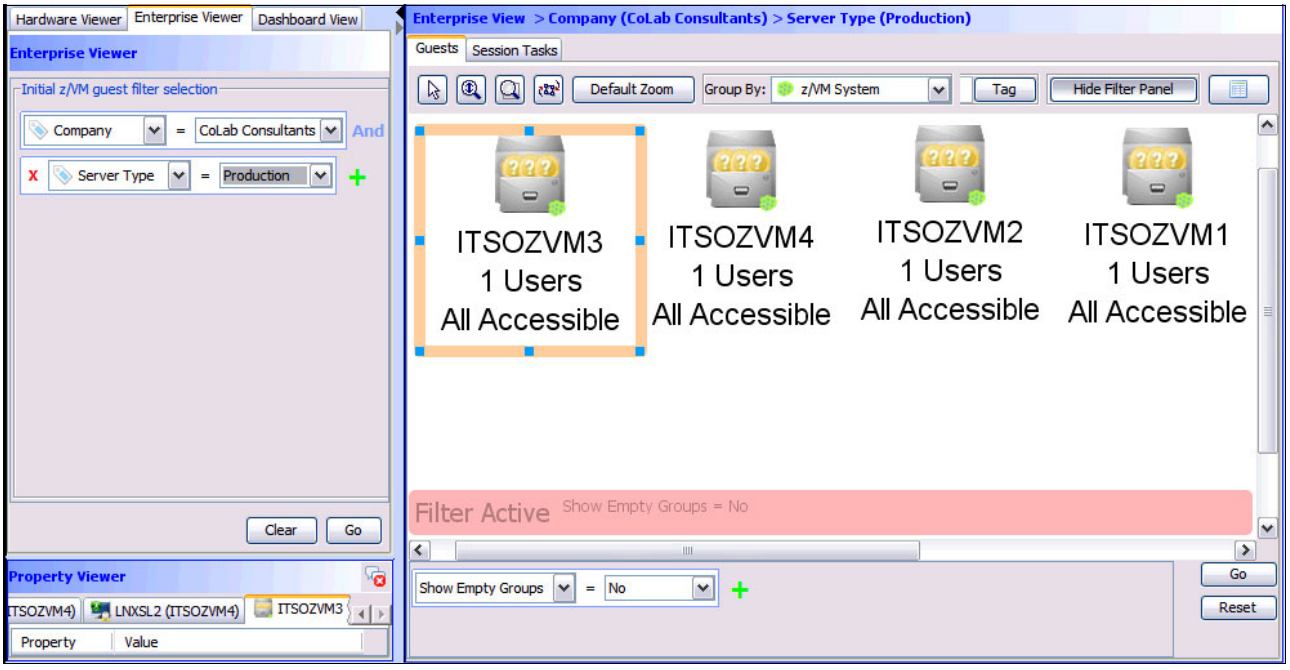


*Figure 5-69 Company Production Servers by z/VM systems*

## 5.8 Creating reports

IBM Wave includes a facility to generate reports that provide information about managed z/VM systems. By using IBM Wave Report Manager, you can generate two types of reports: global reports and private reports. User scope and permissions determine the objects that are included in your reports.

You must have site-level administration authority to save defined global reports. After they are saved, global reports can be used by all IBM Wave users. Private reports can be generated by an IBM Wave user and only the creator of the report can view, edit, and delete the report. If you have site-level administration authority, you can convert global reports to private reports, or convert private reports to global reports.

Creating and generating reports or generating saved reports is done by using the Report Manager. Complete the following steps:

1. Start the Report Manager by clicking **Reports** → **Manage Report Templates** in the menu bar, as shown in Figure 5-70.

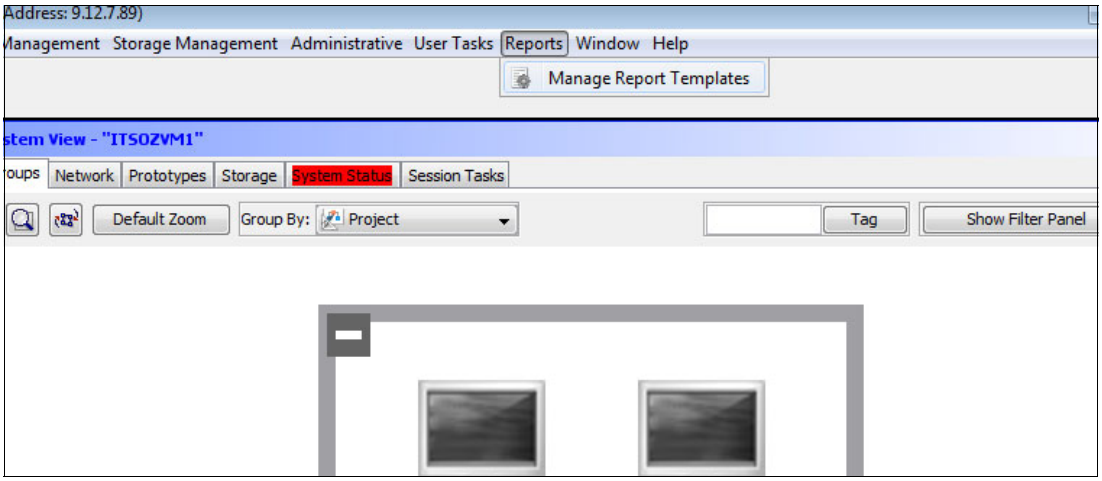


*Figure 5-70 Report Manager Navigation*

Clicking **Manage Report Templates** opens the Report Manager window, as shown in Figure 5-71.

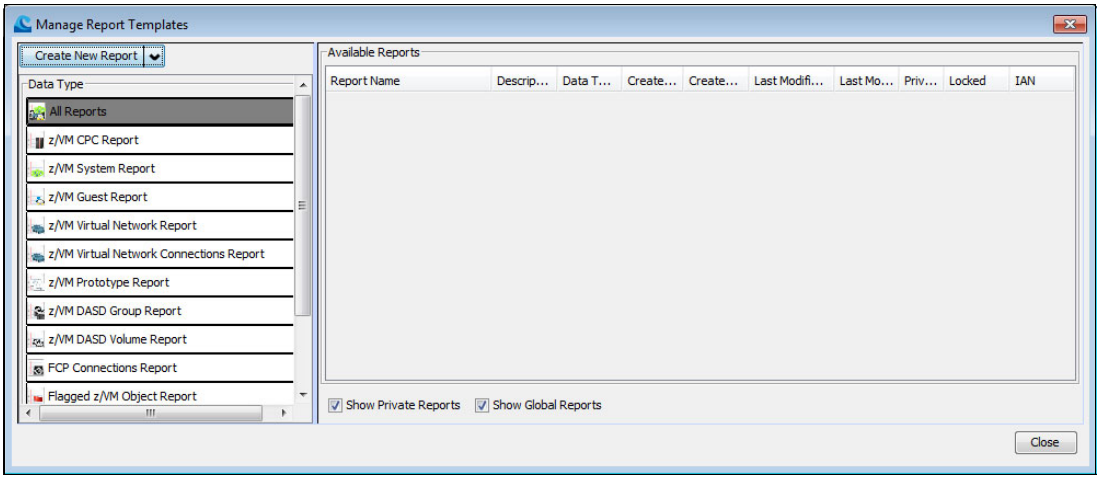


*Figure 5-71 Report Manager window*

The Manage Report Templates window shows, in the left pane, the various reports that can be created, and then lists any saved reports in the body of the window. The following report types are available:

- z/VM CPC Report: Shows information about all CPCs that are under IBM Wave management.
- z/VM system Report: Shows information about all z/VM LPARs that are defined to IBM Wave.
- z/VM guest Report: Shows information about each VM across all IBM Wave managed z/VM systems or a particular z/VM system.
- z/VM Virtual Network Report: Shows virtual switch and guest LAN definition information.
- z/VM Virtual Network Connection Report: Shows network-related information for virtual switches and guest LANs.
- z/VM Prototype report: Shows information that is related to prototypes (such as a VM that is associated with the prototype) and other descriptive information.
- z/VM DASD Group report: Shows information about all DASD groups that are defined across all managed z/VM systems.
- z/VM DASD Volume report: Shows information about all DASD volumes that are available in all managed z/VM systems.
- FCP Connections Report: Shows information about all FCP Connections that are available.
- Flagged z/VM Object report: Shows a list of all objects that require attention in all managed z/VM systems.
- Custom Attributes Report: Shows a list of Customized attributes, such as Company name and other.
- IBM Wave User report: Lists all of the users that are defined to IBM Wave along with the attributes that are associated with each user.
- IBM Wave User Permission report: Lists access level, permissions, and scopes for all IBM Wave users or lists IBM Wave users with access to a particular project or other managed object.

2. To create a report, click **Create New Report** and then select the report from the menu, as shown in Figure 5-72.

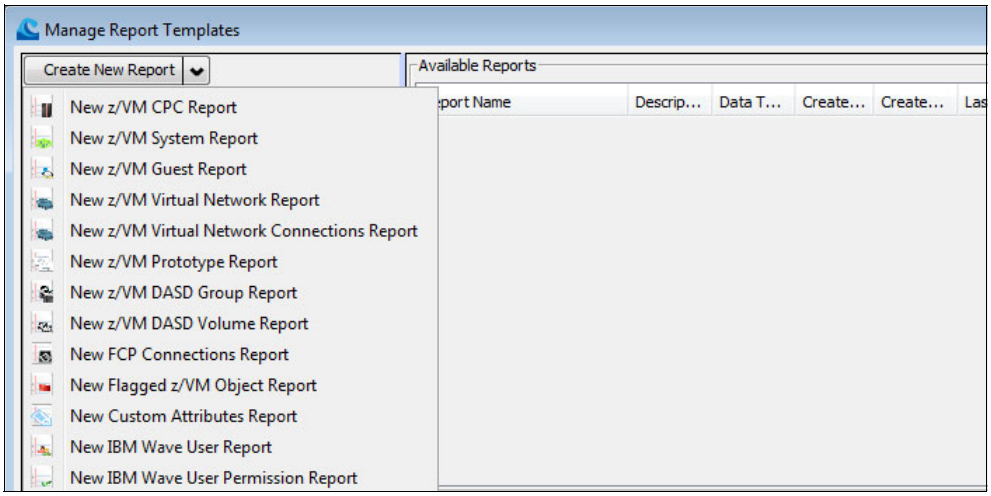


*Figure 5-72 Create z/VM users report*

Clicking the report name opens a window in which you can customize the content of the report. In our example, we select the z/VM User Report type, as shown in Figure 5-73.

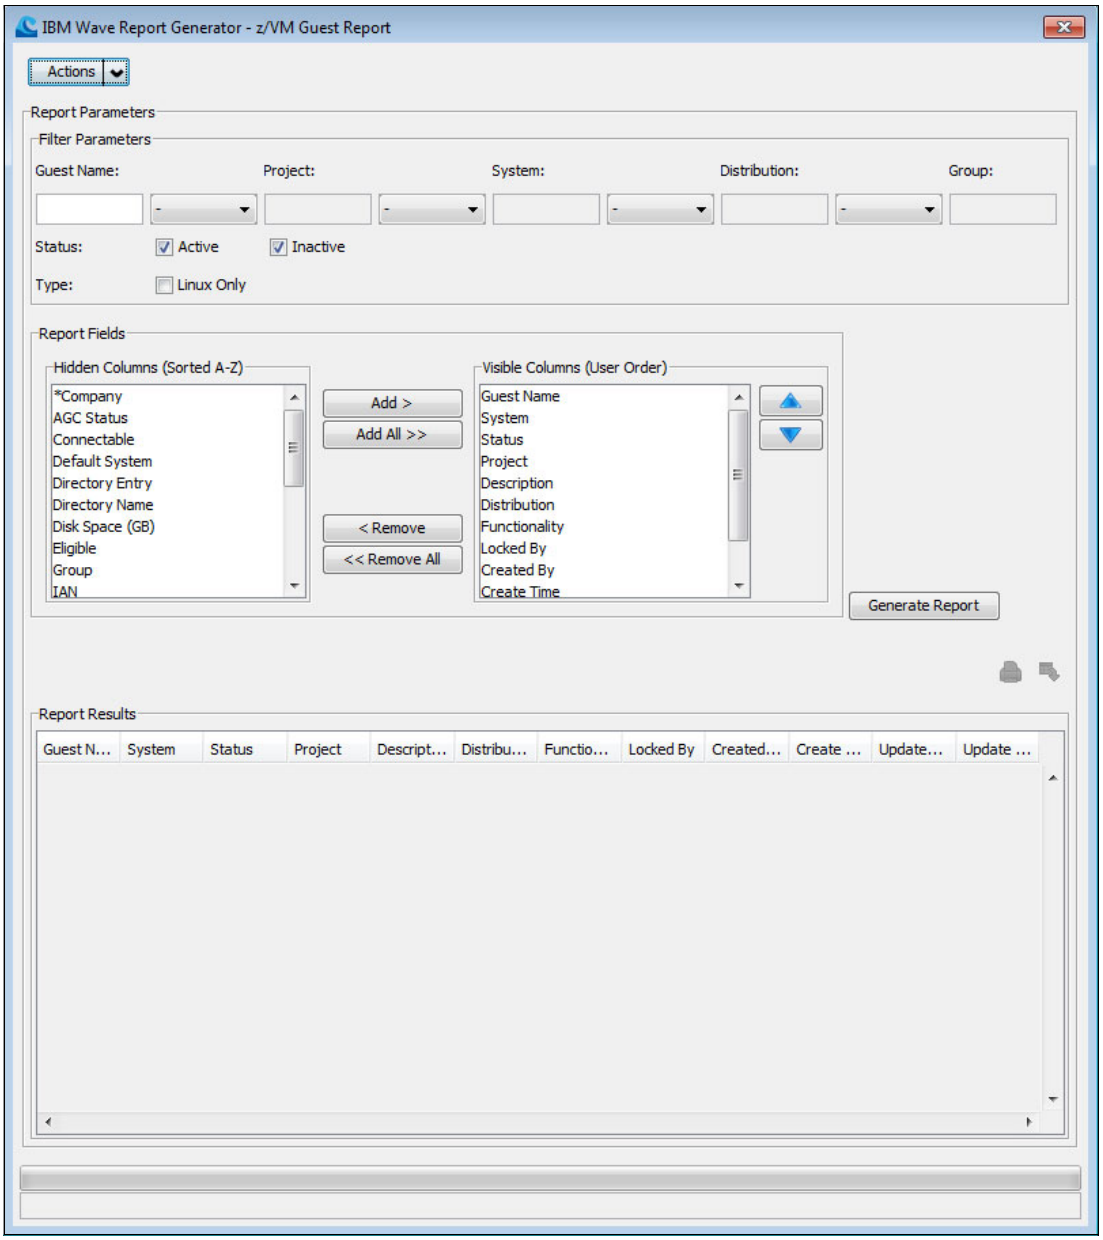


*Figure 5-73 z/VM User Report window*

3. By using this window, you can restrict the report output to particular systems, projects, distributions, or groups. You also can include Active and Inactive VMs, or Linux only VMs. You also can customize the column data by selecting or clearing various items in the Table Fields portion of the window. To see all Active and Inactive VMs across all systems, include an asterisk in the User Name field and click **Generate Report**, as shown in Figure 5-74.

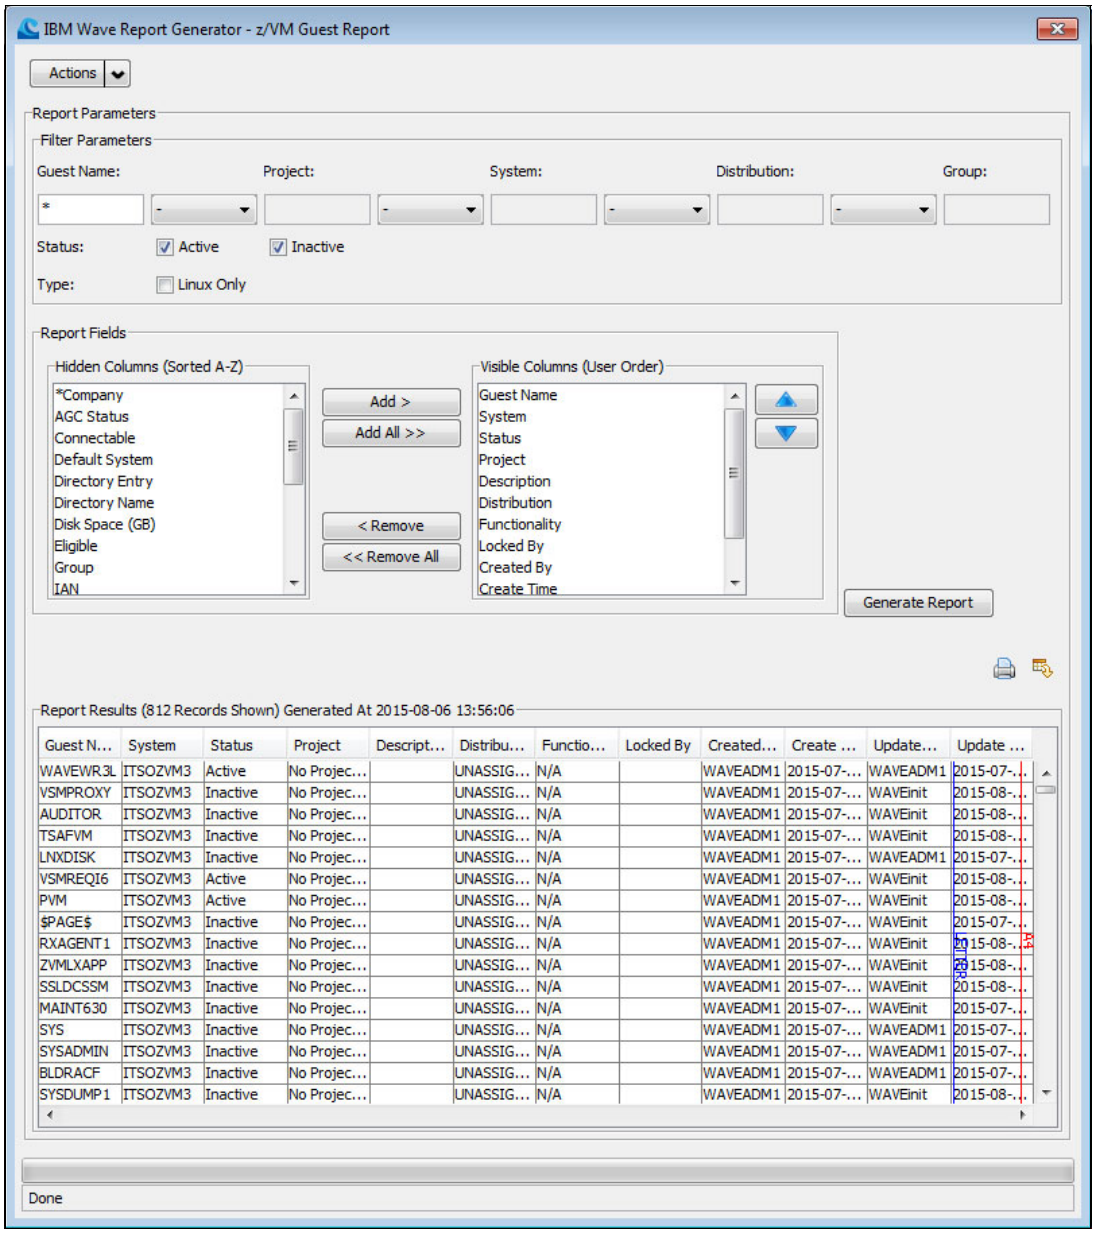


*Figure 5-74 z/VM user report output*

The report output is at the bottom of the window with two lines that show how much of the information is printed with standard letter size paper or A4 paper in landscape mode before the output is folded. You now might want to clear some columns to make the output better fit a printed report or you can leave the report content as is.

4. Clicking **Actions** in the upper left corner displays choices to print the report output or save the report criteria for easy generation at another time. To save the report, click **Actions** → **Save As**. A window opens in which you can enter the name and a description of the report, as shown in Figure 5-75.

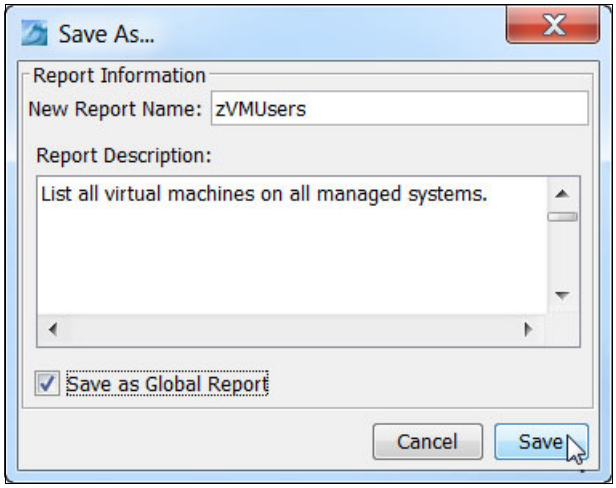


*Figure 5-75 Save z/VM users report*

You must have site-level authority to save the report as a Global Report. By saving the report as a Global Report, it is available for all IBM Wave users to generate. When confirmation is received that the report was saved, it appears in the Report Manager window, as shown in Figure 5-76.

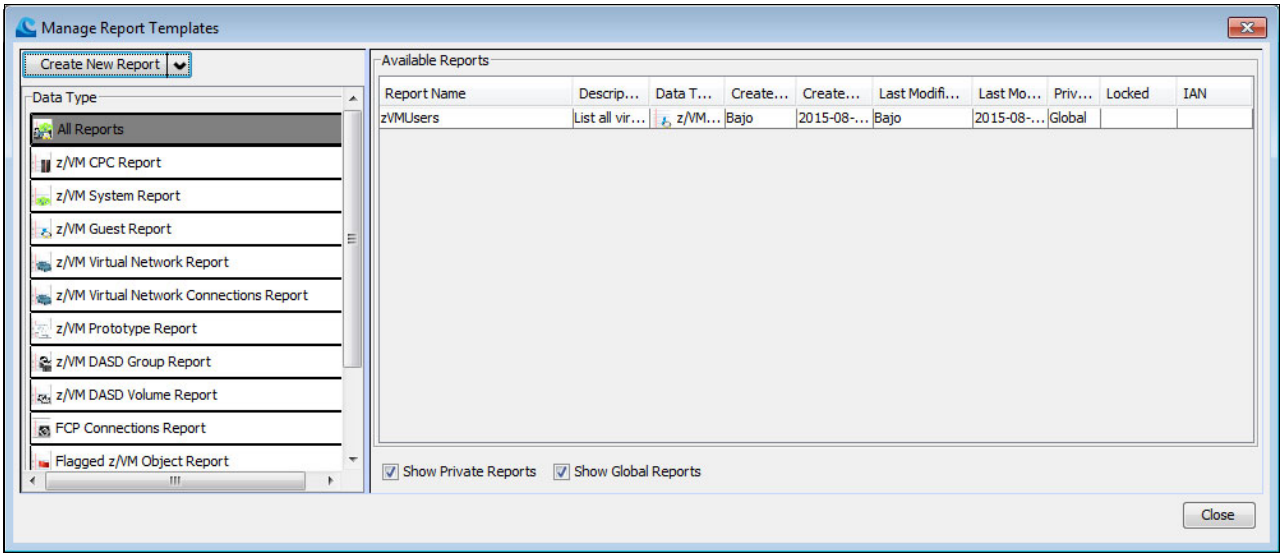


*Figure 5-76 New saved report*

5. To generate the saved report, double-click the line for the report. A window opens that shows all of the selections that were made before the report was saved. When you click **Generate**, the report content is displayed in the center of the window. Figure 5-77 shows the generated report.

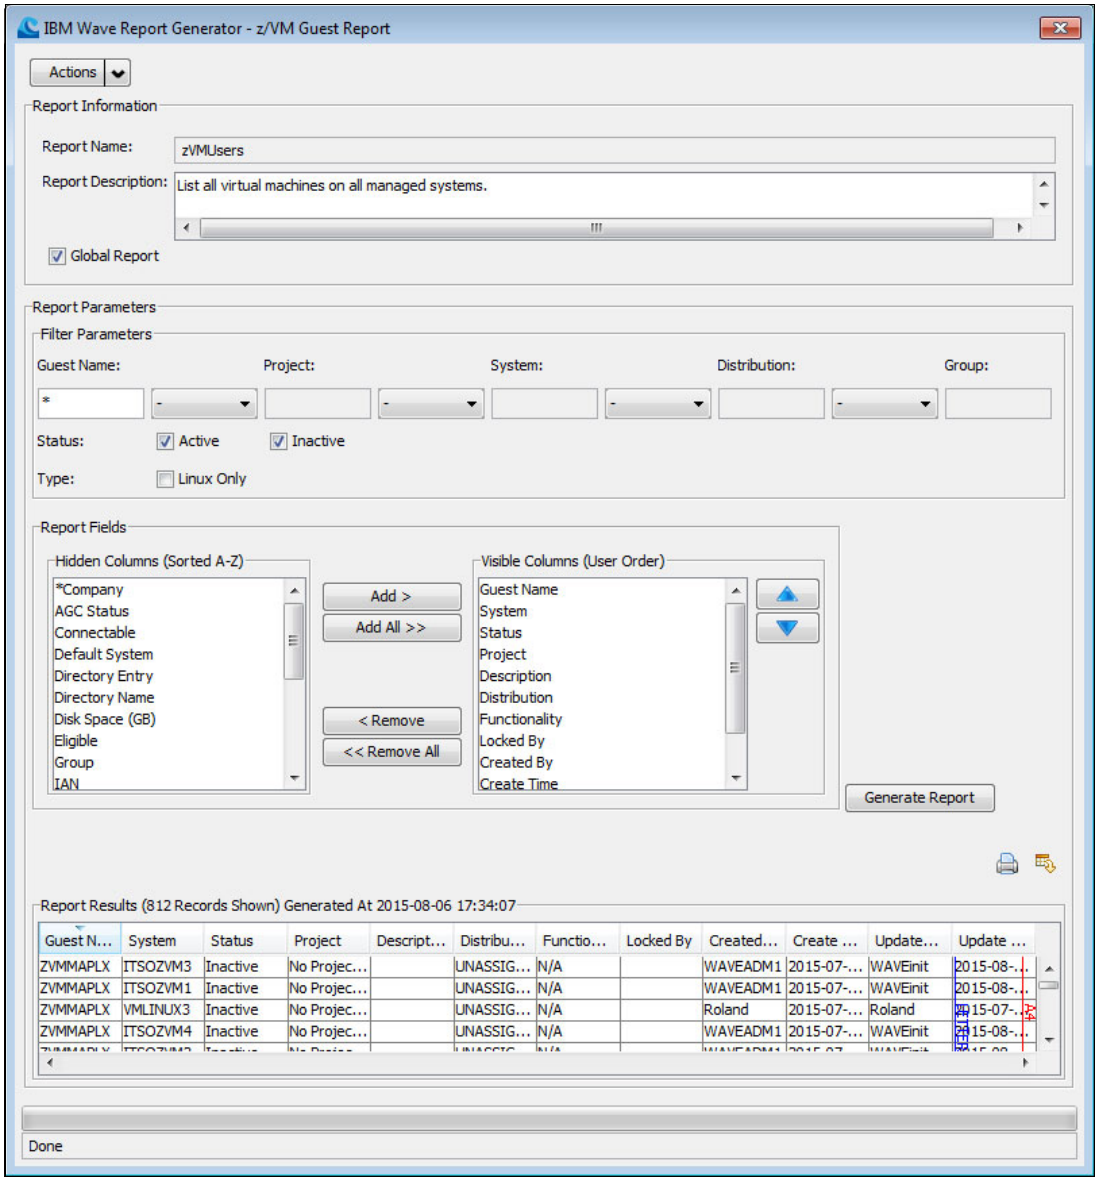


*Figure 5-77 Generate a saved z/VM users report*

6. To print the report, click **Actions** → **Print**. A print preview window opens, as shown in Figure 5-78 on page 135.

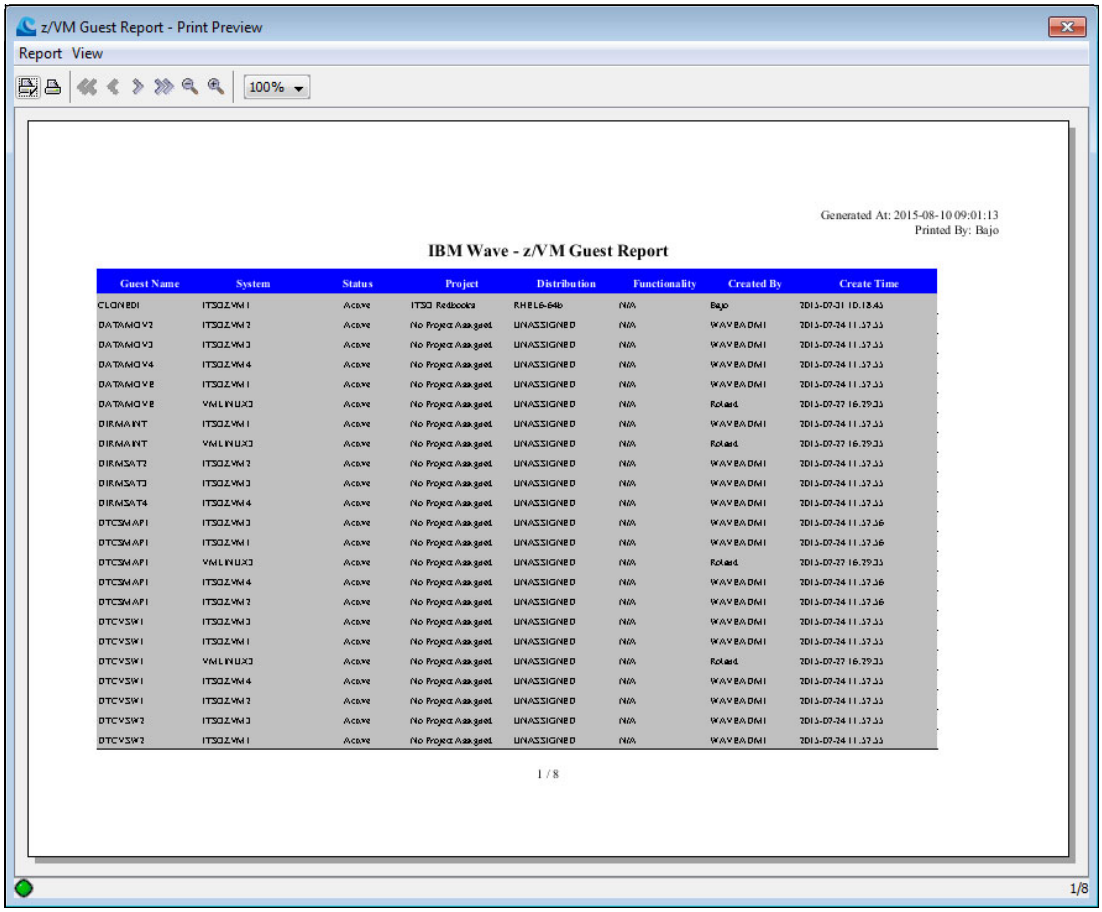


*Figure 5-78 Print Preview window*

7. Click **Report** in the Print Preview window to send the previewed report to a printer.
8. Instead of printing the report, you can export it as a CSV file by clicking **Actions** → **Export to CSV**, as shown in Figure 5-79.

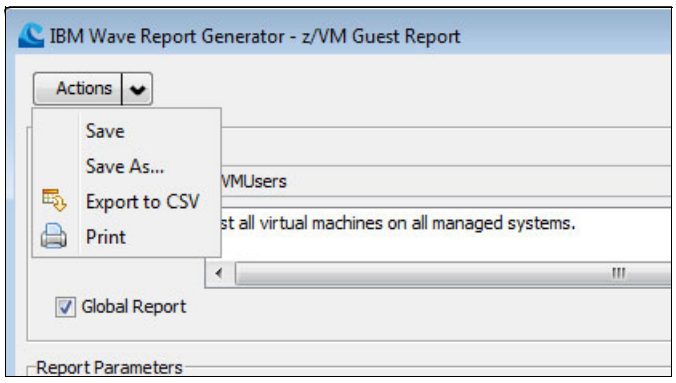


*Figure 5-79 Export to CSV file option*

## 5.9 User management

In a traditional z/VM environment, z/VM systems programmers often carry the burden of administering both hypervisor and guests because traditional 3270-based CMS interfaces often prove too foreign and complex for Linux systems administrators who are working in large multi-hypervisor environments. Although other hypervisors offer methods of easily configuring guests through GUI-based management consoles, z/VM lacked this sort of systems management interface.

This fact led Linux administrators and middleware owners to rely on z/VM systems programmers for tasks that often are handled though self-service systems management GUIs on other hypervisors. Need more processors, memory, or disk? Ask the systems programmer. Need to reset the z/VM user password? Email the systems programmer. For environments with strong z/VM systems programming staff, this was not an issue. However, for new mainframe customers who are consolidating Linux workloads onto z Systems systems, the complexity of traditional z/VM management tools was a stumbling block.

The introduction of IBM Wave changes this dynamic by providing a graphical management tool that provides functions that are suitable for different classes of users. Middleware owners can activate, deactivate, and reset the guests that are running the software for which they are responsible. Linux administrators can provision guests from golden images, resize LVMs, and access the guest from an SSH session into guests to install packages. z/VM systems programmers can configure vSwitches and DASD storage pools, write advanced REXX EXECs, start 3270 sessions for traditional CMS-based tools, and manage the catalog of golden images.

Because IBM Wave offers features that appeal to many classes of users, the functions to create different classes of users and determine permissions and scope are especially critical. After all, do you want to receive an email from your database administrators informing you that they accidentally ran a function that is called “re-IPL this z/VM system” and then all connectivity to the Linux guest was lost?

To prevent these sorts of issues, it is critical to consider the following questions:

- What sorts of different roles (IT or business unit personnel) use Linux guests that are running on z/VM regularly?
- What sort of functions does your company’s Linux systems administrators regularly use in a self-service fashion with other hypervisors?
- How skilled is the IT department at z/VM systems programming?
- What sorts of methods do you have in place to segment guests between projects or business units? Is segmentation provided manually by human administration or by technical means?

The answers to these questions help you determine the classes of users that might exist in your IBM Wave environment, the subset of functions that each class might possess, and the preferred means of segmenting access in multi-tenant environments.

### 5.9.1 IBM Wave User Manager tool overview

IBM Wave administrators manage the creation and modification of users (deletion of users is impossible) and the granting and revocation of scope and privileges through the IBM Wave User Manager. This tool is available from the toolbar by clicking the icon that resembles a pencil that is hovering over two people, as shown in Figure 5-80 on page 137.

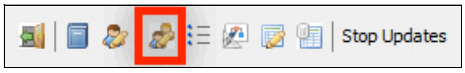


*Figure 5-80 User Manager icon in the IBM Wave toolbar*

It is also available from the menu bar by clicking **Administrative** → **Manage IBM Wave Users**, as shown in Figure 5-81.

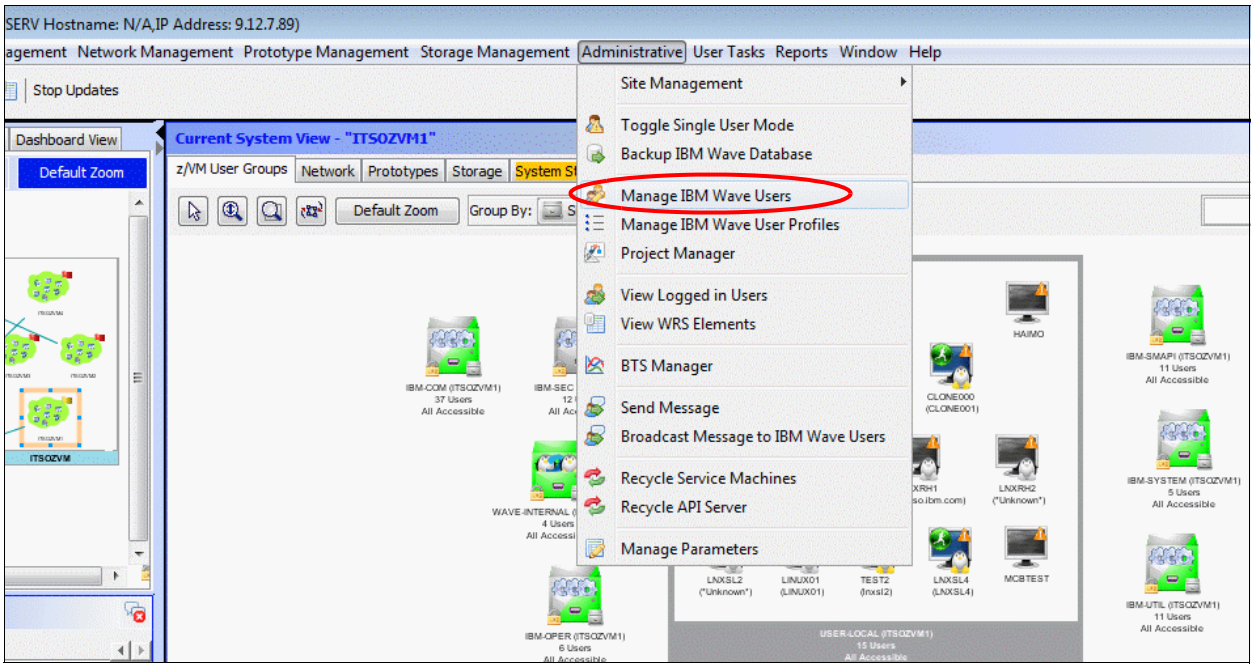


*Figure 5-81 Click Manage IBM Wave Users task*

Either of these methods opens the IBM Wave User Manager. This tool offers the following key functions:

- List all IBM Wave users through a sortable and filterable table on the main IBM Wave User Manager window.
- Create IBM Wave users from scratch by using the Create New IBM Wave User option that is available from the Add New IBM Wave User option.
- Create IBM Wave users by cloning the password, details, user type, scope, and permissions of an existing user through the Clone IBM User tool that is available from the menu through the Clone the IBM Wave User option.
- Grant or revoke roles, permissions, and scope to or from a single IBM Wave user through the Update IBM Wave User window that is available from the menu through the Update IBM Wave User option.
- Grant or revoke roles, permissions, and scope by making a user's roles, permissions, and scope identical to another user through the Copy Selected Scopes and Permissions tool. This tool is available from the menu of the destination user through the Copy scope and permissions to selected user option.
- Audit and revoke permissions across all IBM Wave users through the IBM Wave User Permissions Cleaner tool that is available from the Review Scope and Permissions option.

- Modify a single user's password or security question or force them to change their password upon their next login by using the Update IBM Wave User window that is available from the menu that is opened by right-clicking a user in the table of existing users.
- Suspend or resume a single user through the Update IBM Wave User window that is available from the menu through the Update IBM Wave User option.

These features provide two workflows for administering users. The first workflow creates IBM Wave users one at a time and manually sets the type, scope, and permissions of each user. Although this approach is more fine-tuned, it is time-consuming and does not scale well to larger IBM Wave environments. The second workflow creates multiple IBM Wave users in batch operations by cloning known users or adjusting user scope and permissions to many users.

With these options, it is a preferred practice to create *golden master* users for each particular class of user (z/VM systems Programmers, Linux Systems Administrators, IBM DB2® DBAs, and so on) and use these golden master users to clone individual users of each class.

**Note:** By using IBM Wave, you also can use LDAP for single-sign on and Active Directory groups to predefine the permissions and scope that are associated with classes of users. For more information, see the *IBM Wave Administration Guide*.

To demonstrate how to use IBM Wave to manage users and scope, this section describes the example of creating multiple user IDs for an IBM WebSphere® Application Server middleware team (by using a project). These users have permissions that are limited to the following basic types of IBM Wave tasks with which Linux administrators are concerned:

- Activating, deactivating, and resetting the Linux guests that are running WebSphere Application Server
- Cloning from a Linux golden master that is preinstalled with WebSphere Application Server
- Creating and expanding Logical Volumes (LVMs)
- Accessing Linux guests through SSH

Additionally, this WebSphere Application Server middleware team has a scope that is limited to viewing and interacting with guest VMs in their departmental project.

## 5.9.2 Creating a golden master user

To ease the task of user management, we now create a golden master user that possesses the roles, permissions, and scope that is common to all administrators on the WebSphere Application Server middleware team. No user logs in to IBM Wave with this ID; instead, we clone all instances of WebSphere Application Server administrator users from this golden master.

Complete the following steps:

1. To create this WebSphere Application Server Administrator user, click **Add New IBM WAVE User** from the main IBM Wave User Manager window, as shown in Figure 5-82 on page 139.

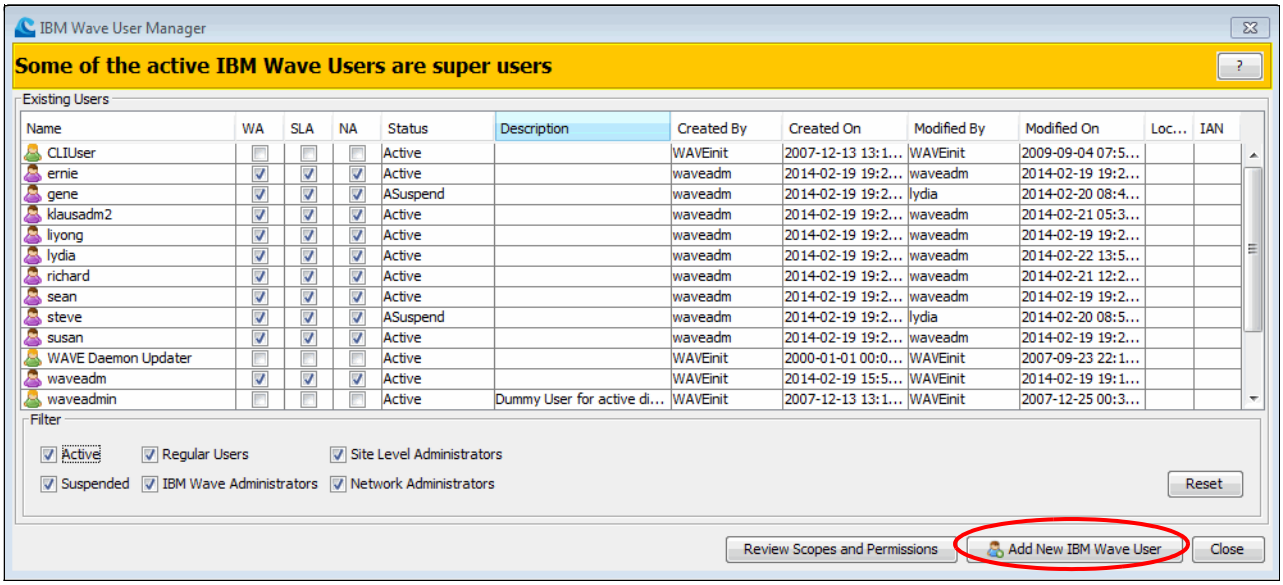


*Figure 5-82 Click Add New IBM Wave User*

The Create New IBM WAVE User window in the General Details tab opens, as shown in Figure 5-83.

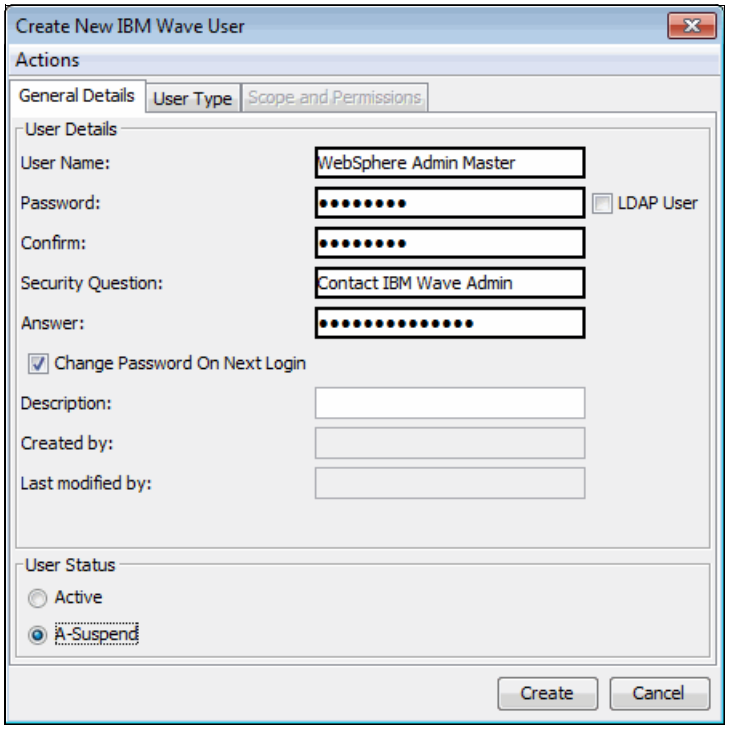


*Figure 5-83 General Details tab of Create New IBM Wave User wizard*

2. To signify that this is a golden master user, give this user the user name WebSphere Admin Master.
3. Assign this user the password “password.” Although this is not a secure password, assigning a simple and memorable password to golden master users is critical because all users that are cloned from this golden master inherit the same password.

However, by keeping this user in the status “A-Suspend” except for the brief periods when you use this guest for cloning actions, this guest can remain unusable. Additionally, by selecting **Change Password On Next Login**, users can be forced to change from the initial password string to something more secure.

The Security Question and Answer values are not completed as normal because you expect to clone numerous users from this golden master user. If you cloned multiple users with the same security question, a malicious user might access other user’s logins by using the security question and answer to reset the password. To prevent this issue, set the Security Question to “Contact IBM Wave Admin” and the password to a secret value that is known only by the IBM Wave administration team. The “Contact IBM Wave Admin” string informs users to contact an IBM Wave administrator for help resetting their IBM Wave password. If a user updates their preferences later, they can change their security question and answer to something personalized to use the Security Question feature as intended.

The Description field is a 255-character area for entering user-readable information to make it easier to organize and identify IBM Wave users. Considering that cloned users have the same description as source users, this field is a good way to associate classes of users. Using this field, you can view all instances of a class of user together by sorting the Description field in the IBM Wave User Manager window. In our example, we enter the string “WebSphere Admin” to describe this class of user.

4. Click the User Type tab to select the type of user. Within this context, the Administrator user type refers to IBM Wave administrators. Linux or middleware administrators are considered Regular Users. The Network Admin and Site Level Admin options are associated with particular permissions.

   In the case of our WebSphere Admin Master user, we ensure that **Regular User** is selected, as shown in Figure 5-84.

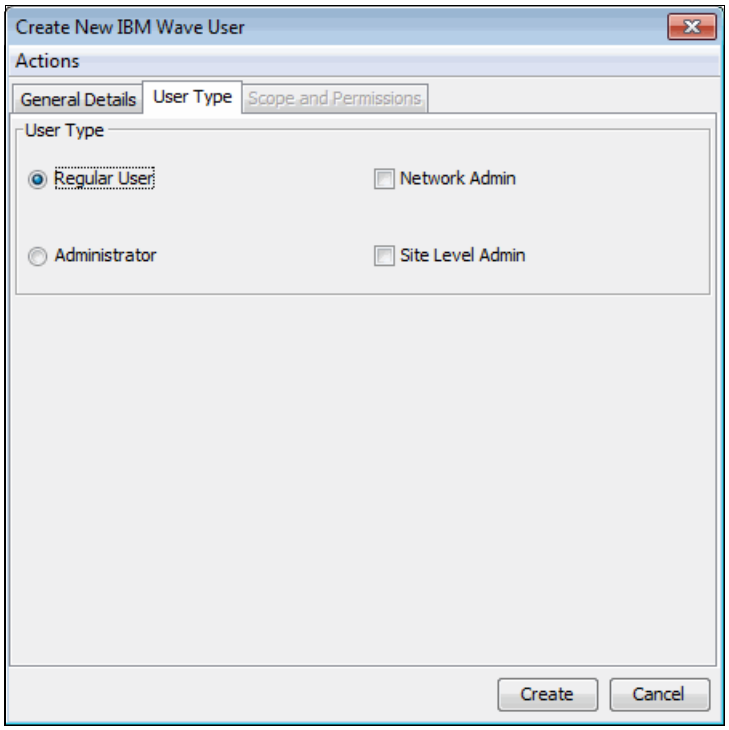


*Figure 5-84 User Type tab of Create New IBM Wave User wizard*

5. The Scope and Permissions tab is disabled. To set Scope and Permissions for the user, create the user, and then access this tab by clicking **Update IBM Wave User**. So, click **Create** to instruct IBM Wave to use the values that were entered to create a user. A confirmation window opens when this process is complete, as shown in Figure 5-85.

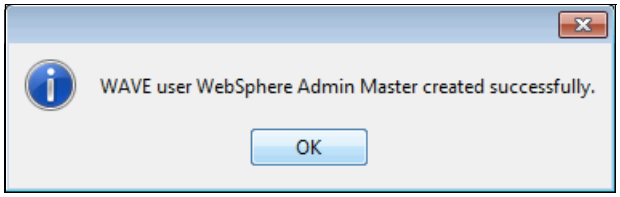


*Figure 5-85 Create an IBM Wave User Confirmation window*

6. After you return to the main IBM Wave User Manager window, you see that the WebSphere Admin Master user is listed within the table of existing users. Right-click this user and select **Update IBM Wave User** from the menu, as shown in Figure 5-86.

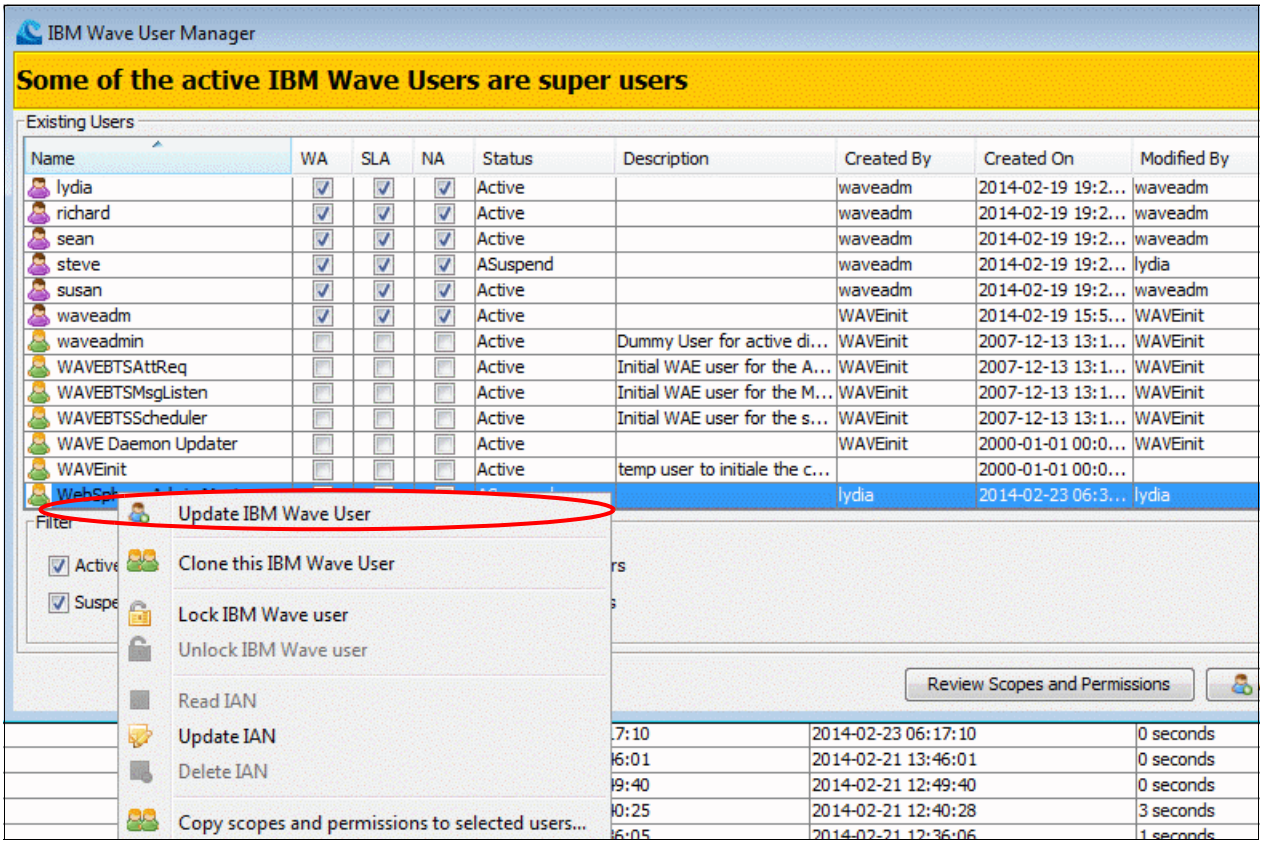


*Figure 5-86 Select Update IBM Wave User from the WebSphere Admin Master user's menu*

The Update IBM Wave User window opens, as shown in Figure 5-87. This window is nearly identical to the Create IBM Wave window that you used to create the WebSphere Admin Master user. However, the User Name field is disabled (IBM Wave user names never can be changed after initial creation) and the Scope and Permissions tab is now accessible.

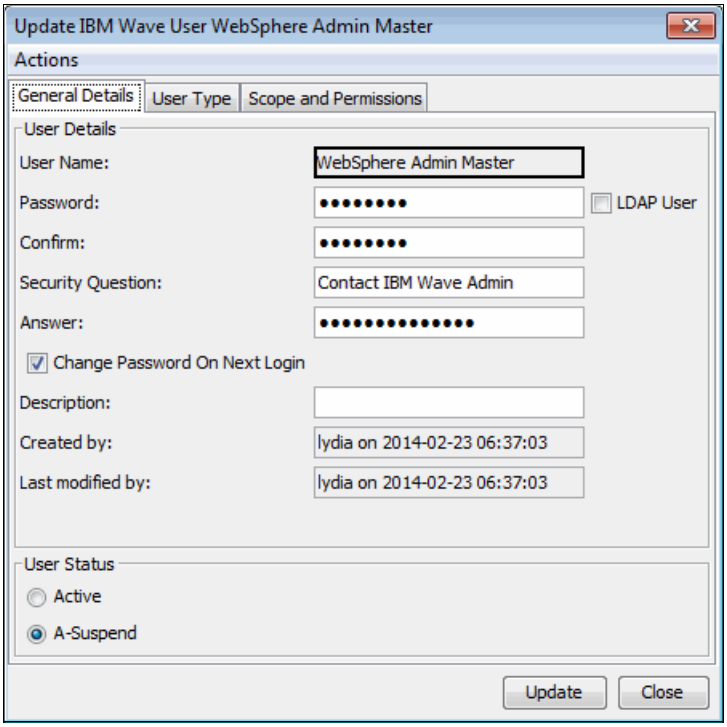


*Figure 5-87 Update IBM Wave user window*

7. Click the **Scope and Permissions** tab and then click **Add Permission**. Grant the user limited permissions to visualize z/VM systems and display the status of the system (Active, Inactive, and so on). All other functions that are related to z/VM systems are restricted. The completed window should look Figure 5-88.

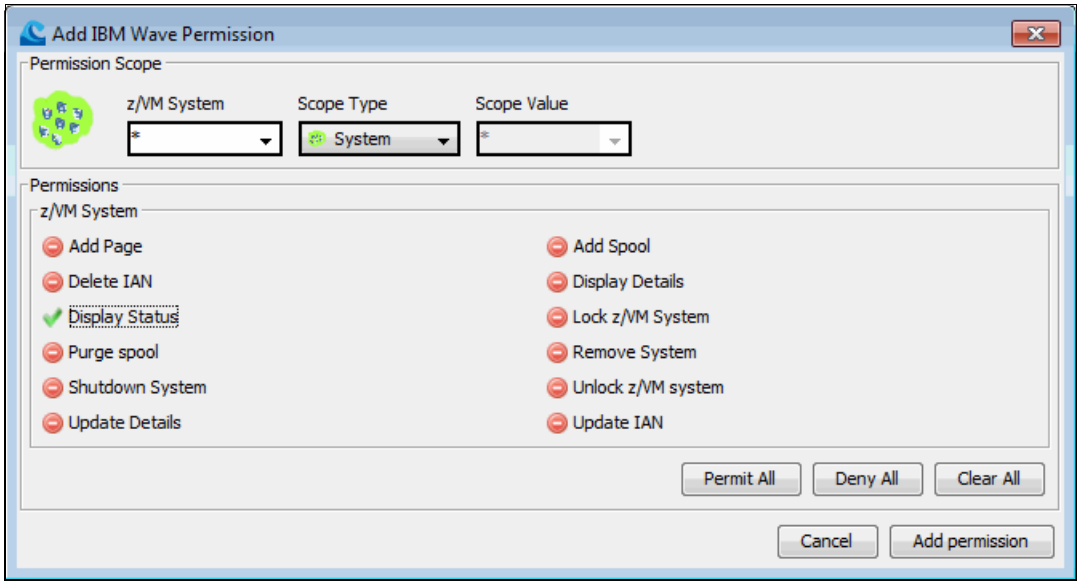


*Figure 5-88 Add permissions for z/VM systems*

8. Click **Add Permission** to return to the Scope and Permissions tab of the Update IBM Wave User window. The permission and scope that you set is listed in the table. Click **Add Permission** to add another permission.
9. Grant the user permission to perform various actions within the scope of the project that is called “WebSphere Middleware Team”, as shown in Figure 5-89. This project contains the Associated Prototypes and Linux guests that are used by the WebSphere middleware team. Because this permission is limited to a particular project, this user cannot see any guests outside of this project. After the information is complete, click **Add Permission**.

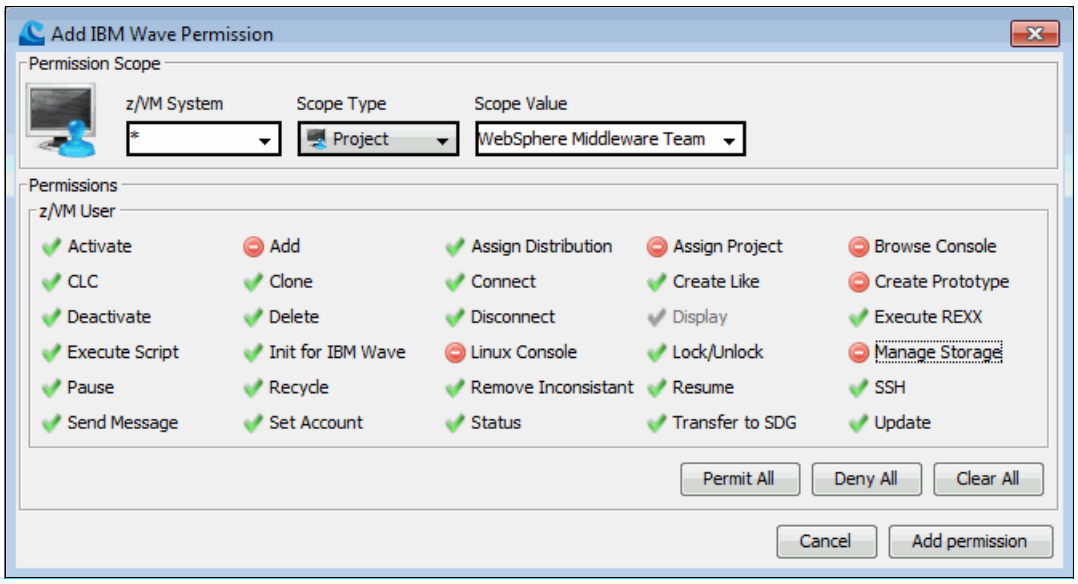


*Figure 5-89 Add permissions for IBM Wave projects*

10. After you return to the Scope and Permissions tab of the Update IBM Wave User window, click **Add Permission**. Add permissions for DASD volumes, as shown in Figure 5-90. Specifically, give the user the permission to add disks from a DASD storage group to a Linux system. If this permission is not granted, it is impossible to create Linux guests or use DASD to expand Linux logical volumes.

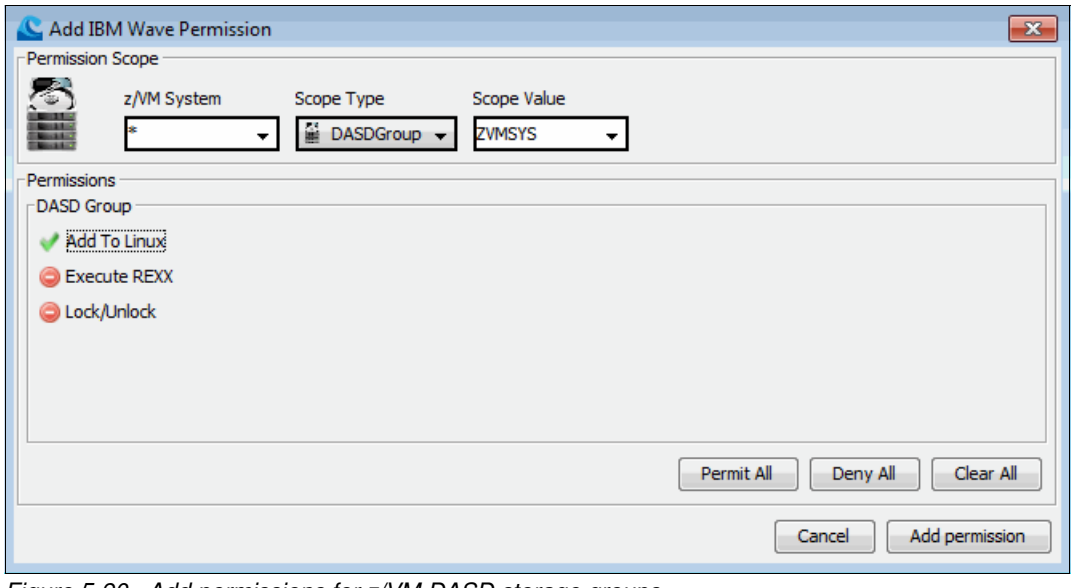


*Figure 5-90 Add permissions for z/VM DASD storage groups*

11. After you complete the information, click **Add Permission**. You return to the Scope and Permissions tab of the Update IBM Wave User window. Close the Update IBM Wave User window to return to the main IBM Wave User Manager window.

You created a golden master user for WebSphere administrators.

## 5.9.3 Cloning individual users from a golden master user

Now that you created a golden master user, you can clone new instances of this type of user.

Recall that you put the WebSphere Admin Master user in a suspended state to prevent users from using it to log in to IBM Wave because it had the string `password` set as the user's password. If you clone users from this golden master user, the clones also are in a suspended state and must be activated individually. This process can be inefficient for large batches of clones. It is more efficient to update the golden master users' status to Active before cloning is done. Then, when cloning is complete, the golden master user can be set back to a suspended state.

Complete the following steps:

1. Right-click the WebSphere Admin Master user and select **Update IBM Wave User**, as shown in Figure 5-91.

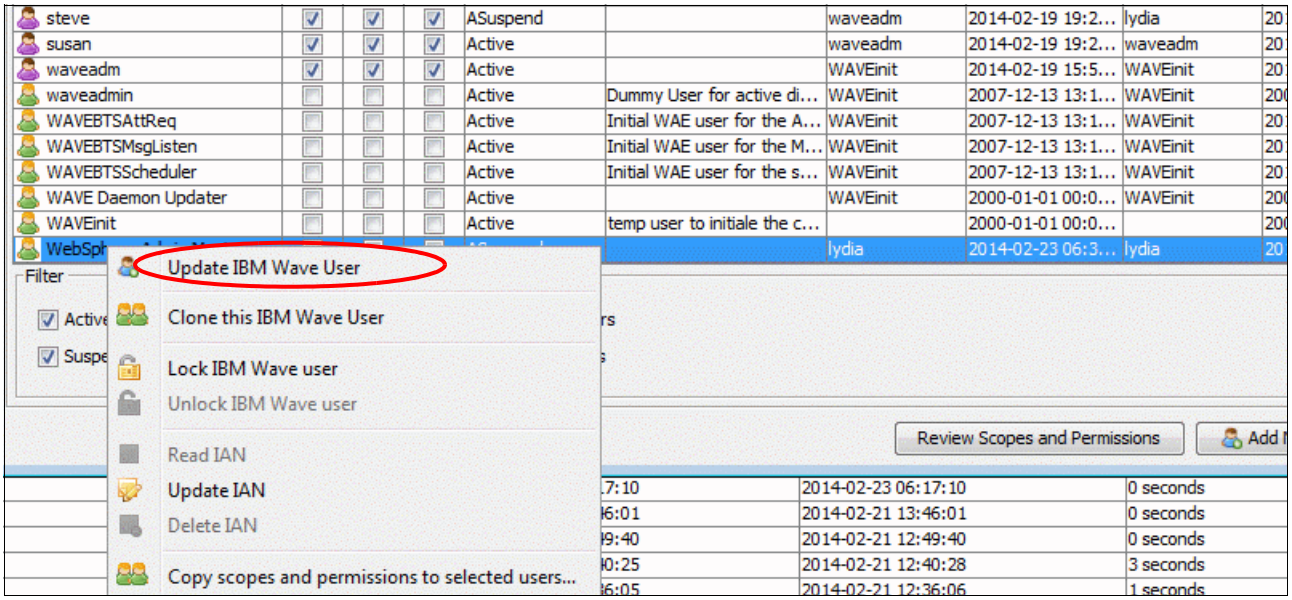


*Figure 5-91 Select Update IBM Wave User from WebSphere Admin Master's menu*

2. In the Update IBM Wave User window, click the **Active** option and then click **Update**, as shown in Figure 5-92 on page 145.

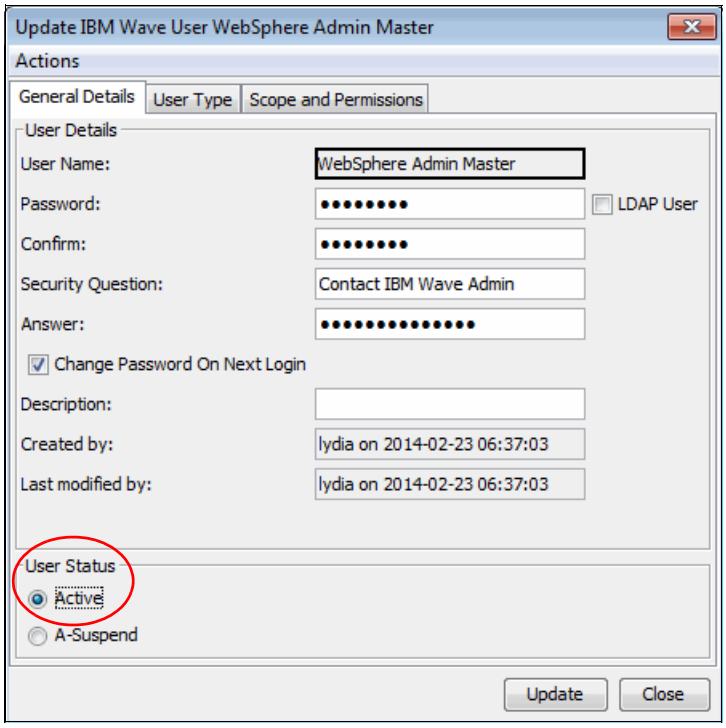


*Figure 5-92 Set WebSphere Admin Master to Active user status*

3. After a confirmation message is shown, you return to the main IBM Wave User Manager window. Right-click the **WebSphere Admin Master** user and select **Clone this IBM Wave User**, as shown in Figure 5-93.

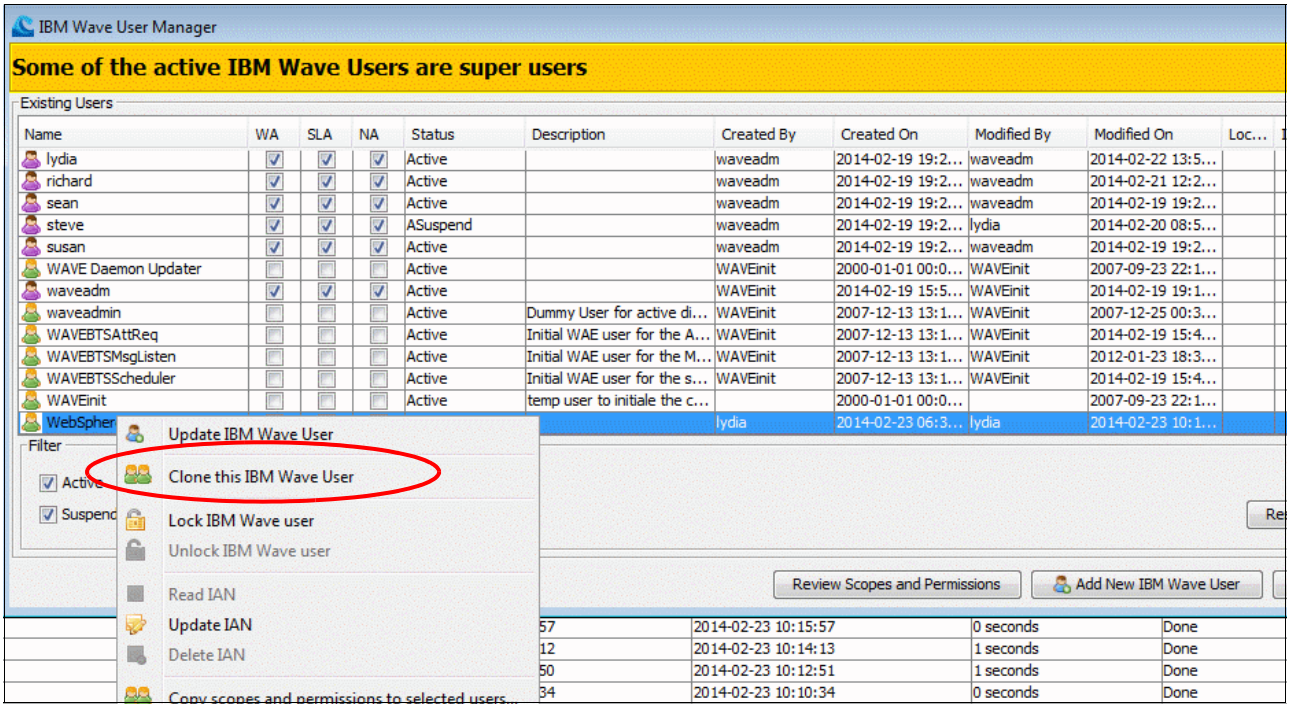


*Figure 5-93 Select Clone this IBM Wave User from WebSphere Admin User's menu*

4. Because you are cloning users for three new WebSphere administrators, enter 3 in the Number of users field. Because you do not want to prefix each of these user names with the same string, leave the Base Name field blank. Your window should resemble what is shown in Figure 5-94. Click **Update**.

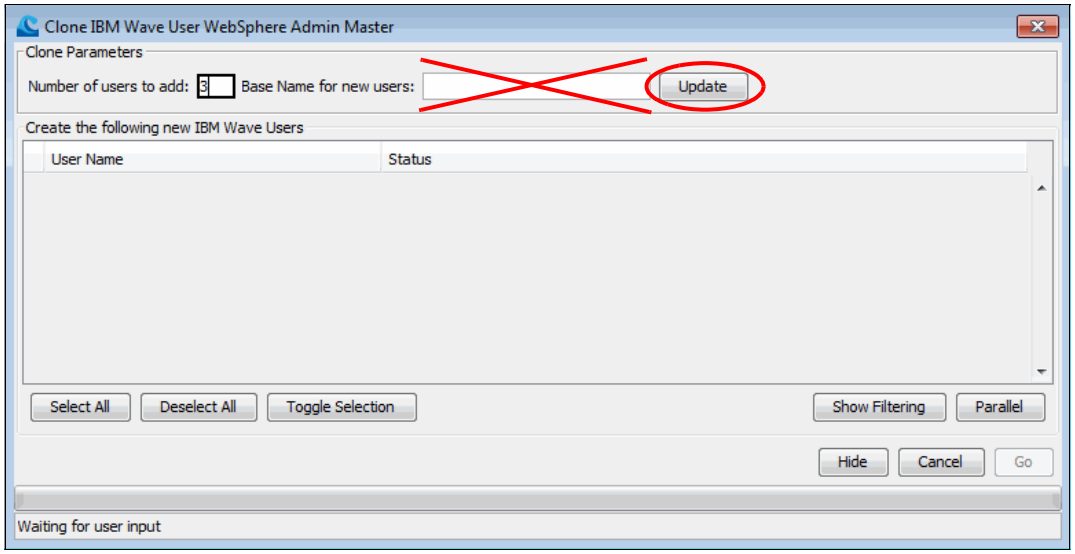


*Figure 5-94 Set the number of User Instances to clone from WebSphere Admin Master*

5. This task populates the table that is in the middle of the window with three rows. Each of the rows has the value <Insert user name here> in the User Name column, as shown in Figure 5-95, which signifies that you must enter a unique user name in each of these fields. Double-click each field, enter a user name, and then press Enter.

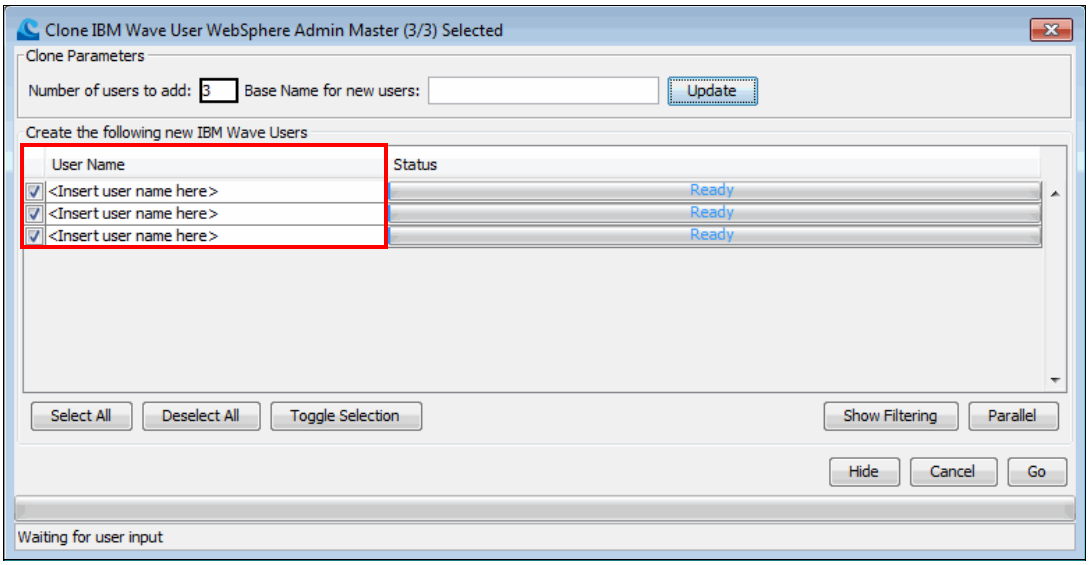


*Figure 5-95 Enter user names for each cloned IBM Wave user*

6. Click **Go** after you enter a name for each user, as shown in Figure 5-96 on page 147.

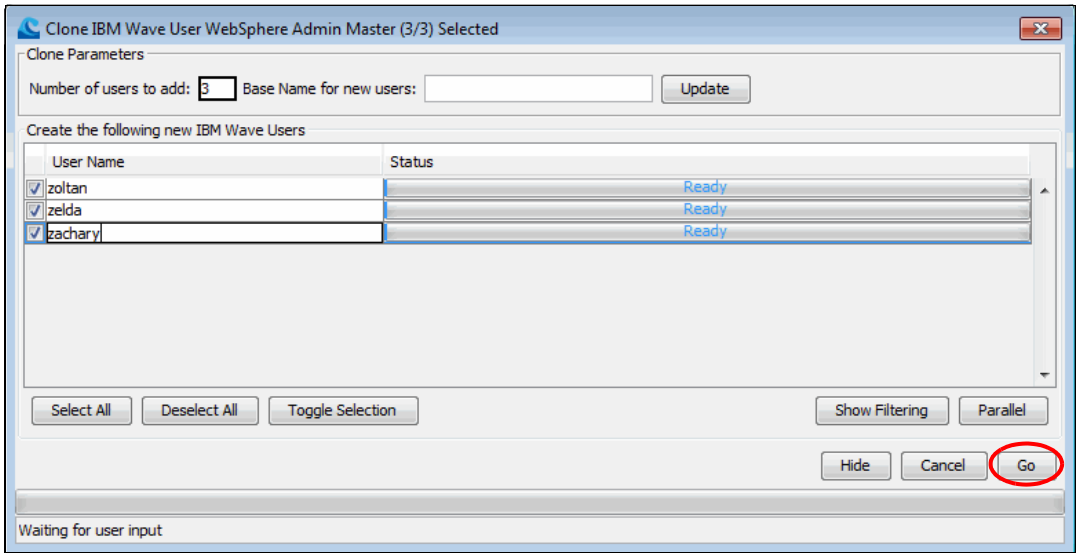


*Figure 5-96 Click Go to begin the IBM Wave user cloning process*

7. The status cell of each row shows the progress of the cloning process. Close the Clone IBM Wave User window when all of the rows show the Done status.
8. Your three new WebSphere admin IDs are created and are ready for their initial login. However, you must put the WebSphere Admin Master user back into a suspended state to prevent security loopholes. Right-click the WebSphere Admin Master and select **Update IBM Wave User**, as shown in Figure 5-97.

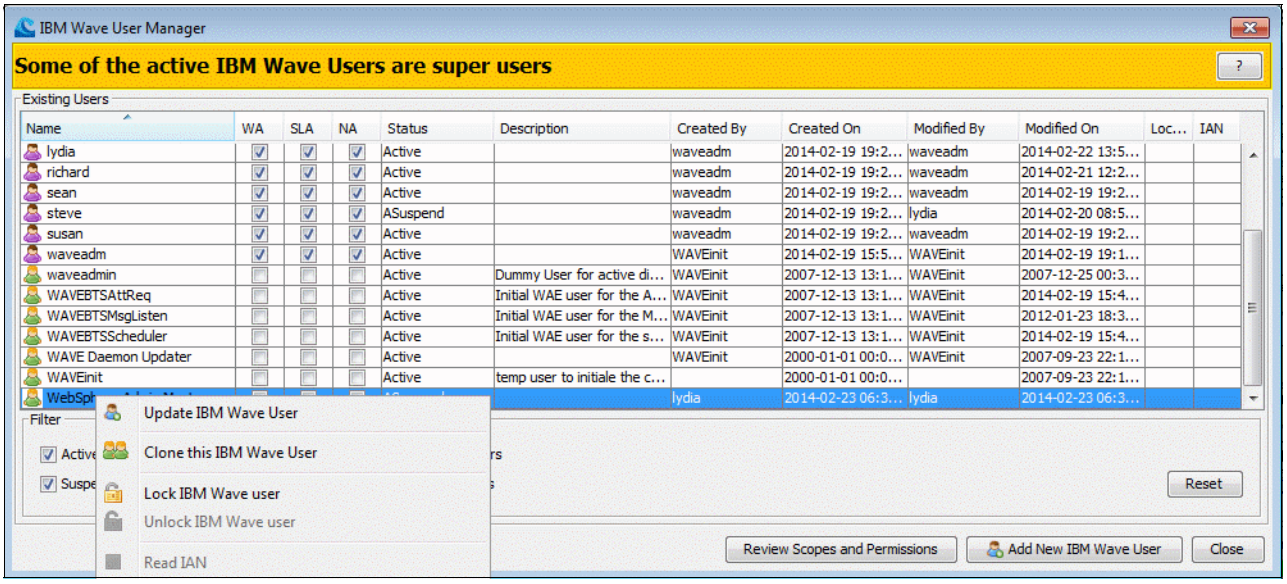


*Figure 5-97 Select the Update IBM Wave User action from WebSphere Admin Master's menu*

9. Click the **A-Suspend** option and then click **Update**, as shown in Figure 5-98.

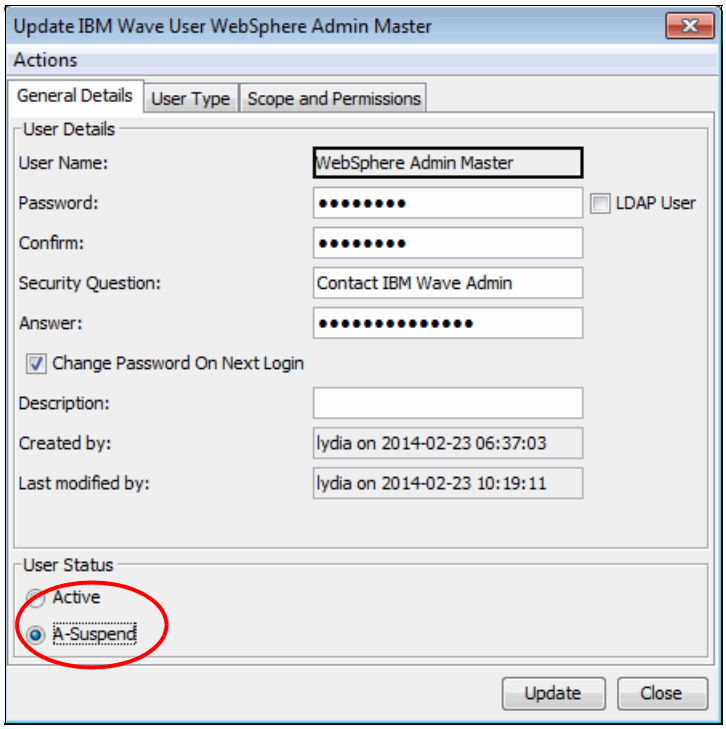


*Figure 5-98 Set WebSphere Admin Master to Suspended User Status*

You created a golden master user that is used explicitly for cloning. You also used this user to clone rapidly three user IDs.

## Revoking permissions from an IBM Wave user

**Note:** To change the permissions of an IBM Wave user, the user ID must not be logged in to an active session.

When an employee changes roles or leaves your company, it is important to revoke permissions from the employee's IBM Wave user. This task is performed from the IBM Wave User Manager window by completing the following steps:

1. Right-click a user and select **Update IBM Wave User**, as shown in Figure 5-99 on page 149.

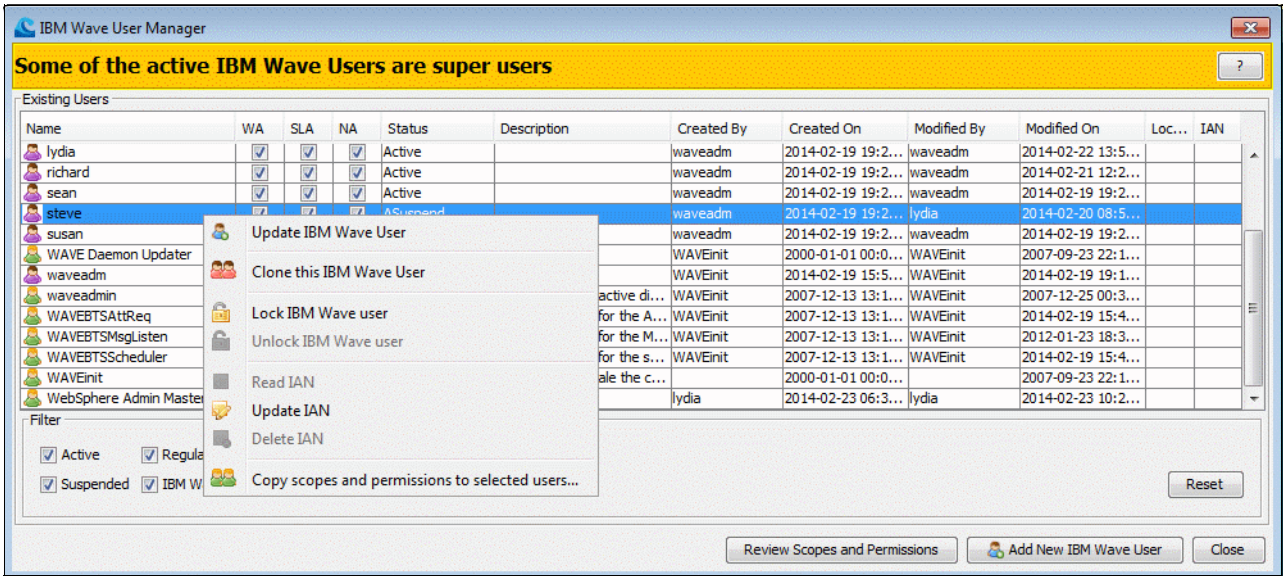


*Figure 5-99 Update IBM Wave User*

2. In the Update IBM Wave User window, click the **Scope and Permissions** tab. A table of the user's current permissions opens. Highlight a permission that you want to delete and click **Delete Permission**, as shown in Figure 5-100.

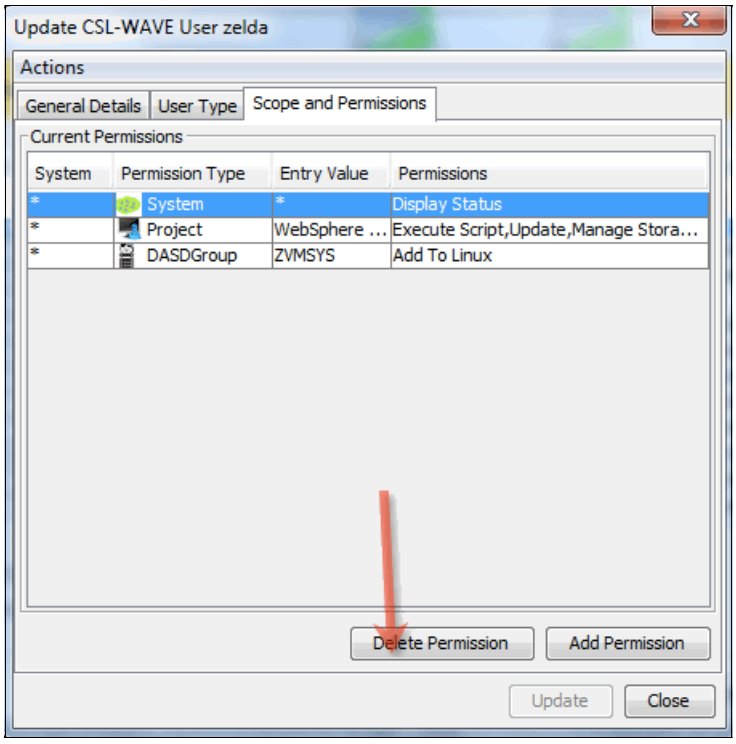


*Figure 5-100 Select and delete IBM Wave user permissions*

3. This process can be repeated for all of the permissions that are available to the user. After this process is complete, click **Close** to return to the main IBM Wave User Manager window.

## Changing the class of IBM Wave user by cloning another golden master

Section 5.9.2 “Creating a golden master user” on page 138 demonstrates how to create a golden master template user that is preconfigured with the scope and permission that is needed by a particular class of users. You named this user WebSphere Admin Master to signify its status as a golden master template. You then cloned numerous instances of users with the class of WebSphere admin by cloning from WebSphere Admin Master.

However, perhaps one of these users changes roles in the company and moves from the role of WebSphere Admin to Database Admin.

As you used golden master templates to create instances of a particular class of user, you can use these templates to modify a particular user ID from one class of user to another by selectively cloning the scopes and permission of another user template.

**Note:** Assuming that a user is changing jobs, all existing permissions must be deleted, as described in “Revoking permissions from an IBM Wave user” on page 148, before the scopes and permissions of the new user template are cloned because the clone action appends other permissions rather than replacing existing permissions. A user can then have the permissions of multiple classes of users.

To grant an existing IBM Wave user with all of the scope and permissions of another class of users, complete the following steps:

1. Right-click a user and select **Copy scopes and permission to selected users**, as shown in Figure 5-101.

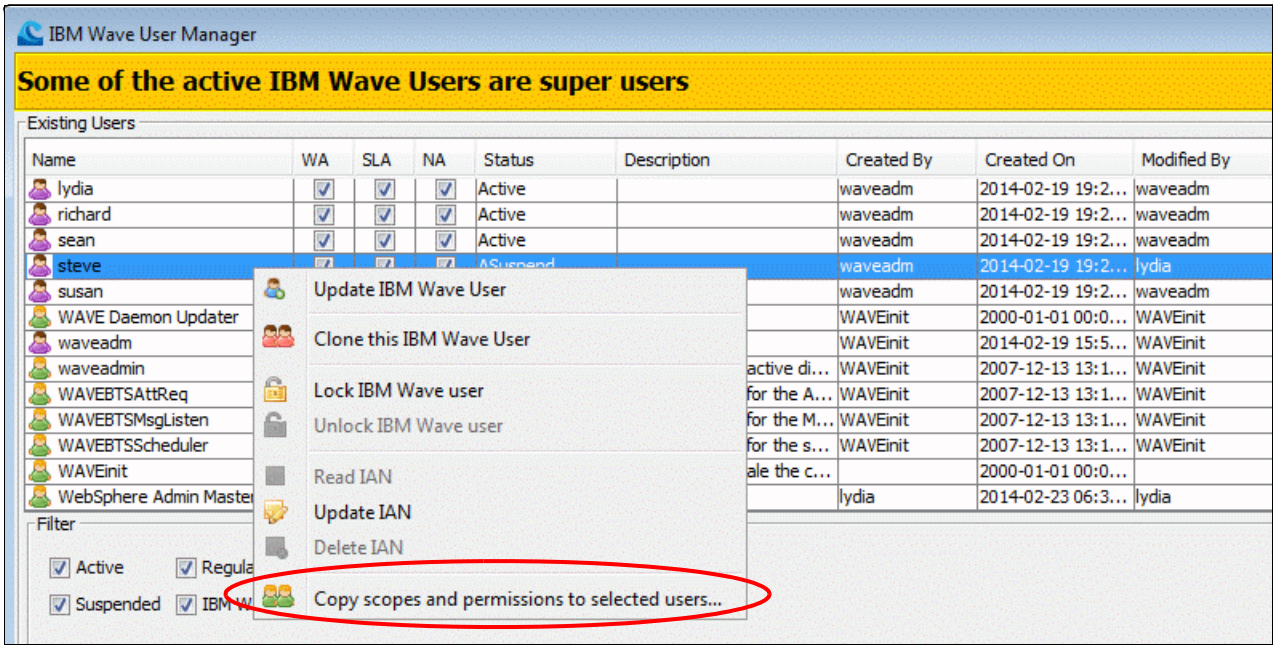


*Figure 5-101 Select Copy scopes and permissions to selected users*

2. In the Copy Selected Scopes and Permissions window, ensure that the target user is selected. From the **Select IBM Wave User to copy scopes and permissions from** menu, select the source user. As shown in Figure 5-102 on page 151, you copy the scopes and permissions from the Database Admin Master to user “steve”.

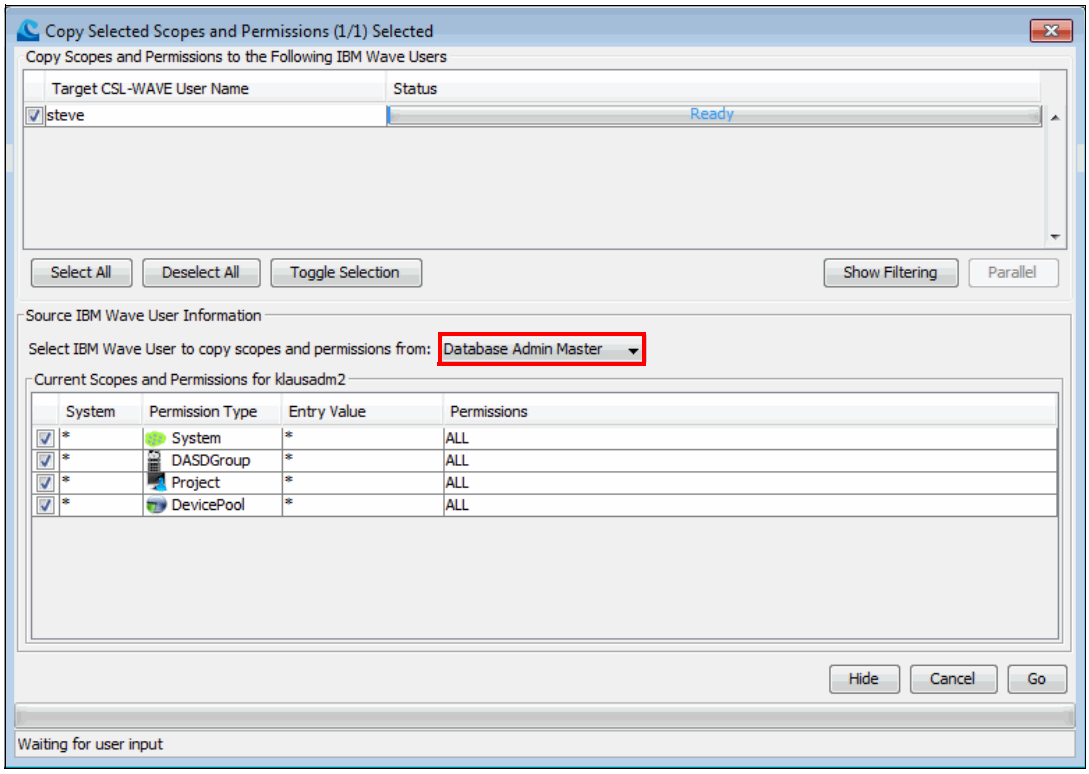


*Figure 5-102 Select IBM Wave User to copy scopes and permissions from*

Your user now is granted all of the scopes and permissions that are associated with this new class. However, unlike when the cloning action is used to create IBM Wave users, this cloning action does not affect the other attributes of the user, including password, challenge question, and description.

3. Because the previous examples show the use of the Description field to organize IBM Wave users by class, it is now important to modify this field to match the class from which you cloned the scopes and permissions. Right-click the user to which you just added a class of permission and select **Update IBM Wave User**. In the Update IBM Wave User window, change the Description field to match the name of the newly assigned class of users and click **Update**, as shown in Figure 5-103.

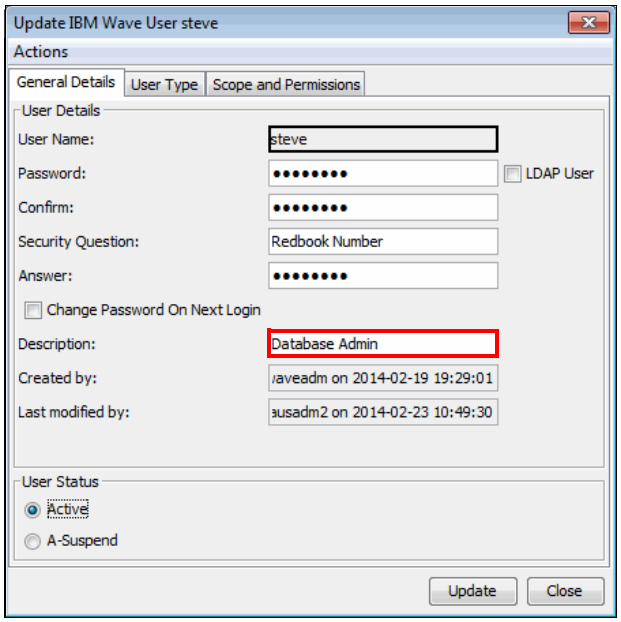


*Figure 5-103 Update the User's Description field to match the new scopes and permissions*

4. The main IBM Wave User Manager window now correctly reflects the newly assigned class of user, as shown in Figure 5-104. You can now click the **Description** column header to sort the IBM Wave users by class and see that your user that is listed alongside the template and other users of that class.

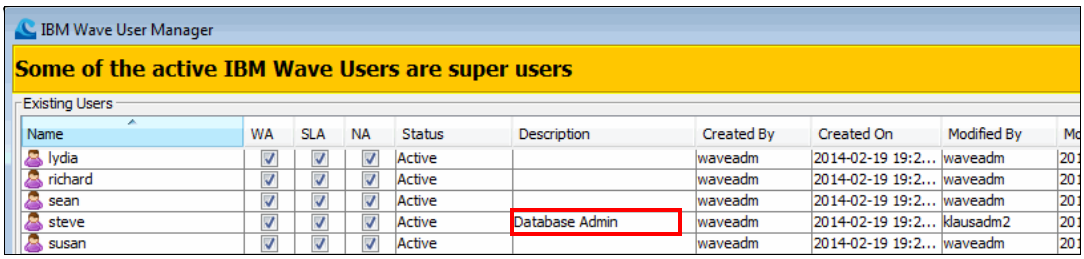


*Figure 5-104 New value in the Description field*

## Suspending an IBM Wave user

Occasionally, your organization has employees that choose to leave or move to positions that do not require access to a z/VM systems management tool. In these cases, it is important to revoke the employee's IBM Wave user ID. Because of auditing requirements in certain industries, IBM Wave makes it impossible for an IBM Wave user to be entirely deleted. Instead, a user is administratively suspended indefinitely when a particular ID is no longer needed.

Before the user is suspended, complete the following steps:

1. Delete all permissions to ensure that the IBM Wave user cannot be used if it is somehow reactivated.
2. Change the Description field of the user to Suspended to help sort and differentiate indefinitely suspended users from all other IBM Wave users.
3. Administratively suspend the user by right-clicking a user that is listed in the IBM Wave User Manager and select **Update IBM Wave User**. In the Update IBM Wave User window, change the User Status to A-Suspend, as shown in Figure 5-105, and then click **Update**.

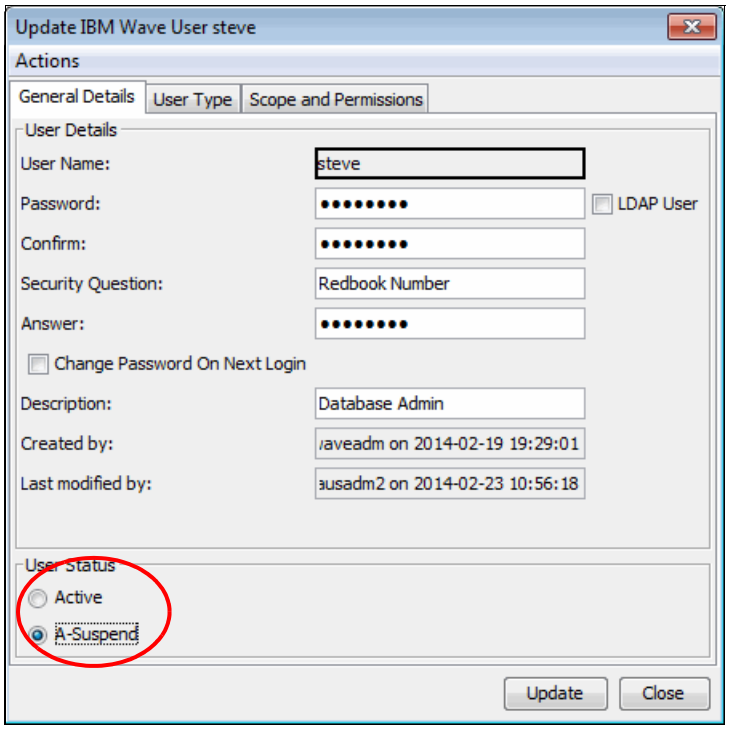


*Figure 5-105 Changing s user status to A-Suspend*

The IBM Wave user is now suspended.

## 5.10 Backup option for IBM Wave

Because IBM Wave runs on a Linux guest (typically running on z/VM), this guest can be backed up and restored by using the same processes as any other Linux guest. For more information about general backup techniques for z/VM and Linux guests, see *Set up Linux on IBM System z for Production*, SG24-8137, and *Using z/VM v 6.2 Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8039.

In addition to traditional backup techniques, IBM Wave administrators can use IBM Wave to back up the IBM Wave database by using SQL scripts. The IBM Wave backup option is useful for capturing snapshots of the IBM Wave metadata to reverse changes to pure-metadata objects, such as projects, site-defined groups, and IBM Wave users. This section describes this level of backup.

## 5.10.1 Entering Single User Mode

To back up the IBM Wave database, your IBM Wave session must enter Single User Mode. Under this mode, IBM Wave allows only a single user session to be active at one time. If other IBM Wave users are active when you attempt to enter Single User Mode, you are prompted to send messages to the other active IBM Wave users to instruct them to log off. You then can request a logoff option on their behalf, which opens the logoff confirmation window on their workstation.

To enter Single User Mode, complete the following steps:

1. Click **Administrative** → **Toggle Single User Mode** from the menu bar, as shown in Figure 5-106.

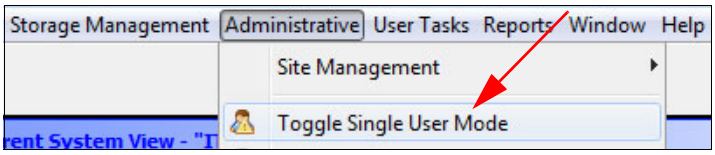


*Figure 5-106 Toggle Single User Mode*

2. If no other users are active, you immediately enter Single User Mode, which prevents other IBM Wave users from logging in. However, if other IBM Wave users are in active sessions, you see the Request Users Logoff window, as shown in Figure 5-107. Click **Send Message**.

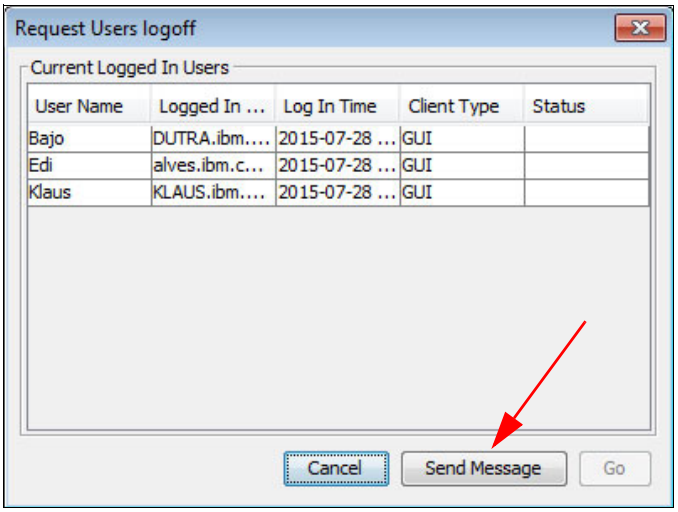


*Figure 5-107 Request Users Logoff window*

3. The Broadcast Message window opens, in which you can send a message to all other active users. Enter your message in the text area and then click **Go**, as shown in Figure 5-108 on page 155.

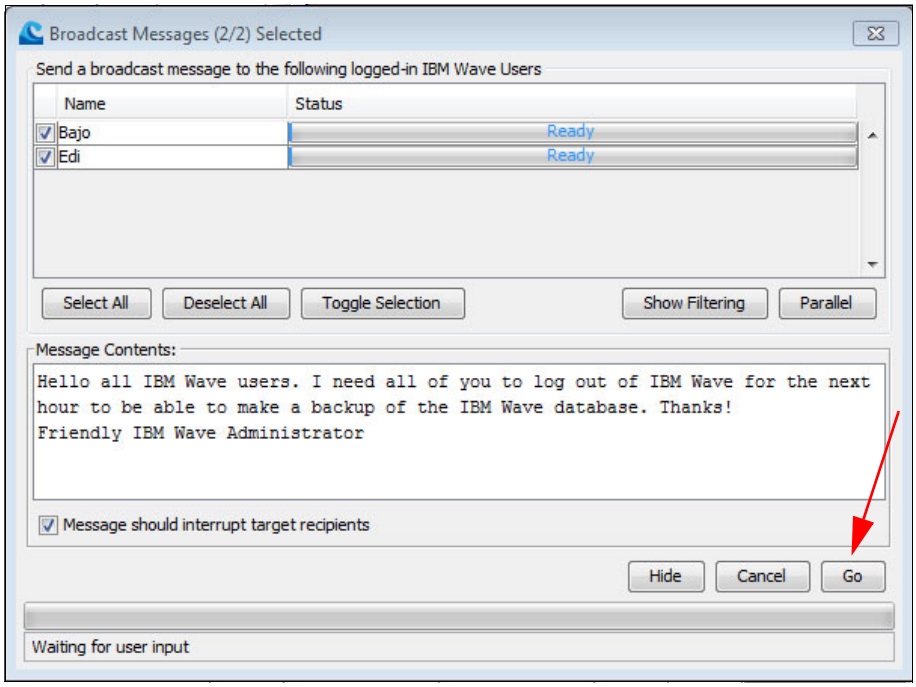


*Figure 5-108 Broadcast Messages window*

The other IBM Wave users receive a notification, as shown in Figure 5-109.

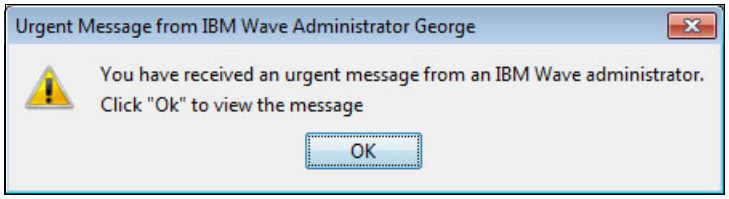


*Figure 5-109 Users receive a message from the administrator*

Then, a new window opens and shows the message in Figure 5-110.

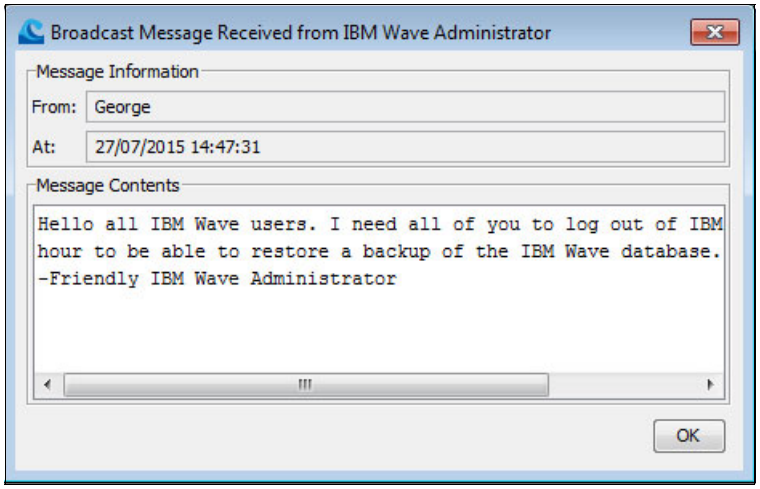


*Figure 5-110 Administrator message requesting a logoff*

The other IBM Wave users, after reading this message, learn that they must finish their work and begin the process of logging out. As the users log out, the users disappear from the Current Logged In Users table in the Request User Logoff window. However, it is possible that one or more users are not at their workstation and left their IBM Wave session unattended. Alternatively, it is possible that the user has a hung session if connectivity between the user's workstation and the IBM Wave server was interrupted before a graceful logoff.

IBM Wave anticipates this scenario by automatically replacing the Send Message option with a Request Logoff option after you send a message to the other IBM Wave users. For example, in Figure 5-111, the IBM Wave user "Bajo" did not log off as requested by the message from the IBM Wave administrator because they were not at their desk. By examining the Workunits tab, the IBM Wave administrator determines that Bajo is not active on IBM Wave. In this situation, an IBM Wave administrator can click the **Request Logoff** option.

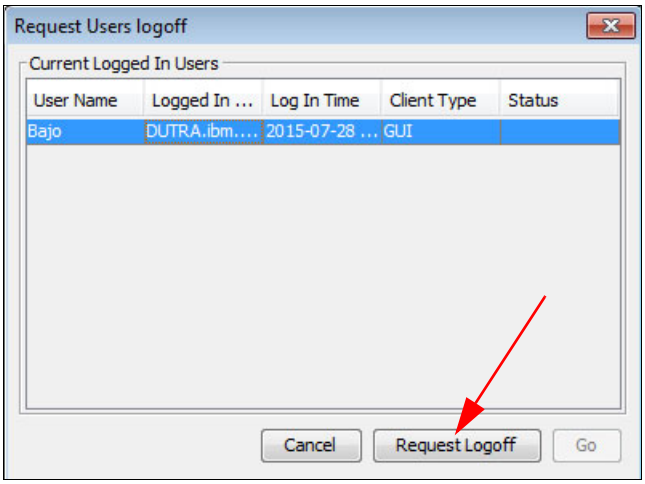


*Figure 5-111 Force an IBM Wave User off through the Request Logoff option*

The IBM Wave server then attempts to contact the selected user's IBM Wave client. If successful, it instructs the Wave client to start the logoff process. The user's IBM Wave session is remotely ended and a message is displayed, as shown in Figure 5-112. After the user clicks **OK**, the IBM Wave client closes.

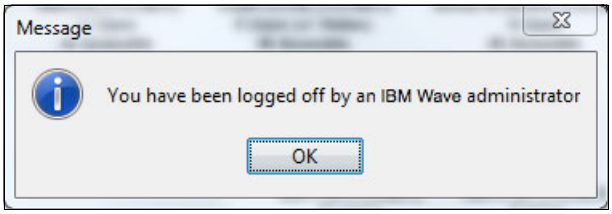


*Figure 5-112 Logged off message*

In the case of a hung session where the IBM Wave server cannot reach the IBM Wave client, the Request Logoff process does not end the user session. The session is not ended because it is possible that tasks might be in progress when the session hung and ending it in such a state might lead to unpredictable results.

Nevertheless, if the IBM Wave administrator still must end the session, the Request Logoff option is replaced by a Force logoff option if the Request Logoff action failed. Starting the Force Logoff process causes the IBM Wave server to unilaterally end a selected IBM Wave session. Because this process might be destructive, IBM Wave presents a warning message to the IBM Wave administrator, as shown in Figure 5-113.

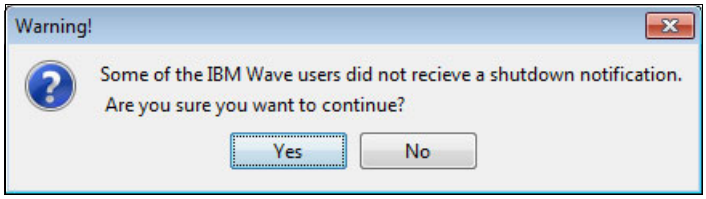


*Figure 5-113 Force Logoff Warning*

4. Click **Yes** to continue.

   The IBM Wave administrator can now successfully end all other user sessions, as shown in Figure 5-114.

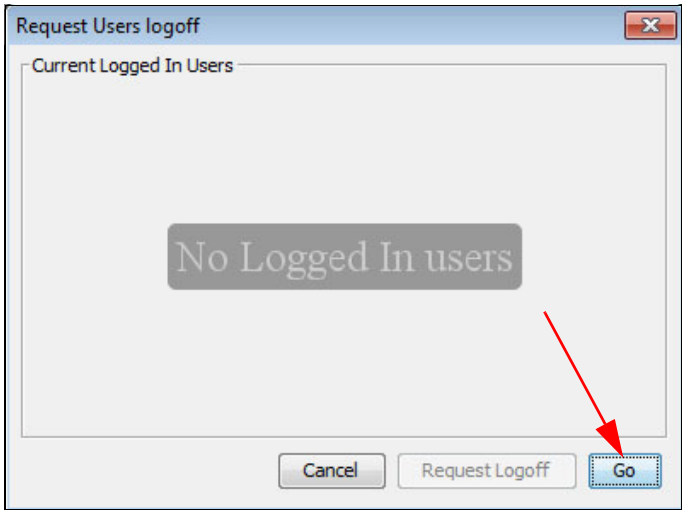


*Figure 5-114 All users are logged off*

Now that all other users are logged off, the IBM Wave administrator can close the Request Users logoff window. Closing this window automatically places IBM Wave into Single User Mode and a message is displayed to the administrator, as shown in Figure 5-115.

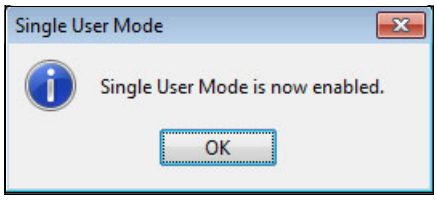


*Figure 5-115 Single User Mode is now enabled message*

## 5.10.2 Backing up the IBM Wave database

To back up the IBM Wave database, click **Backup IBM Wave for z/VM Database** from the Administrative menu, as shown in Figure 5-116.

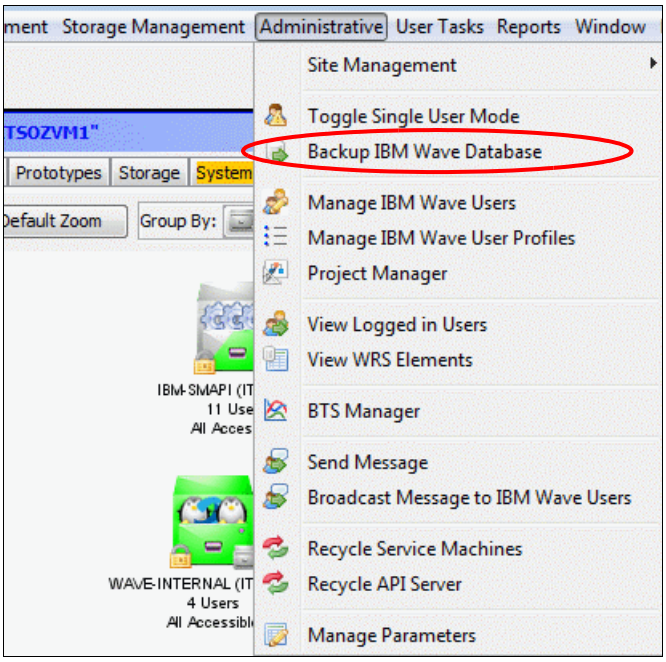


*Figure 5-116 Select the backup option*

When you select **Backup IBM Wave for z/VM Database**, the IBM Wave server examines its database and determines the size of the SQL script that is generated from this action. It then checks the capacity that is available in the volume that is holding the `/usr/wave/DBBackup` directory. This step is critical because this script can be larger than 1 GB.

After the backup is done, IBM Wave displays a message, as shown in Figure 5-117.

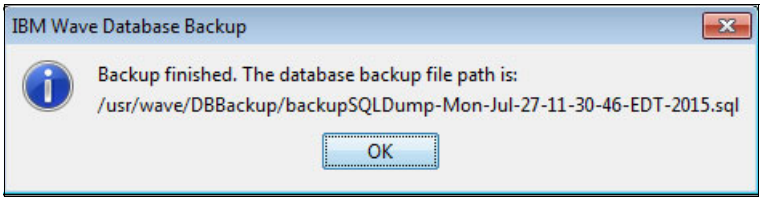


*Figure 5-117 Backup finished*

If the file system does not have enough available space, IBM Wave displays an error message, as shown in Figure 5-118, and ends the backup.

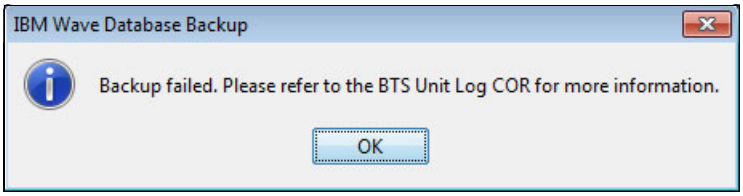


*Figure 5-118 Error message*

Open the BTS Work Unit log to see the details of the error, as shown in Figure 5-119.

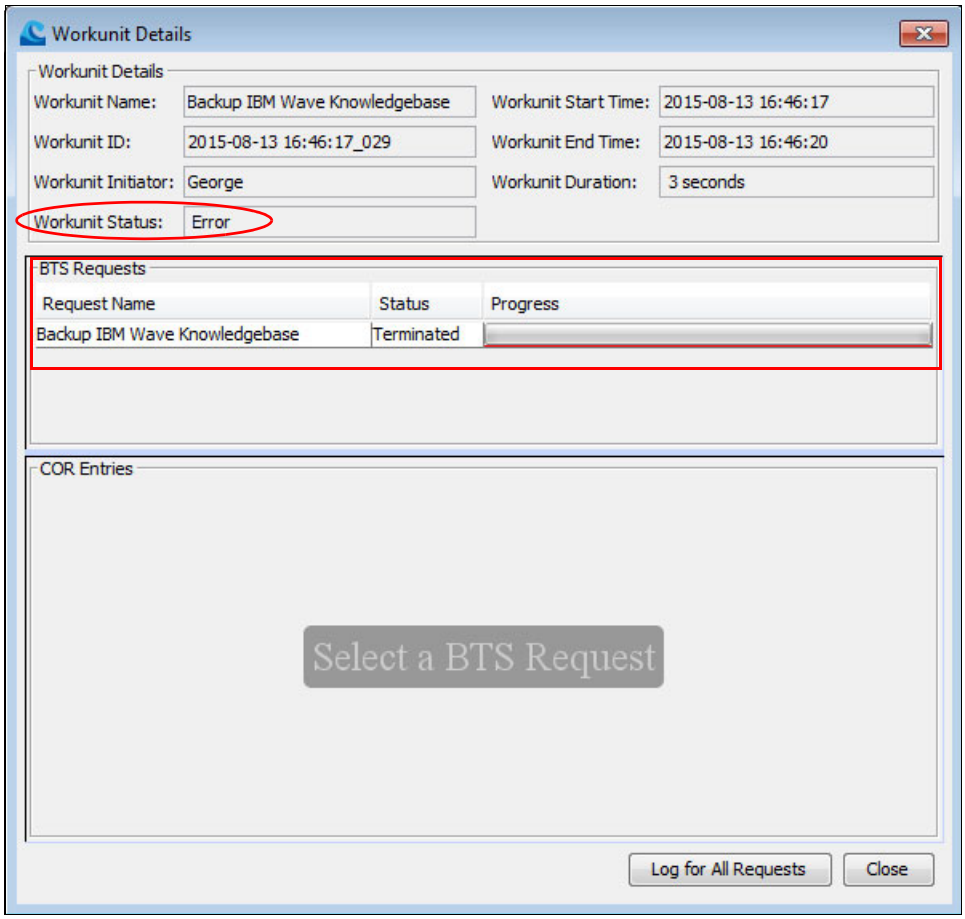


*Figure 5-119 BTS Work Unit Log of backup error*

Clicking twice opens a new window, which shows detailed messages, as shown in Figure 5-120.

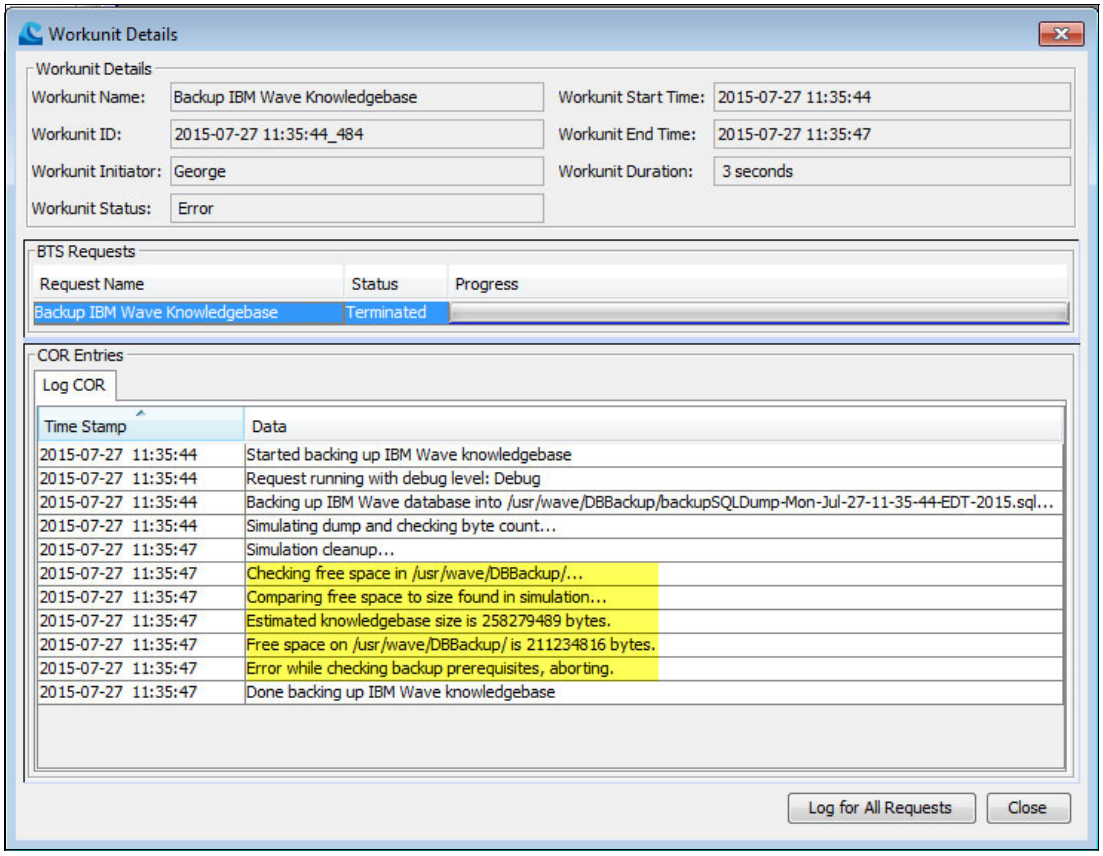


*Figure 5-120 BTS Work Unit Log - details of the backup error*

When this error occurs, an IBM Wave administrator must free sufficient space by resizing the volume that is holding the `/usr/wave/DBBackup` directory or deleting older backup files through an SSH session. The administrator then reruns the IBM Wave database backup.

If IBM Wave determines that the appropriate Wave server volume contains sufficient space, it generates an SQL script in the `/usr/wave/DBBackup` directory that is named in the following format:

`backupSQLDump-<Day>-<month>-<Day of Month>-<Hour>-<minute>-<Second>-<Year>.sql`

**Note:** If the IBM Wave database is backed up more than once in 24 hours, IBM Wave prompts the user if the current day's backup should be overwritten.

These files can be accessed by using SSH or FTP. Figure 5-1 shows an example of the `/usr/wave/DBBackup` directory with a backup file.

*Example 5-1 IBM Wave backup*

```
wavesr:/usr/wave/DBBackup # ls -latr
total 251932
drwxr-xr-x 10 waveuser root      4096 Jul 27 08:59 ..
drwxr-xr-x  2 root     root      4096 Jul 27 11:30 .
-rw-r--r--  1 root     root 257711010 Jul 27 11:30 backupSQLDump-Mon-Jul-27-11-30-46-EDT-2015.sql
```

## 5.11 Using the IBM Wave Command-Line Interface

If you set up your server or workstation to use the IBM Wave Command-Line Interface (CLI), as described in 4.3 “IBM Wave Command-Line Interface installation” on page 82, you can use it for some basic tasks, such as activating, deactivating, and querying the setup or status.

With IBM Wave V1.2 Fix Pack 0.5, the following IBM Wave commands are available:

- **`activate <Guest Name>`**

  Activates a specific z/VM guest.

- **`deactivate <Guest Name> [IMMED | WITHIN 1-32767]`**

  Deactivates a specific z/VM guest. **`IMMED`** causes the z/VM to be forced, and **`WITHIN`** ***`xxx`*** causes a z/VM signal to be sent to the z/VM guest.

  **Note:** CMS guests should be forced because they cannot handle a z/VM signal. Linux guests should be signaled.

- **`details <Guest Name>`**

  Returns details about the specific z/VM guest.

- **`status <BTS Workunit ID>`**

  Retrieves the status of a BTS Workunit ID.

- **`signalActivationDone <Guest Name>`**

  Signals to the BTS that the activation of the specific z/VM guest is complete.

### Wave CLI parameters

Here are the IBM Wave CLI parameters:

- **`-u <IBM Wave User>`**

  This mandatory parameter specifies the IBM Wave user through which the IBM Wave CLI commands are run. All generated BTS Work Units and log output is attributed to the specified user.

- **`-bts <BTS IP Address>`**

  This mandatory parameter specifies the IP address of the WAVESRV on which the BTS is located.

- **`-p <IBM Wave User Password>`**

  This optional parameter can be used to specify the IBM Wave user's password. If this parameter is not specified, the IBM Wave CLI prompts interactively for the user’s password to be entered in a secure manner.

- **`-tvp <z/VM system Name>`**

  This parameter is mandatory for several IBM Wave CLI commands and specifies the z/VM system name to which the IBM Wave CLI commands refer.

- **`-f <filename>`**

  The file name specifies the full path of a local file containing a list of IBM Wave CLI commands that are run one after the other, serially, as in batch mode.

**Note:** Starting the IBM Wave CLI with the -f (file name) 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.

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:

- Each line can hold one command at most.
- Empty lines are ignored.
- Use white space characters to separate tokens, as shown in the following example:

```
DETAILS WAVESRV
ACTIVATE     TRAUNER
DEACTIVATE CMSGUEST    IMMED
DEACTIVATE LINGUEST WITHIN 10
```

- **`-domain`**

  The Active Directory domain, if the supplied IBM Wave user is an (Windows) AD user.

- **`-forceLogin`**

  Used if a previous IBM Wave CLI call failed to the extent that the IBM Wave resource serialization (WRS) entry for the CLI session is retained in the IBM Wave knowledge base.

- **`-ignoreIAN`**

  Specifies that the command is run regardless of whether an Intelligent Active Note (IAN) exists for the specified object.

- **`-v`**

  Specifies that verbose output is used.

- **`-port`**

  Specifies the port to use to connect to the BTS. If left unspecified, the default port of 3300 is used.

- **`-noCOR`**

  If specified, the COR entry is not displayed.

- **`-forceTLS`**

  Establish a connection with the BTS only if TLS encryption is available.

## 5.11.1 Using CLI in Windows

Calling the IBM Wave CLI is a Java call, as shown in Figure 5-121 on page 163.

```
java -cp WAVECommon120.jar;log4j-1.2.16.jar;mariadb-java-client-1.1.1.jar;
jdom.jar;commons-collections-3.2.jar;commons-dbcp-1.2.2.jar;
commons-pool-1.4.jar;jsch-0.1.42.jar;WAVECLI120.jar com.CSL.WAVE.CLI.WAVECLI -u
roland -bts 9.12.7.89 -tvp ITSOZVM1 -v details WAVEWRKS > roland1.txt
```

*Figure 5-121 Java Wave CLI call*

Figure 5-121 displays the details of the WAVEWRKS z/VM guest that uses the IBM Wave user ID `roland` on the Wave Server at `9.12.7.89`. The VM is `ITSOZVM1`, and we used the verbose (**-v**) option. We redirected the output to a file that is named `roland1.txt`. The password Input for the Wave user ID is prompted.

**Note:** There is no space after the >> ; << in the java JAR file listing. We added a blank to better display the command.

To shorten this command, you might use and adapt the `wavecli.bat` file. We modified the wavecli.bat file for the needs of our project, as described in “Wave CLI helper file” on page 284.

## Details command

To display details about a z/VM guest, run the command that is shown in Figure 5-122.

```
wavecli -u roland -bts 9.12.7.89 -tvp ITSOZVM1 details trauner > trauner1.txt
```

*Figure 5-122 CLI displays details about z/VM user trauner - redirect output to file*

The command displays a request to enter the password, as shown in Figure 5-123.

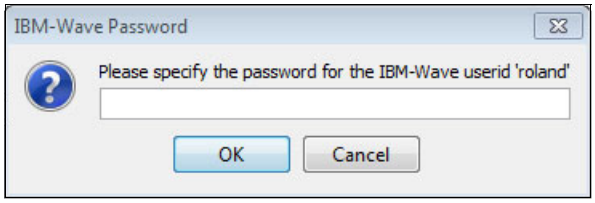


*Figure 5-123 WAVECLI password prompt*

The output is written to `trauner1.txt`, as shown in Figure 5-124.

```
IBM Wave CLI> BTS workunit completed: 2015-07-29 12:04:03_689
IBM Wave CLI> ************ COR Entry Start ************
IBM Wave CLI>
IBM Wave CLI> 2015-07-29  00:04:03 BTS Workunit: CLI Command Details, Request: Get z/VM guest details for trauner
(ITSOZVM1) started
IBM Wave CLI> 2015-07-29  00:04:03 Request running with debug level: Debug
IBM Wave CLI> 2015-07-29  00:04:03 Retrieving data from IBM Wave knowledgebase for z/VM guest trauner in z/VM
system ITSOZVM1...
IBM Wave CLI> 2015-07-29  00:04:03
IBM Wave CLI> 2015-07-29  00:04:03 Guest name: TRAUNER, z/VM system: ITSOZVM1, SDG: USER-LOCAL
IBM Wave CLI> 2015-07-29  00:04:03 Status: Inactive, Connectable: false
IBM Wave CLI> 2015-07-29  00:04:03
IBM Wave CLI> 2015-07-29  00:04:03 Project: ITSOProj1
IBM Wave CLI> 2015-07-29  00:04:03 Functionality: N/A
IBM Wave CLI> 2015-07-29  00:04:03 Description: N/A
IBM Wave CLI> 2015-07-29  00:04:03
IBM Wave CLI> 2015-07-29  00:04:03 Memory: (Minimum) 64, (Maximum) 1024
IBM Wave CLI> 2015-07-29  00:04:03 CPU Count: 1
IBM Wave CLI> 2015-07-29  00:04:03 Disk Space: 104222720 Bytes (0.1 GB)
IBM Wave CLI> 2015-07-29  00:04:03
IBM Wave CLI> 2015-07-29  00:04:03 Created By: Roland on 2015-07-29 11:49:40
IBM Wave CLI> 2015-07-29  00:04:03 Last Modified: WAVEinit on 2015-07-29 11:58:12
IBM Wave CLI> 2015-07-29  00:04:03
IBM Wave CLI> 2015-07-29  00:04:03 Network Details:
IBM Wave CLI> 2015-07-29  00:04:03     IP: UNKNOWN         Network: 10.60.100.0     Netmask: 255.255.255.0
Broadcast: 10.60.100.255   with vlanID: N/A  to z/VM vSwitch: DTCSMAPI through NIC: F000
IBM Wave CLI> 2015-07-29  00:04:03
IBM Wave CLI> 2015-07-29  00:04:03 ------------------ Start Guest Directory Entry ------------------
IBM Wave CLI> 2015-07-29  00:04:04 USER TRAUNER ******** 64M 1024M G
IBM Wave CLI> 2015-07-29  00:04:04    CPU 00
IBM Wave CLI> 2015-07-29  00:04:04    IPL CMS
IBM Wave CLI> 2015-07-29  00:04:04    MACHINE ESA 4
IBM Wave CLI> 2015-07-29  00:04:04    CONSOLE 0009 3215
IBM Wave CLI> 2015-07-29  00:04:04    NICDEF F000 TYPE QDIO LAN SYSTEM DTCSMAPI DEVICES 3
IBM Wave CLI> 2015-07-29  00:04:04    SPOOL 000C 3505 A
IBM Wave CLI> 2015-07-29  00:04:04    SPOOL 000D 3525 A
IBM Wave CLI> 2015-07-29  00:04:04    SPOOL 000E 1403 A
IBM Wave CLI> 2015-07-29  00:04:04    LINK MAINT 0190 0190 RR
IBM Wave CLI> 2015-07-29  00:04:04    LINK MAINT 019D 019D RR
IBM Wave CLI> 2015-07-29  00:04:04    LINK MAINT 019E 019E RR
IBM Wave CLI> 2015-07-29  00:04:04    MDISK 0192 3390 9706 140 LXD023 MW
IBM Wave CLI> 2015-07-29  00:04:04 *DVHOPT LNK0 LOG1 RCM1 SMS0 NPW1 LNGAMENG PWC20150729 CRCXX
IBM Wave CLI> 2015-07-29  00:04:04 ------------------ End Guest Directory Entry ------------------
IBM Wave CLI> 2015-07-29  00:04:04 BTS Workunit: CLI Command Details, Request: Get z/VM guest details for trauner
(ITSOZVM1) ended
IBM Wave CLI> ************ COR Entry End ************
IBM Wave CLI>
IBM Wave CLI> Details RC: 0
IBM Wave CLI> Final run RC: 0
```

*Figure 5-124 CLI output from details - trauner WAVECLI command*

## Activate command

To activate user `trauner` and redirect output to `trauner2.txt`, run command that is shown in Figure 5-125.

```
wavecli -u roland -bts 9.12.7.89 -tvp ITSOZVM1 activate trauner > trauner2.txt
```

*Figure 5-125 Activate user trauner and redirect output to trauner2.txt*

IBM Wave then activates z/VM guest `trauner`, as shown in Figure 5-126 on page 165.

```
IBM Wave CLI> BTS workunit completed: 2015-07-29 12:20:29_332
IBM Wave CLI> ************ COR Entry Start ************
IBM Wave CLI>
IBM Wave CLI> 2015-07-29  00:20:29 BTS Workunit: CLI Command Activate, Request: Activate z/VM guest trauner
(ITSOZVM1) started
IBM Wave CLI> 2015-07-29  00:20:29 Request running with debug level: Debug
IBM Wave CLI> 2015-07-29  00:20:29 Retrieving data from IBM Wave knowledgebase for z/VM guest trauner in z/VM
system ITSOZVM1...
IBM Wave CLI> 2015-07-29  00:20:29 Activating z/VM guest trauner in z/VM system ITSOZVM1...
IBM Wave CLI> 2015-07-29  00:20:29 z/VM guest trauner in z/VM system ITSOZVM1 activated successfully.
IBM Wave CLI> 2015-07-29  00:20:29 Performing volatile DEQs for Request...
IBM Wave CLI> 2015-07-29  00:20:29 BTS Workunit: CLI Command Activate, Request: Activate z/VM guest trauner
(ITSOZVM1) ended
IBM Wave CLI> ************ COR Entry End ************
IBM Wave CLI>
IBM Wave CLI> Activate RC: 0
IBM Wave CLI> Final run RC: 0
```

*Figure 5-126 CLI activate results for z/VM guest trauner*

If you run another **`details`** command, as shown in Figure 5-122 on page 163, the Status displays `Active`.

## Deactivate command

To deactivate user `trauner` and redirect output to `trauner4.txt`, run the command that is shown in Figure 5-127.

```
wavecli -u roland -bts 9.12.7.89 -tvp ITSOZVM1 deactivate trauner IMMED
```

*Figure 5-127 Deactivate user trauner and redirect output to trauner4.txt*

IBM Wave then forces off the z/VM guest `trauner`, as shown in Figure 5-128.

```
IBM Wave CLI> BTS workunit completed: 2015-07-29 12:33:03_258
IBM Wave CLI> ************ COR Entry Start ************
IBM Wave CLI>
IBM Wave CLI> 2015-07-29  00:33:03 BTS Workunit: CLI Command Deactivate, Request: Deactivate z/VM guest trauner
(ITSOZVM1) started
IBM Wave CLI> 2015-07-29  00:33:03 Request running with debug level: Debug
IBM Wave CLI> 2015-07-29  00:33:03 Retrieving data from IBM Wave knowledgebase for z/VM guest trauner in z/VM
system ITSOZVM1...
IBM Wave CLI> 2015-07-29  00:33:03 Validating the guest status...
IBM Wave CLI> 2015-07-29  00:33:03 Validating z/VM guest trauner in z/VM system ITSOZVM1...
IBM Wave CLI> 2015-07-29  00:33:03 Deactivating z/VM guest trauner in z/VM system ITSOZVM1...
IBM Wave CLI> 2015-07-29  00:33:03 Performing volatile DEQs for Request...
IBM Wave CLI> 2015-07-29  00:33:03 BTS Workunit: CLI Command Deactivate, Request: Deactivate z/VM guest trauner
(ITSOZVM1) ended
IBM Wave CLI> ************ COR Entry End ************
IBM Wave CLI>
IBM Wave CLI> Deactivate RC: 0
IBM Wave CLI> Final run RC: 0
```

*Figure 5-128 CLI deactivate results for z/VM guest trauner*

## Status command

To get information about a specific BTS Workunit ID, run the command that is shown in Figure 5-129.

.

```
wavecli -u roland -bts 9.12.7.89 -tvp ITSOZVM1 status 2015-07-29 15:30:56_909
```

*Figure 5-129 CLI status command for BTS Workunit ID 2015-07-29 15:30:56_909*

The Workunit ID can be found at the beginning of the output of **wavecli** commands, as shown in Figure 5-128 on page 165, where you see the BTS Workunit ID of 2015-07-29 12:33:03_258. BTS Workunit IDs can also be found in the GUI.

The output in this case shows that the BTS work unit completed successfully, as shown in Figure 5-130.

```
IBM Wave CLI> BTS workunit completed: 2015-07-29 15:40:47_194
IBM Wave CLI> ************ COR Entry Start ************
IBM Wave CLI>
IBM Wave CLI> 2015-07-29  15:40:47 BTS Workunit: CLI Command status, Request: Query Workunit Status started
IBM Wave CLI> 2015-07-29  15:40:47 Request running with debug level: Debug
IBM Wave CLI> 2015-07-29  15:40:49 BTS Workunit 2015-07-29 15:30:56_909 has ended. RC = 0
IBM Wave CLI> 2015-07-29  15:40:49 BTS Workunit: CLI Command status, Request: Query Workunit Status ended
IBM Wave CLI> ************ COR Entry End ************
IBM Wave CLI>
IBM Wave CLI> status RC: 0
IBM Wave CLI> Final run RC: 0
```

*Figure 5-130 CLI status command results*

### signalActivationDone command

This command (Figure 5-131) signals to the BTS that the activation of a z/VM guest is complete. It releases the lock for a user object in BTS. This command is best used for problem determination purposes.

```
wavecli -u roland -bts 9.12.7.89 -tvp ITSOZVM1 signalActivationDone trauner
```

*Figure 5-131 CLI signalActivationDone command*

Figure 5-131 shows the release of the lock.

```
IBM Wave CLI> BTS workunit completed: 2015-07-29 14:37:36_151
IBM Wave CLI> ************ COR Entry Start ************
IBM Wave CLI>
IBM Wave CLI> 2015-07-29  14:37:36 BTS Workunit: CLI Command Signal Activation Done, Request: Notify Sync Object
started
IBM Wave CLI> 2015-07-29  14:37:36 Request running with debug level: Debug
IBM Wave CLI> 2015-07-29  14:37:36 Releasing lock for object trauner...
IBM Wave CLI> 2015-07-29  14:37:36 Lock released for object trauner from 9.12.5.132
IBM Wave CLI> 2015-07-29  14:37:36 BTS Workunit: CLI Command Signal Activation Done, Request: Notify Sync Object
ended
IBM Wave CLI> ************ COR Entry End ************
IBM Wave CLI>
IBM Wave CLI> Signal Activation Done RC: 0
IBM Wave CLI> Final run RC: 0
```

*Figure 5-132 CLI signalActivationDone results*

## 5.12 Communication-Less Connection technology

Communication-Less Connection (CLC) is an IBM Wave patented technology that can help you solve problems. In a non-virtualized environment, when a server fails to load one of its operating system components, the Linux system administrator can physically go to the server and plug in a screen and keyboard to solve the problem. Use the 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.

Without CLC, the only way to solve a problem in z/VM was to open a 3270 session to the failing virtual server. IBM Wave provides standard 3270 access to the z/VM LPAR. However, 3270 access is not helpful when the applications you need to help solve the problem cannot be run.

The IBM Wave CLC technology is dynamic. There is no need for Linux virtual server configuration or definitions. You can access CLC directly from the user interface. The only requirement to use CLC is that the z/VM LPAR is reachable through TCP/IP. When the entire z/VM LPAR is unreachable, CLC is not available.

**Note:** Editors such as `vi` are not available in CLC.

Details about using CLC can be found in “CLC Access” on page 100 in 5.3.2 “Current System View” on page 96.

6

# Managing IBM z/VM

This chapter describes how to use the IBM Wave for z/VM (IBM Wave) interface to provision, deprovision, and manage virtual hardware, including systems, disk, and networking equipment. Typically, these tasks are the responsibility of a z/VM systems programmer.

This chapter includes the following topics:

- Managing z/VM user IDs
- Managing Device Pools
- Managing DASD
- Managing the network
- Managing z/VM SSI clusters

# 6.1 Managing z/VM user IDs

This section describes the following topics:

- Defining a z/VM user ID
- Viewing the details of a z/VM user ID
- Activating a z/VM user ID
- Logging in to Conversational Monitoring System as new z/VM user ID
- Deactivating a z/VM user
- Removing a z/VM user ID

**Note:** For the purposes of this publication, consider *z/VM user ID* to be synonymous with the terms *guest* and *virtual machine* (VM).

## 6.1.1 Defining a z/VM user ID

To define a z/VM user ID, complete the following steps:

1. Click **Hardware Viewer** to display your physical systems and logical partitions. Then, select the z/VM logical partition in which you want to provision a new z/VM user ID, as shown in Figure 6-1.

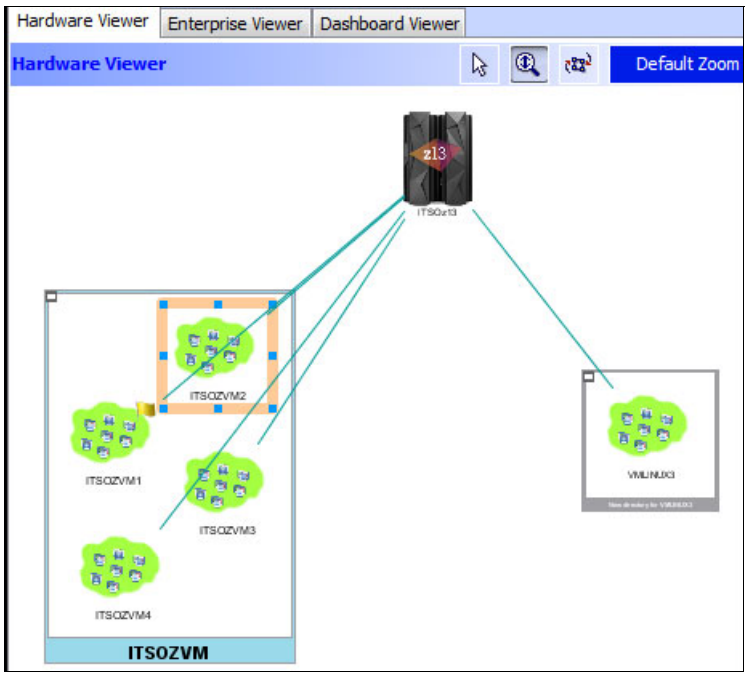


*Figure 6-1 Select a z/VM LPAR from the Hardware Viewer*

2. Ensure that the Group By menu is set to **Site Defined Group** to display icons that represent each of the Site Defined Groups, as shown in Figure 6-2.

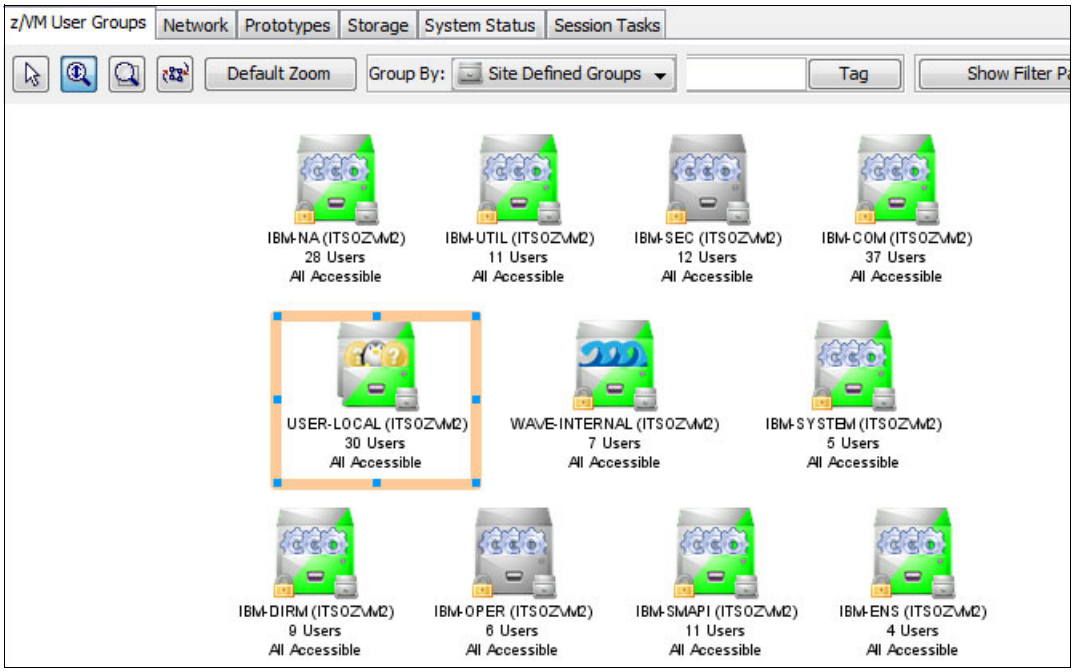


*Figure 6-2 Select Site Defined Groups from the Group By menu*

**Note:** A Site Defined Group is an IBM Wave concept with which site administrators can define logical groupings of related z/VM guests. During auto-discovery, IBM Wave creates and locks numerous Site Defined Groups that contain IBM and ISV service machines. For example, all of the service machines that are associated with DirMaint are stored in IBM-DIRM. All other z/VM user IDs that are found during auto-discovery are placed in a Site Defined Group that is called USER-LOCAL.

3. Double-click **USER-LOCAL** or a custom Site Defined Group that you created to open that group. The group is now depicted by a box that contains icons that represent each of the z/VM user IDs that are contained in this group. To create a z/VM user ID within this Site Defined Group, right-click within this box and select **z/VM User Actions** → **Create New z/VM User**, as shown in Figure 6-3.

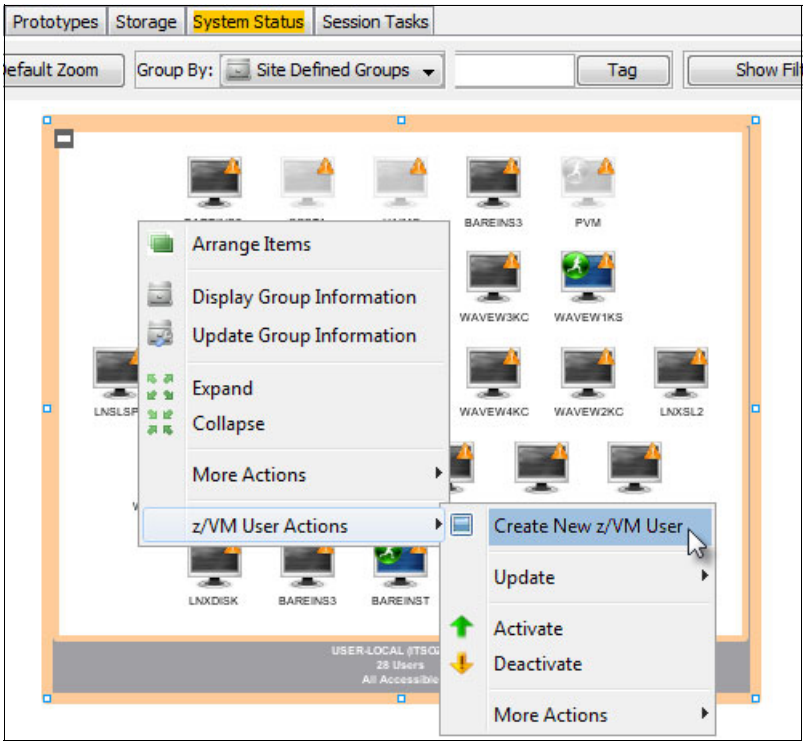


*Figure 6-3 Select Create New z/VM user*

A window opens that includes the following fields that are needed to provision a new z/VM user:

- User Name: The z/VM user ID of your new guest. This name is case-insensitive and limited to up to eight characters.
- Password: The password for logging in to z/VM as the new z/VM user ID.
- Storage: The field and menu identify the size and units of the storage that you want to define and allocate to your new guest. The menu to the right of the label “from” is used to identify the storage group from which you intend to allocate your storage.

> **Note:** The way that IBM Wave allocates and attaches minidisks might appear unusual to an experienced z/VM systems programmer. After the Create New z/VM User window is completed, IBM Wave determines the size of direct access storage device (DASD) volumes that are available in the selected storage group and allocates one or more minidisks with a capacity that totals the amount that you selected. As a user of IBM Wave, there is no way to know the size and number of individual minidisks that your guest is allocated. To have greater control over types of DASD (for example, Mod-9 versus Mod-27), create and label storage groups that separate out the available DASD volumes by type (for example, MOD9GRP and MOD27GRP). Additionally, IBM Wave attaches these new minidisks at addresses starting at 192 and incrementing by one. For example, if IBM Wave attaches three minidisks to your z/VM User, they are attached at 192, 193, and 194. After the user is created, the minidisks are in an unformatted state.

– Memory Min/Max (MB): These fields identify the minimum and maximum amount of memory that your system can be allocated by the z/VM Control Program. Depending on how your guest is configured, this setting might be useful because your guest might dynamically resize this value as needed. However, because the CP command **DEFINE STORAGE** that is responsible for this value results in a system restart, dynamic resizing of memory often cannot be done with Linux guests. Also, there is no function within IBM Wave to resize this memory value. Therefore, size the Memory Min (MB) value for the memory requirements of your workload. You cannot use IBM Wave to perform a bare metal installation of Linux if your Memory Min (MB) is smaller than 1000 MB.
– Virtual CPUs: This field sets the number of virtual processors that are available to the guest.
– Project: By using this menu, the guest can be assigned to a predefined project. The detailed information that is associated with each project can be viewed by selecting a project and clicking **Details** to the right of the menu. If your site did not define any projects, you might not be able to select any projects. If so, projects can be defined and assigned later.
– Functionality: This menu is used to express the interdependencies between Linux guests. For example, if an application server depends on a database server, the database server has a higher activation level to ensure that it is available when the application server starts. If your site is not using this function, you have an option that is called N/A (Application Level 1) only, which means that all Linux guests have the same priority for start.
– Description: This is a field that is used to insert a textual description of the function that is performed by the guest. This information is useful because of the z/VM limitation that a user name can be no longer than eight characters prevents the user name from effectively conveying the purpose of the guest. This field is stored as metadata in the database of IBM Wave.
– Default z/VM system: This field defines the default member of an single-system image (SSI) cluster that the guest loads if the entire cluster is restarted from scratch.

– Network Information: This area displays a list of available virtual network segments for your guest. The number of available segments can be zero to many, depending on the network configuration of your z/VM environment. If you defined a virtual switch, you likely see at least one segment that is automatically generated to connect your guest to your virtual switch. The values for this example are shown in Figure 6-4.

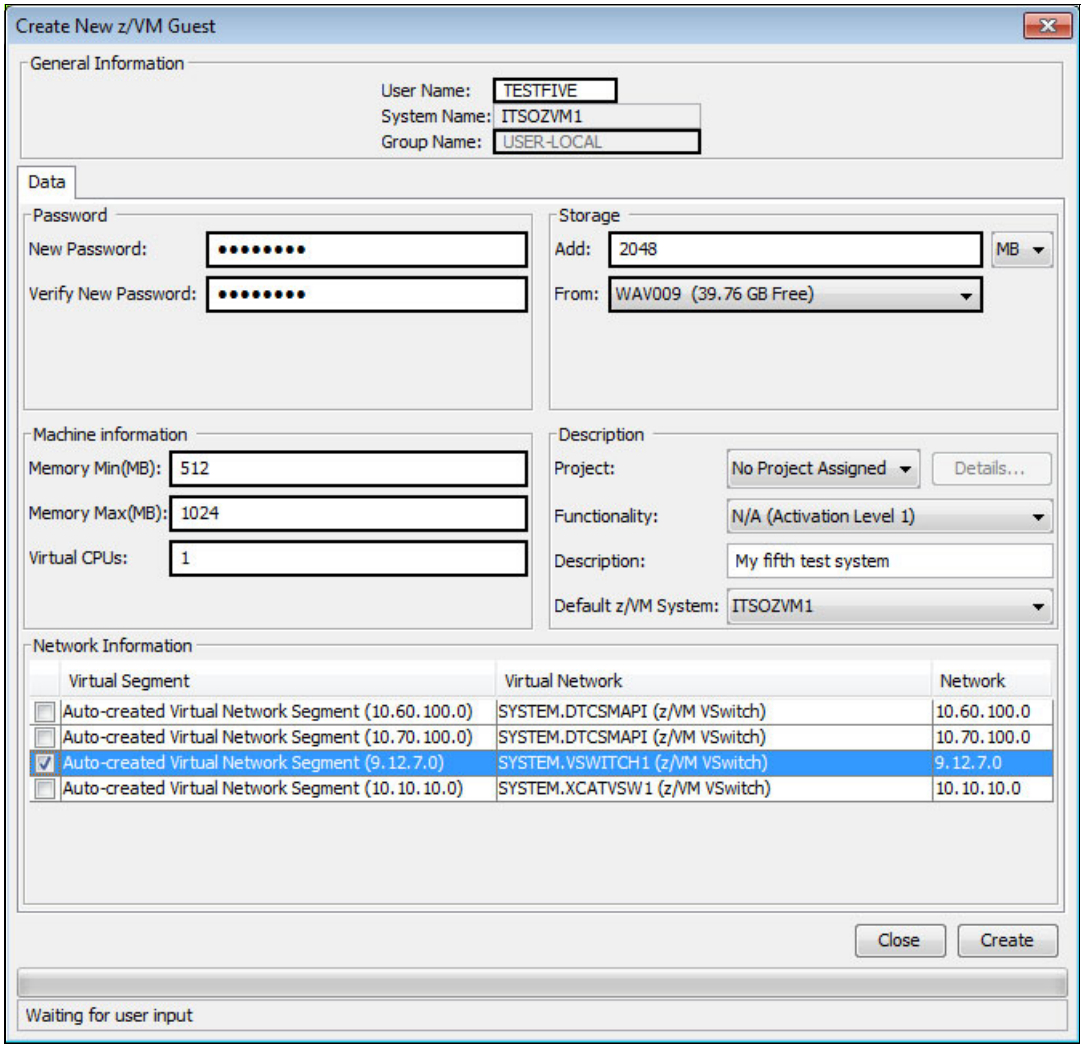


*Figure 6-4 Complete the Create New z/VM guest window*

**Note:** IBM Wave has various parameters that set the default and maximum values for Memory Min (MB), Memory Max (MB), CPUs, and Disk Space. To change these values, click **Administrative** → **Manage Parameters**, as shown in Figure 6-5 on page 175. All of the fields are under the first Threshold tab and are shown in Figure 6-5 on page 175.

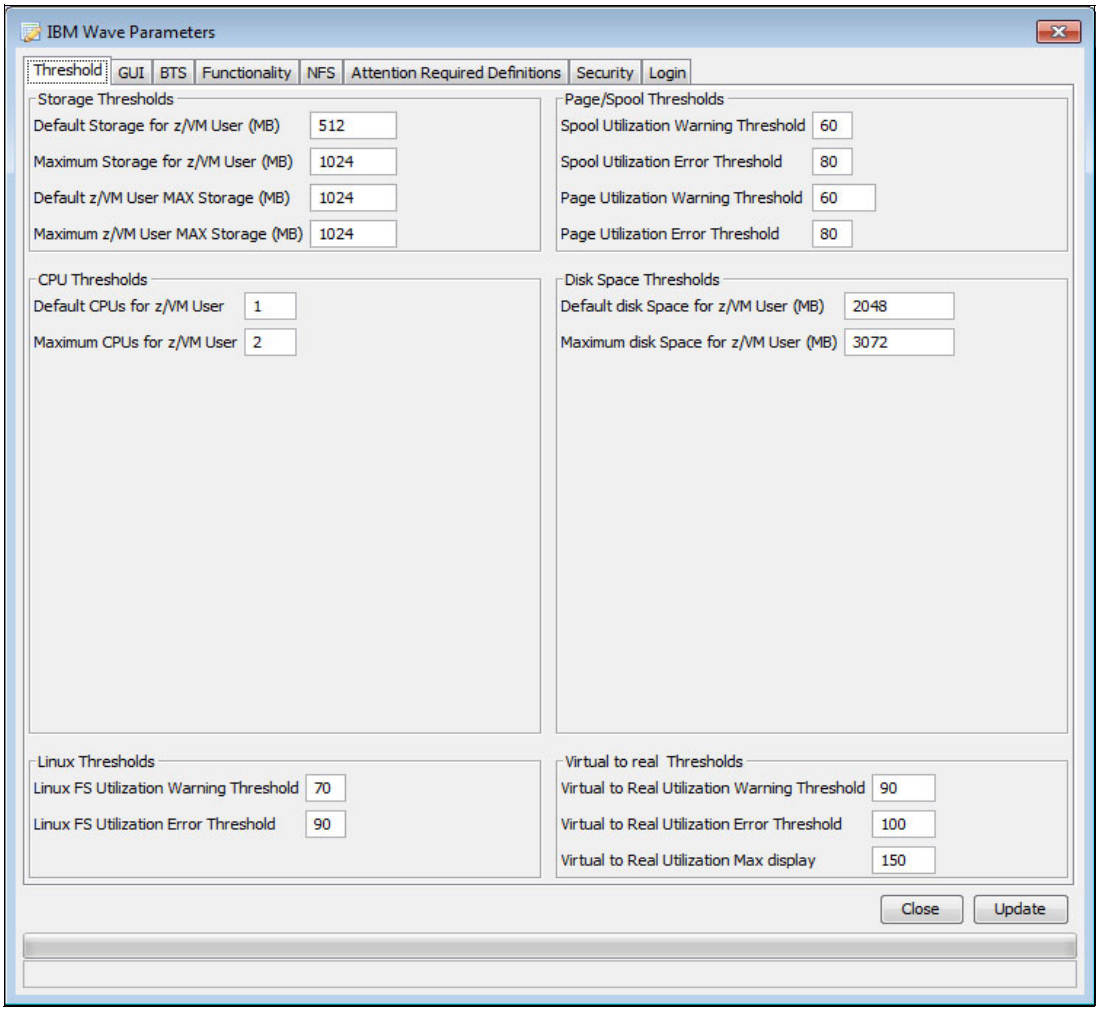


*Figure 6-5 IBM Wave Parameters Threshold tab*

4. When you complete the information in the window, click **Create**. If you have an invalid entry in any of the required fields, you receive an error message that forces you back to the window to correct the field. If you have a missing optional field (such as the Description field), you receive a message that prompts you to confirm that you want to create the guest without the optional field, as shown in Figure 6-6.

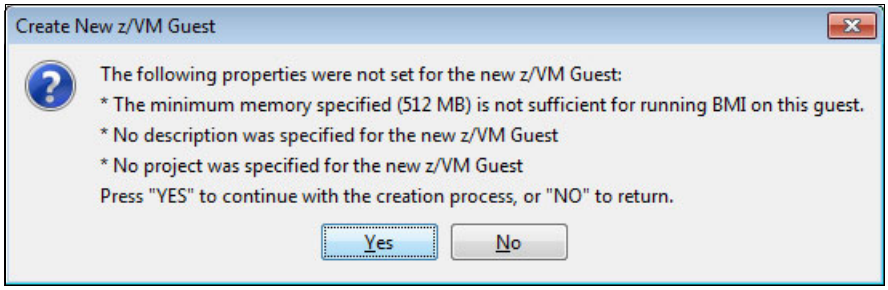


*Figure 6-6 Create New z/VM guest confirmation*

5. Click **Yes** to continue or **No** to stop and review the fields in the window. After the confirmation, a Background Task Scheduler (BTS) is started. You can see its progress by opening the BTS Work Unit Log, as shown in Figure 6-7.

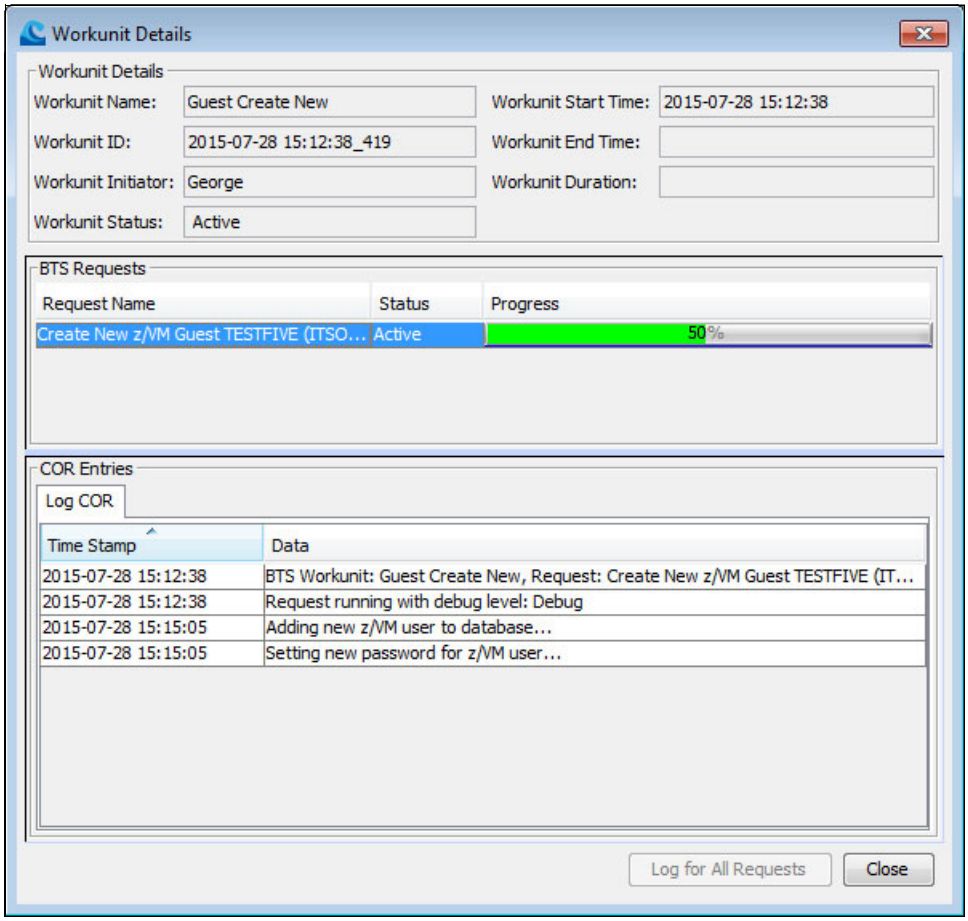


*Figure 6-7 Create New z/VM User Workunit details*

Your new z/VM user ID now appears in the selected Site Defined Group as an Inactive guest, as shown in Figure 6-8 on page 177.

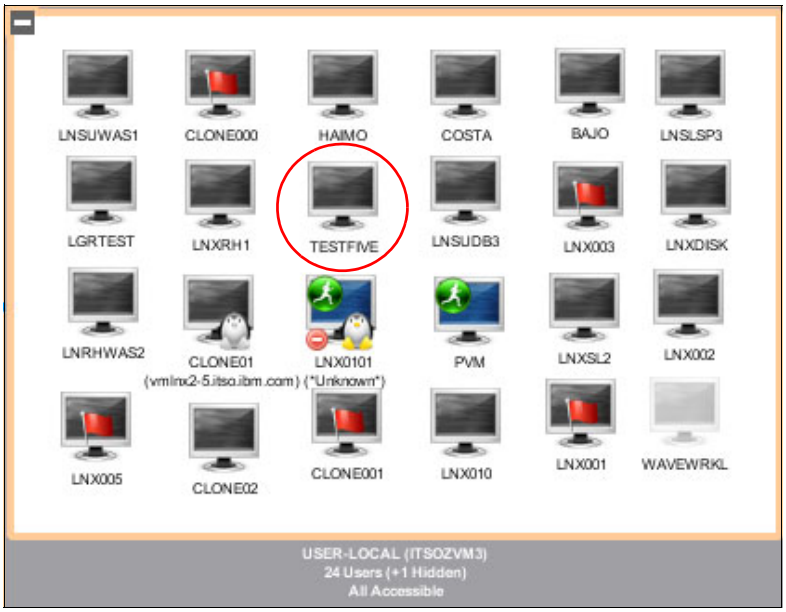


*Figure 6-8 New z/VM user is now visible in Site Defined Group*

## 6.1.2 Viewing the details of a z/VM user ID

Double-click the icon of the new z/VM user ID to display its details. The most useful information is found in the z/VM view, which displays the user ID's complete directory entry, as shown in Figure 6-9.

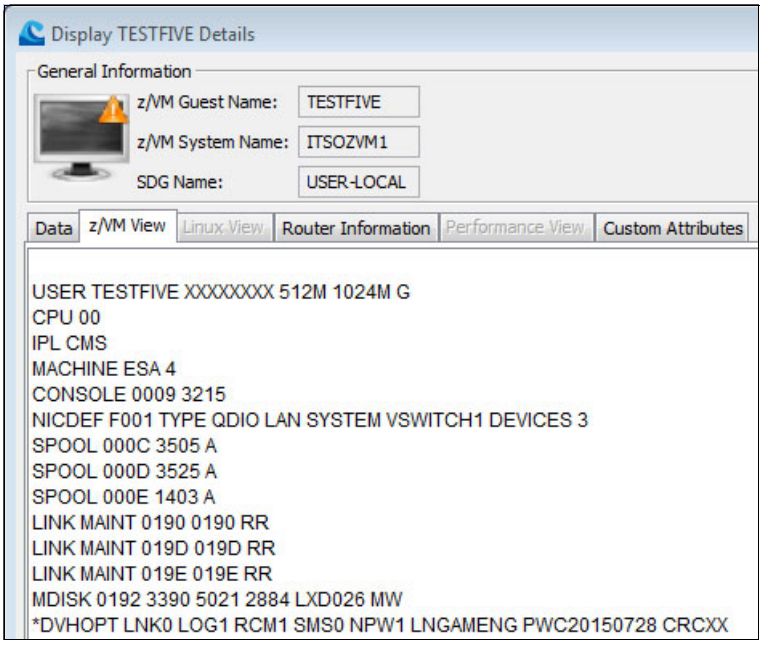


*Figure 6-9 Directory Entry in z/VM View of Display Guest Details window*

The new user's directory entry has an IPL Conversational Monitoring System (CMS) statement. This statement means that the guest automatically starts in the CMS operating system when it is activated, even if the IBM Wave user installs Linux on this guest. To start the Linux kernel, the user must manually log in to the z/VM user ID to enter the CMS environment and then enter the command IPL 0192 (assuming that the Linux kernel is installed to the 0192 minidisk). You can change this by logging in to a 3270 terminal emulator session and running the command that is shown in Example 6-1.

*Example 6-1 Change the boot disk*

```
Dirm for <<userid>> IPL 192
```

<<userid>> is the new z/VM user ID that you created.

The new user links to MAINT's 0190, 019D, and 019E minidisks in read-only mode. By using this configuration, the new guest can initially load CMS off MAINT's 0190 minidisk without having a local CMS file system. These statements also make the new user ineligible for Live Guest Relocation (LGR) to another member of an SSI environment.

To have your guest automatically start in Linux and be eligible for LGR, the z/VM user ID's directory entry must be manually edited through traditional 3270 methods, as shown in Example 6-2.

*Example 6-2 Example of changing the IPL CMS statement by using DirMaint*

```
dirm for testfive ipl 192
DMSWSF006E No read/write file mode accessed for TESTFIVE NETLOG
DVHCMD1186W File mode A is not an R/W disk or directory.  Your DIRMAINT
DVHCMD1186W NEWMAIL date/time stamp will not persist across an IPL.
DVHXMT1181R Enter the current logon password of TESTFIVE at ITSOZVM1 for
DVHXMT1181R authentication.  It will not be displayed on the
DVHXMT1181R terminal.  To exit without processing the command, just
DVHXMT1181R press ENTER.
Author note: the password was typed in and enter was pressed
DVHXMT1185W File mode A is not an R/W disk or directory.  Your request
DVHXMT1185W counter will not persist across an IPL.
DVHXMT1191I Your IPL request has been sent for processing to DIRMAINT at
DVHXMT1191I ITSOZVM1.
Ready; T=0.01/0.01 15:03:18
 DVHREQ2288I Your IPL request for TESTFIVE at * has been accepted.
 DVHBIU3450I The source for directory entry TESTFIVE has been updated.
 DVHBIU3423I The next ONLINE will take place via Diagnose 84.
 DVHRLA3891I Your DSATCTL request has been relayed for processing.
 DVHRLA3891I Your DSATCTL request has been relayed for processing.
 DVHRLA3891I Your DSATCTL request has been relayed for processing.
 DVHBIU3428I Changes made to directory entry TESTFIVE have been placed
 DVHBIU3428I online.
 DVHREQ2289I Your IPL request for TESTFIVE at * has completed; with RC
 DVHREQ2289I = 0
```

Open the Network View to see a graphical representation of the network segments that are connecting your user ID to a vSwitch. The example network view is shown in Figure 6-10.

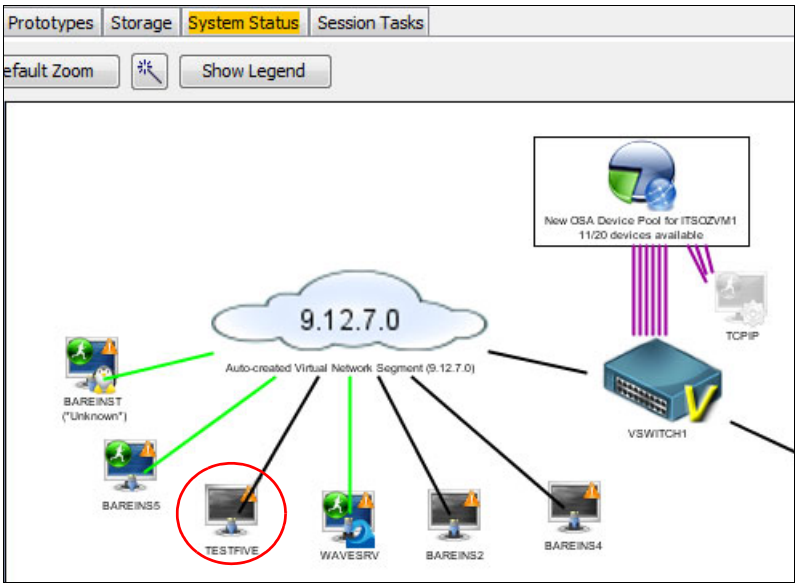


*Figure 6-10 New z/VM guest in Network View*

### 6.1.3 Activating a z/VM user ID

After a z/VM user is created, the user must be activated. This process is analogous to starting a physical server. You activate a user by completing the following steps:

1. Right-click the icon that represents your z/VM user and selecting **Activate**, as shown in Figure 6-11. Multiple z/VM users can be activated simultaneously by highlighting more than one icon before you right-click and select **Activate**.

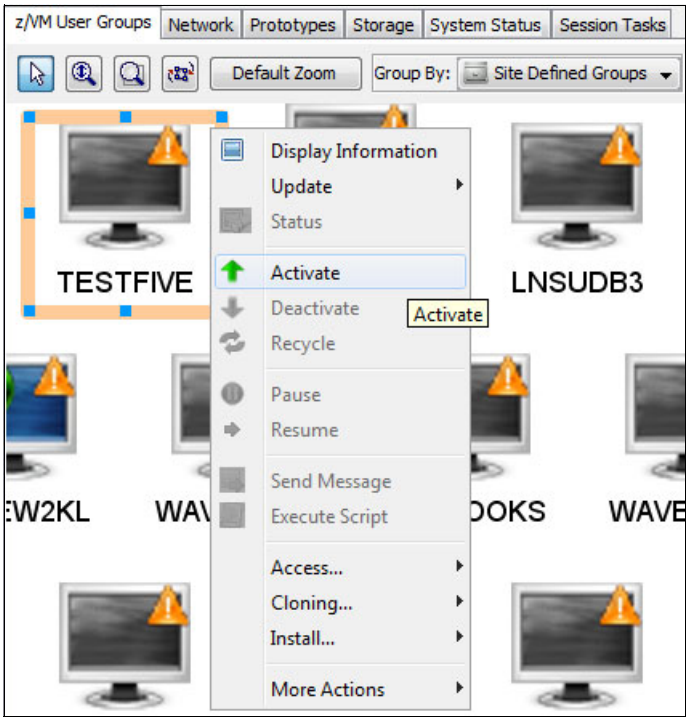


*Figure 6-11 Select Activate on a z/VM user*

2. The z/VM User Activation window opens, as shown in Figure 6-12 on page 181. Verify that the z/VM users that you want to activate are selected and then click **Go**.

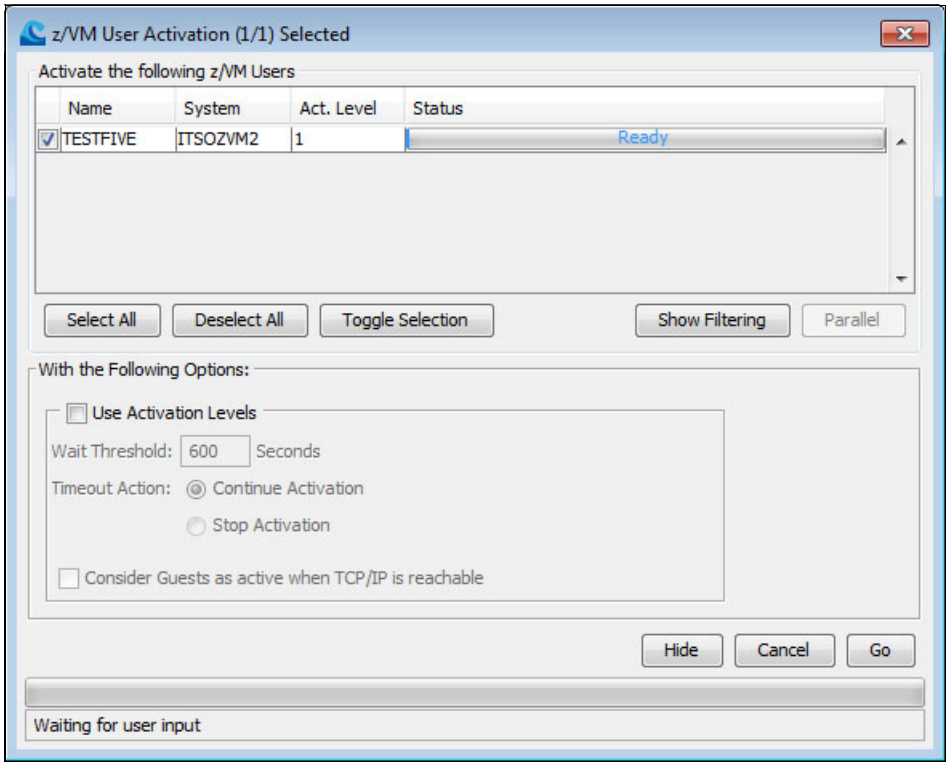


*Figure 6-12 z/VM User Activation window*

3. A window opens that indicates that the Activate z/VM User task was dispatched to the BTS. Click **OK** and close the z/VM User Activation window to return to the Current System View in the main IBM Wave window.

After the z/VM user is running, the icon that represents the new z/VM user is partially covered by a green circle that contains a running man graphic, as shown in Figure 6-13.


*Figure 6-13 User is activated and is running*

## 6.1.4 Logging in to Conversational Monitoring System as new z/VM user ID

After a z/VM user is active, an IBM Wave user can log in to CMS as this user. By logging in as such, an IBM Wave user can perform z/VM systems management and programming tasks by using the traditional 3270-based CMS interface.

To enter this environment, complete the following steps:

1. Right-click the icon of your new running z/VM user and select **Access** → **IBM Wave 3270 Console** to open an integrated 3270 terminal, as shown in Figure 6-14.

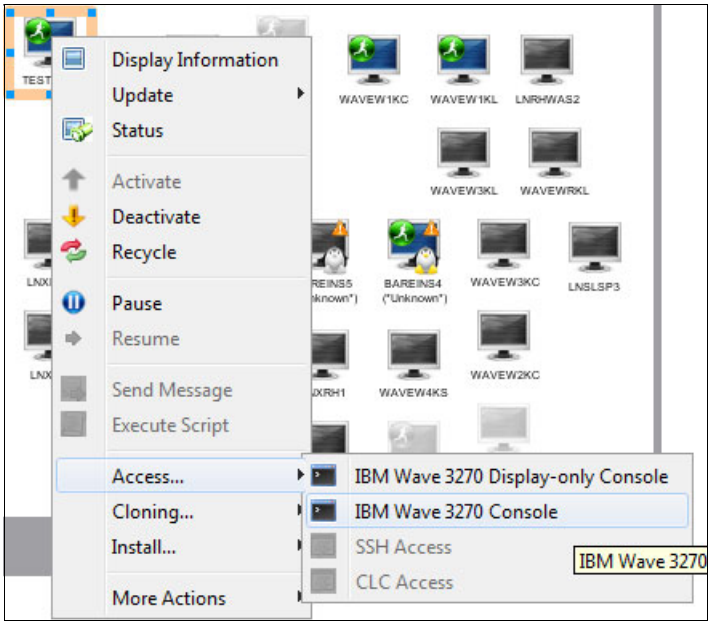


*Figure 6-14 Select the 3270 Console*

2. A window opens that prompts you for the password that is associated with this z/VM user. Enter the password and click **Login**, as shown in Figure 6-15.

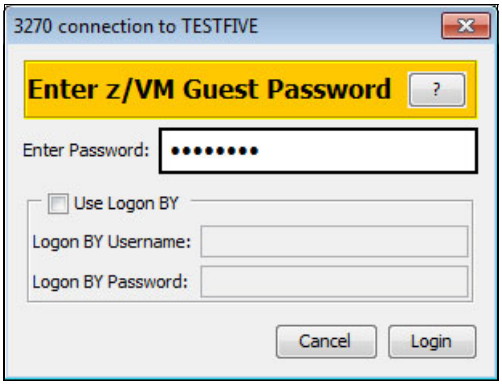


*Figure 6-15 Password prompt for 3270 connection*

A 3270 terminal session opens through an integrated Java based 3270 terminal emulator and you are logged in as the z/VM user that you selected.

After you are logged on, you can access the traditional CMS environment by running the **`IPL CMS`** command. However, because IBM Wave is designed with the assumption that z/VM users that are created through the GUI are used for Linux guests, your new z/VM user lacks a 191 minidisk at file mode A, which means that there is no local read/write file system that is formatted for CMS. Because of this fact, numerous CMS utilities issue an error similar to the error that is shown in Example 6-3 on page 183.

*Example 6-3 Error that is issued when a 191 minidisk is missing*

```
DMSACP113S A(191) not attached or invalid device address
```

By running a **QUERY DASD** command (as shown in Example 6-4), you can see that your user has a large DASD at 0192, which is equal to the disk size that you requested in the Create New z/VM guest window.

*Example 6-4 Results from the QUERY DASD command*

```
query dasd
DASD 0190 3390 ZVM4RS R/O        214 CYL ON DASD  991B SUBCHANNEL = 0007
DASD 0192 3390 ZVM4U1 R/W       2884 CYL ON DASD  DB53 SUBCHANNEL = 000A
DASD 019D 3390 ZVM4RS R/O        292 CYL ON DASD  991B SUBCHANNEL = 0008
DASD 019E 3390 ZVM4RS R/O        500 CYL ON DASD  991B SUBCHANNEL = 0009
```

The 192 disk is unformatted because IBM Wave assumes that this disk is formatted as a Linux root file system through the IBM Wave Linux installation wizard. Therefore, it is impossible to access the 0192 disk as a CMS file mode. Example 6-5 shows the type of error that you receive if you attempt to access the disk.

*Example 6-5 Disk access error*

```
acc 0192 b
DMSACP112S B(192) device error
```

Despite these limitations, the CMS environment can still be used to edit manually configuration files through **XEDIT**. Utilities such as the TCP/IP utilities on TCPMAINT's 592 disk are also potentially accessible and usable.

### 6.1.5 Deactivating a z/VM user

To deactivate a z/VM user, complete the following steps:

1. Right-click the user's icon and select **Deactivate**, as shown in Figure 6-16. You can simultaneously adjective multiple z/VM users by selecting multiple icons, right-clicking, and then selecting **Deactivate**.

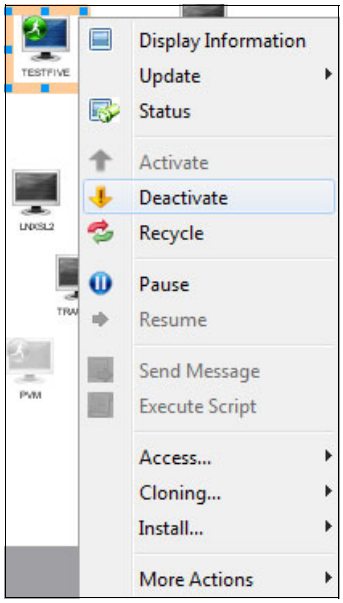


*Figure 6-16 Deactivate a z/VM user*

The Deactivate z/VM Users window opens, as shown in Figure 6-17 on page 185. Ensure that the users that you want to deactivate are selected. Additionally, if this is a Linux guest, you can run a **`shutdown -h`** command to the Linux guest's shell or force off the guest through a z/VM user logoff. A z/VM user logoff is the equivalent of powering off a machine, which is destructive to a Linux guest that was not shut down, which means that if you want to deactivate a z/VM user that is running Linux, you must run the Deactivate z/VM User task twice. It must first be run with the **Issue shutdown -h with timeout** option selected, and then a second time with the **Force z/VM User Logoff** option selected.

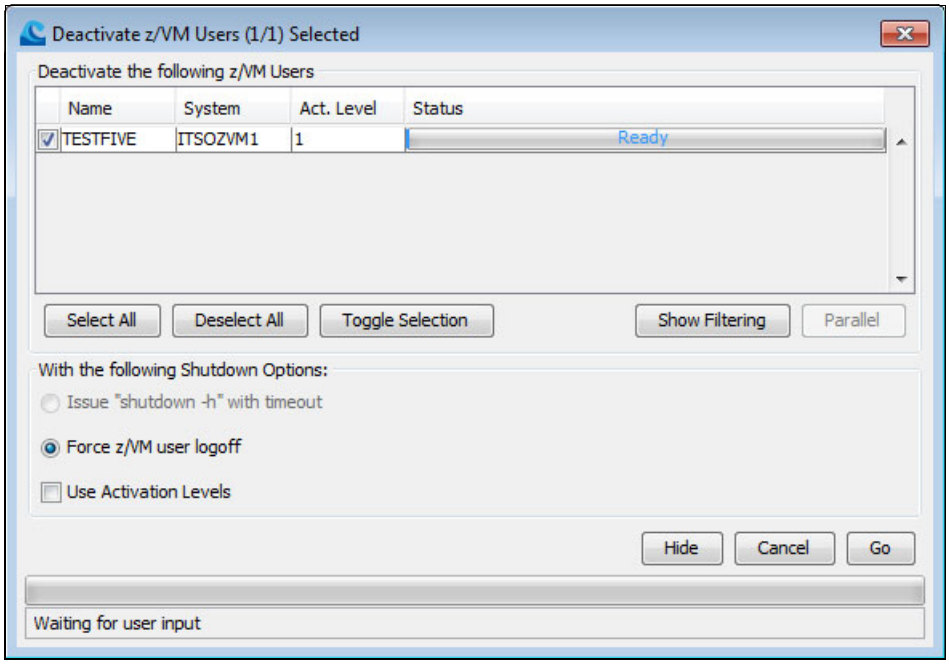


*Figure 6-17 Deactivate z/VM Users window*

2. After you select **Force z/VM User Logoff** and click **Go**, a window opens in which you are prompted to confirm that you want to force to deactivation, as shown in Figure 6-18.

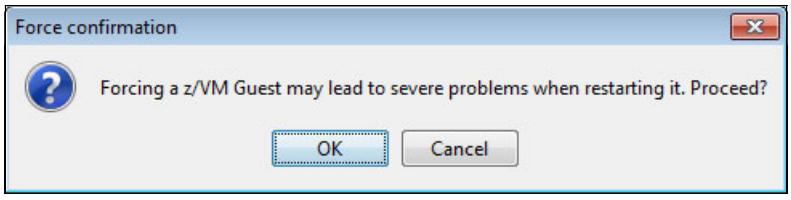


*Figure 6-18 Force confirmation*

## 6.1.6 Removing a z/VM user ID

There are instances in which you must remove a z/VM user ID. To do this, complete the following steps:

1. Right-click one or more user IDs.
2. Select **More Actions** → **Delete**, as shown in Figure 6-19.

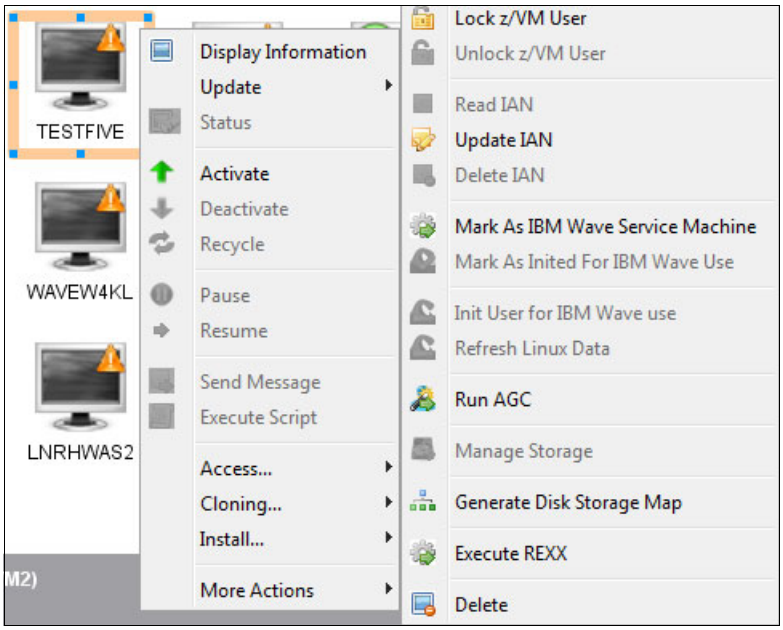


*Figure 6-19 Select Delete*

3. Confirm that each of user IDs that you want delete is selected, as shown in Figure 6-20.

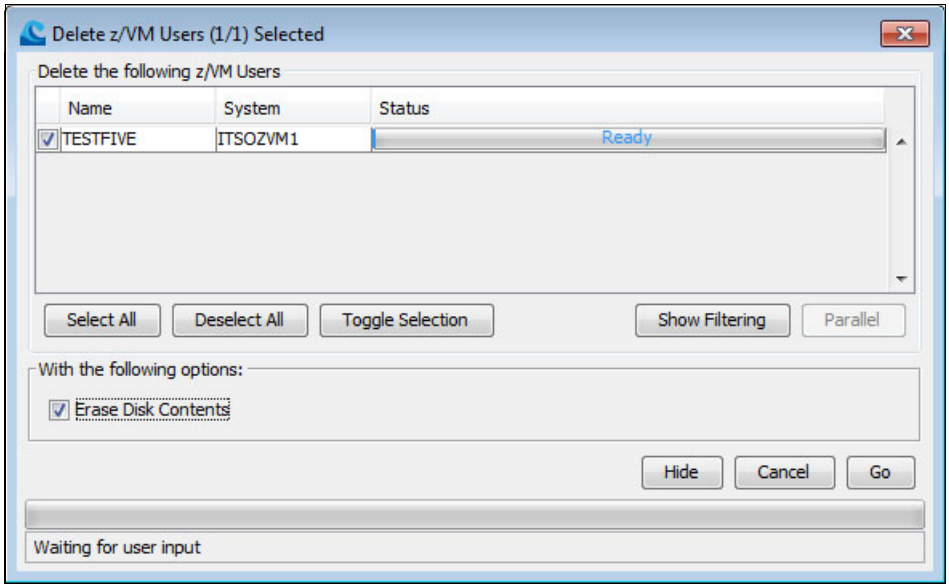


*Figure 6-20 Delete a z/VM user*

4. Select the **Erase Disk Contents** option. It is unknown whether this option only zeroes out the disks or whether it returns the disk to the storage group for reallocation.
5. Click **Go**. A confirmation window opens, as shown in Figure 6-21 on page 187.

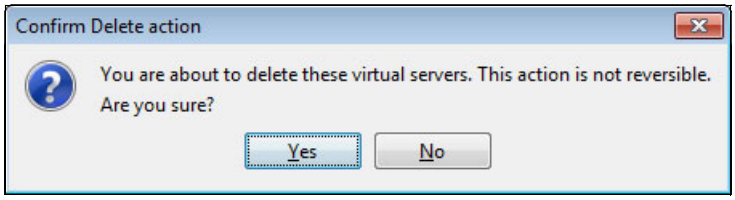


*Figure 6-21 Confirm Delete action window*

6. Click **Yes** if you are sure that you want to proceed. You receive a message that shows that your request was submitted to the BTS.

   If you want to see the progress of the work unit, open the BTW Work Units tab and double-click **Delete z/VM guests**, as shown in Figure 6-22.

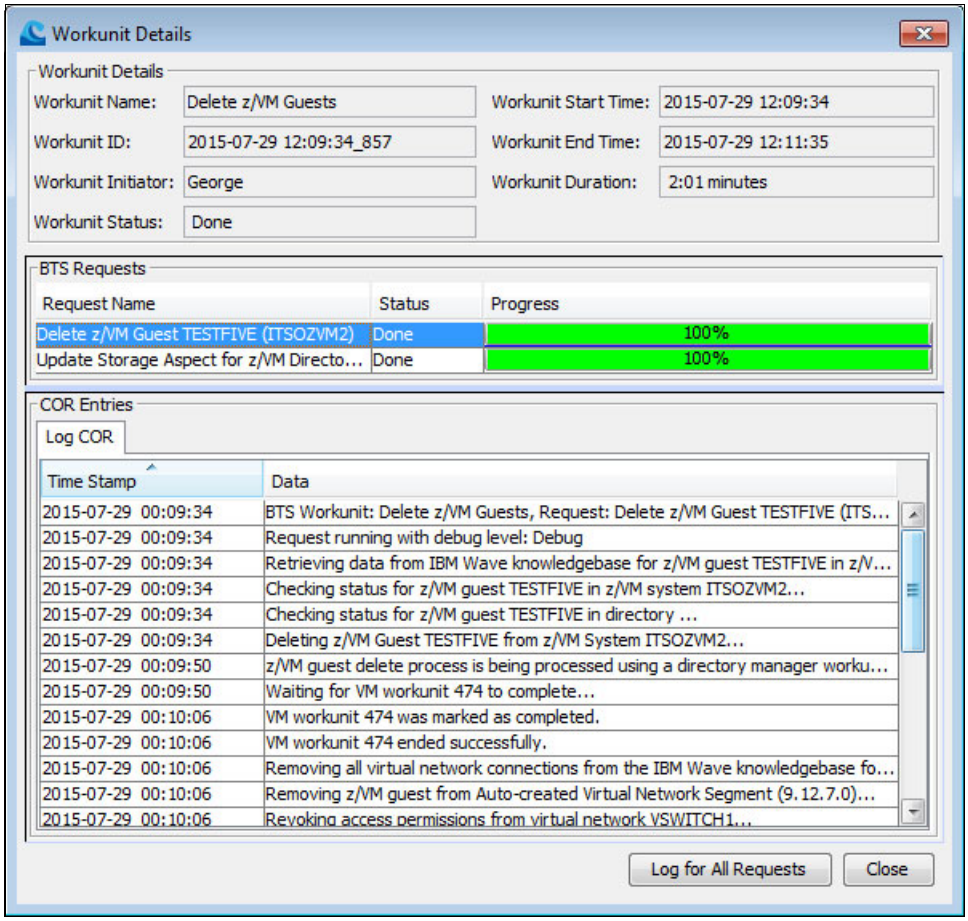


*Figure 6-22 Workunit details for delete z/VM guests task*

When the task is complete, the z/VM user icon disappears from its Site Defined Group.

## 6.2 Managing Device Pools

By using IBM Wave, you can create, modify, and update Device Pools. It is also possible to transfer Real Devices from one Device Pool to another.

This is an important function of IBM Wave, and you should pay attention when doing changes. The DASD, OSA, and HiperSockets resources are discovered by IBM Wave, which creates the Devices Pool grouping by system, type and CHPID.

The Manage Devices Pool function displays their resources in an organized manner, and is useful when you need to allocate those resources by IBM Wave.

To open the Manage Device Pools window, click **Administrative** → **Site Management** → **Manage Device Pools**, as shown in Figure 6-23.

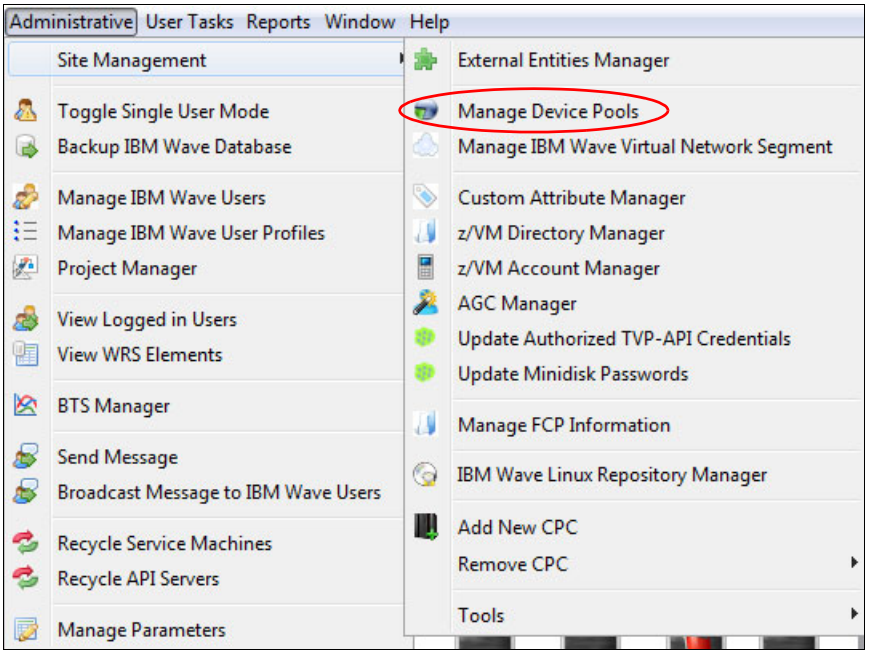


*Figure 6-23 Manage Device Pools menu option*

A window opens in which all of the Device Pools that are accessible to the current IBM Wave user are listed, as shown in Figure 6-24 on page 189. Only Device Pools that are in the IBM Wave user's scope are shown here. The window is a display-only version of the Device Pool Manager, which is accessible from the Administrative menu.

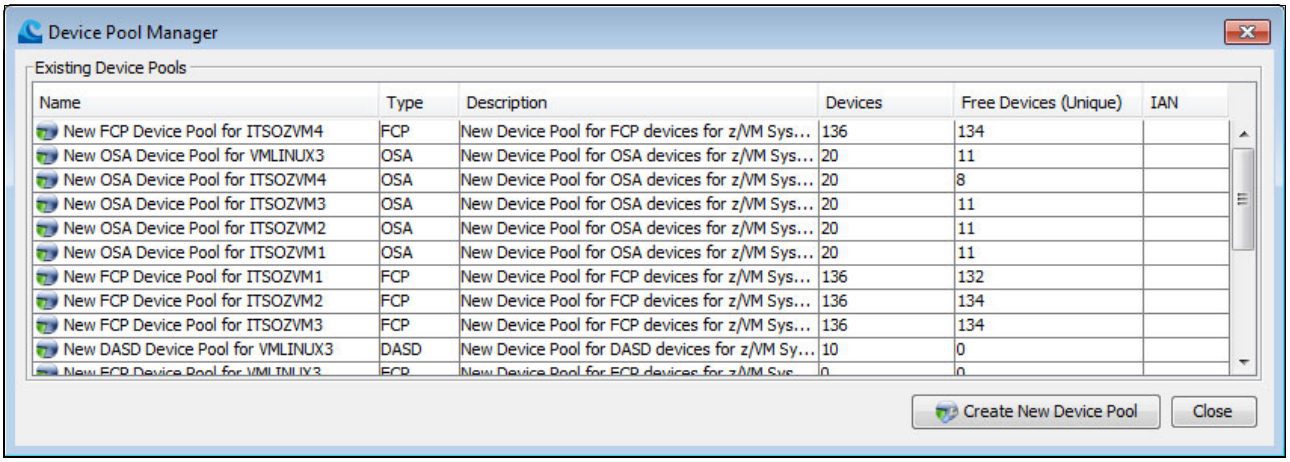


*Figure 6-24 Device Pool manager*

The Device Pools Manager uses the following device-class pools:

- DASD (FCP and IBM ECKD™)
- HiperSockets
- OSAs

It can manage multiple pools for each device class (Test, PROD, SYS, and so on).

Looking at the Device Pool Manager, you get to the Device Pools, which list information about each device in the pool, as shown in Figure 6-25.

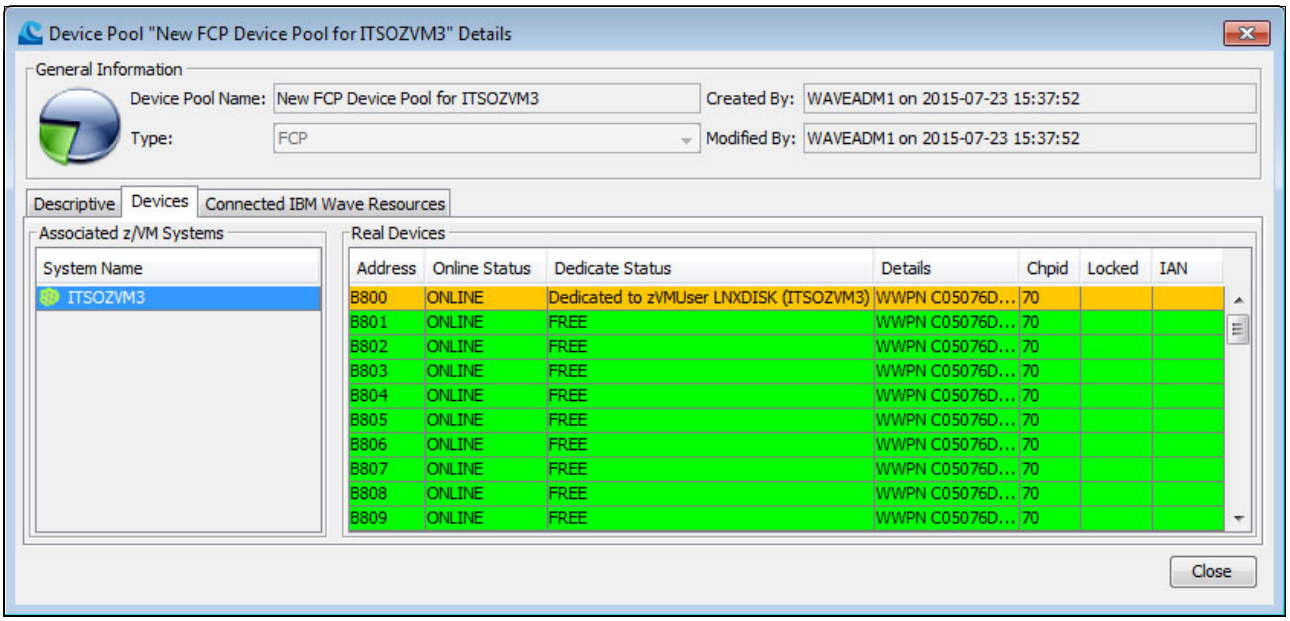


*Figure 6-25 Device Pool manager details*

The following fields are included at the top of the window:

- Name: Name of the Device Pool.
- Type: Type of the Device Pool.
- Created By/Last Modified By: The IBM Wave user who created or last modified this Device Pool and a time stamp.

The following information is available under the Descriptive tab:

- Default Virtual Device: An optional device address (four hexadecimal digits) that specifies the default virtual device to use when you are dedicating real devices from this device pool to z/VM guests.
- A Description: An optional description for the Device Pool.

In the Devices tab, you see all the devices that are owned by the Device Pool and the following information:

- System Name: This table displays all z/VM systems that are associated with the Device Pool. When a z/VM system is selected, the Real Devices table populates with Real Devices from the selected z/VM system that is owned by the Device Pool.
- Real Devices: This table lists the Real Devices and their System-Addresses that belong to the selected z/VM system and are owned by the Device Pool. Each line in the table provides information about one Real Device, including its address, online status, and dedicated status.

The next tab is the Connected IBM Wave Resources tab. This tab displays a table that lists all of the IBM Wave Managed Entities that are connected to the Device Pool.

## 6.2.1 Creating a Device Pool

To create a Device Pool, complete the following steps:

1. Open the Device Pools Manager by clicking **Administrative** → **Manage Device Pools** from the Site Management submenu.
2. Click **Create New Device Pool**, as shown in Figure 6-26.

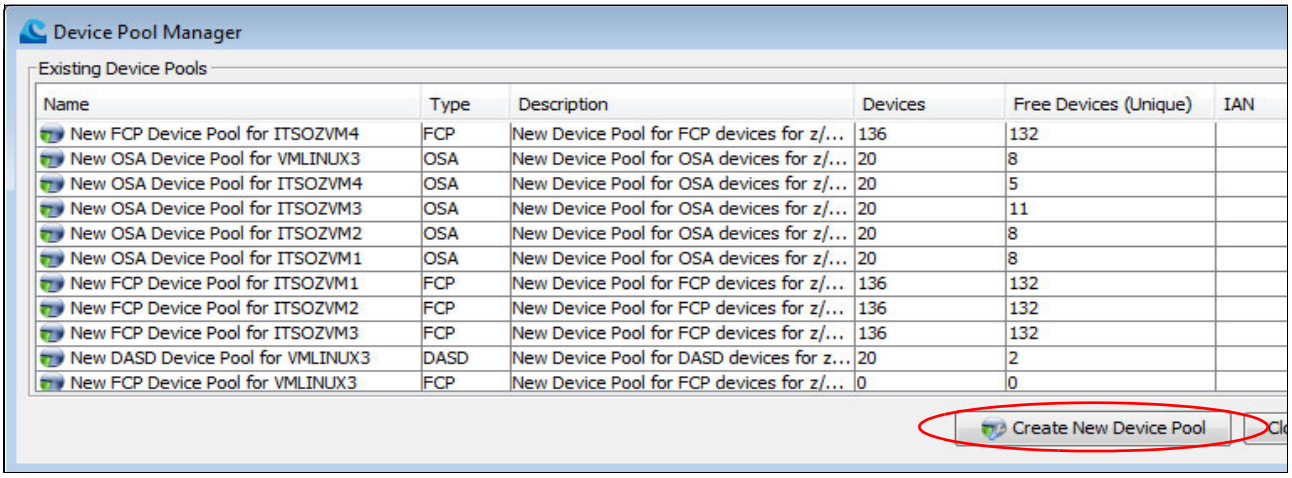


*Figure 6-26 Click Create New Device Pool*

3. In the Create New Device Pool window, enter a name for the Device Pool and select the type of Device Pool (for example, OSA). A description is optional. Optionally, you can specify a Default Virtual Device, as shown in Figure 6-27. Click **Create**.

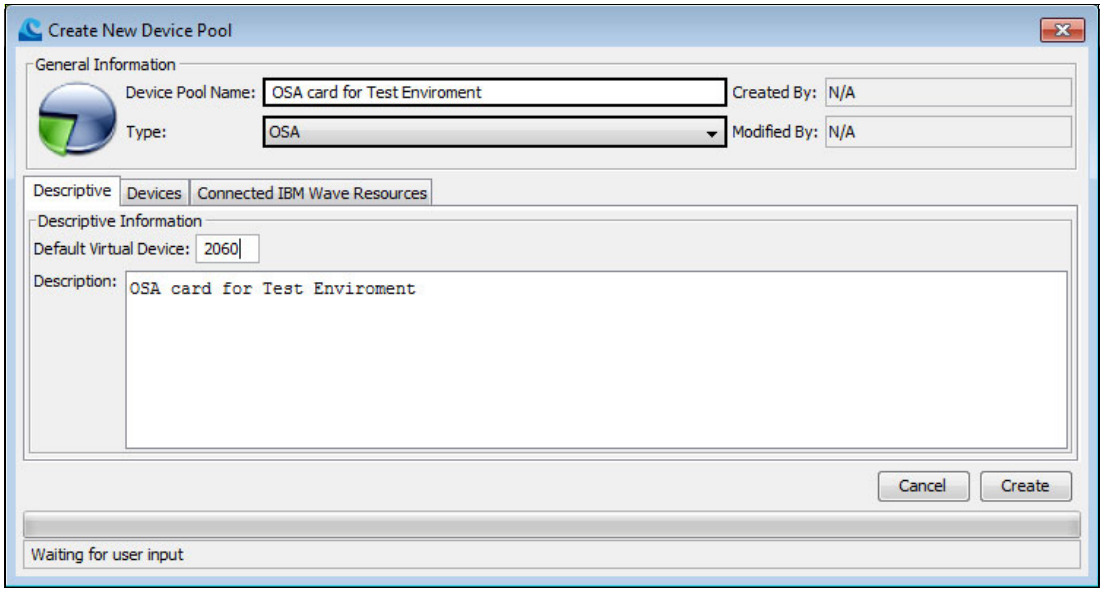


*Figure 6-27 Create New Device Pool window*

4. A window opens in which it is indicated that the work unit was dispatched to the BTS. You can then close the Create New Device Pool window. In a clockwise fashion, Figure 6-28 shows a composite of the work unit being submitted to the BTS, the work unit details, and the Device Pool manager with the new Device Pool.

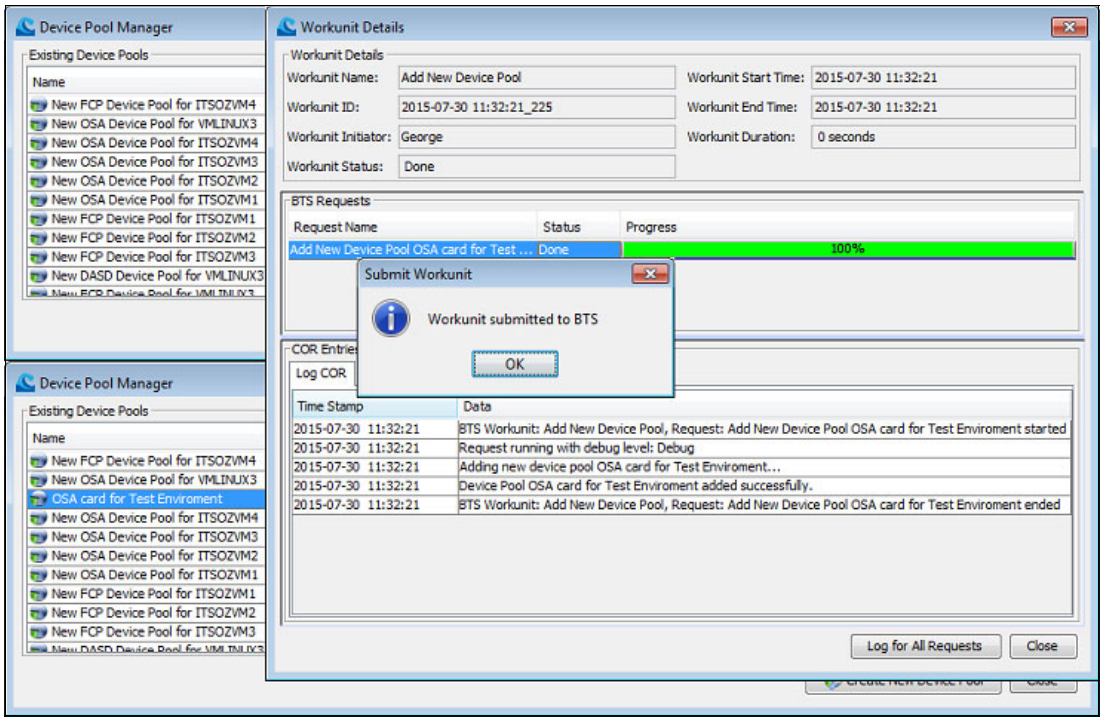


*Figure 6-28 Device Pool creation process*

5. Now that you created an empty Device Pool, you must associate it with a z/VM system before you can transfer devices from one pool to another. A Device Pool can be associated with several z/VM systems at the same time. To do this task, right-click the newly created pool entry and select **Update Details**, as shown in Figure 6-29.

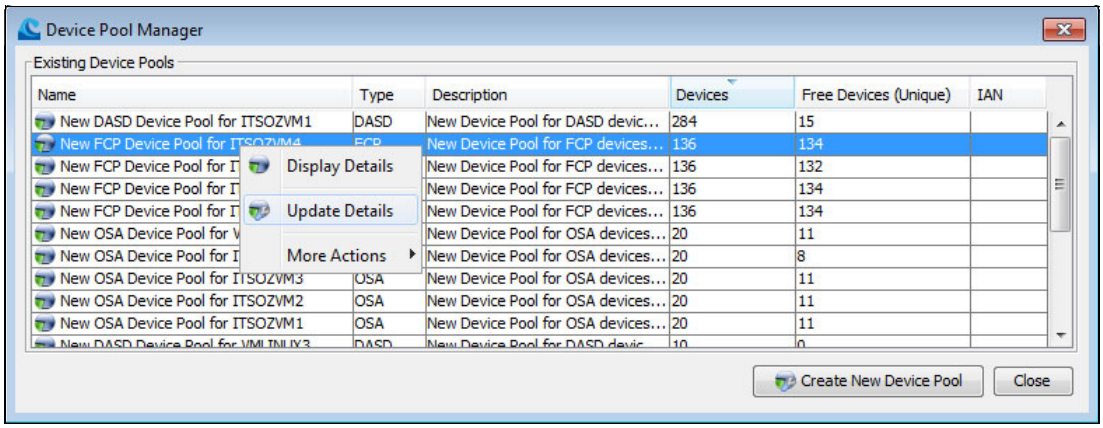


*Figure 6-29 Select Update details option*

6. A window opens in which you can update the selected Device Pool, as shown in Figure 6-30. Select the **Devices** tab, and then click **Associate System**.

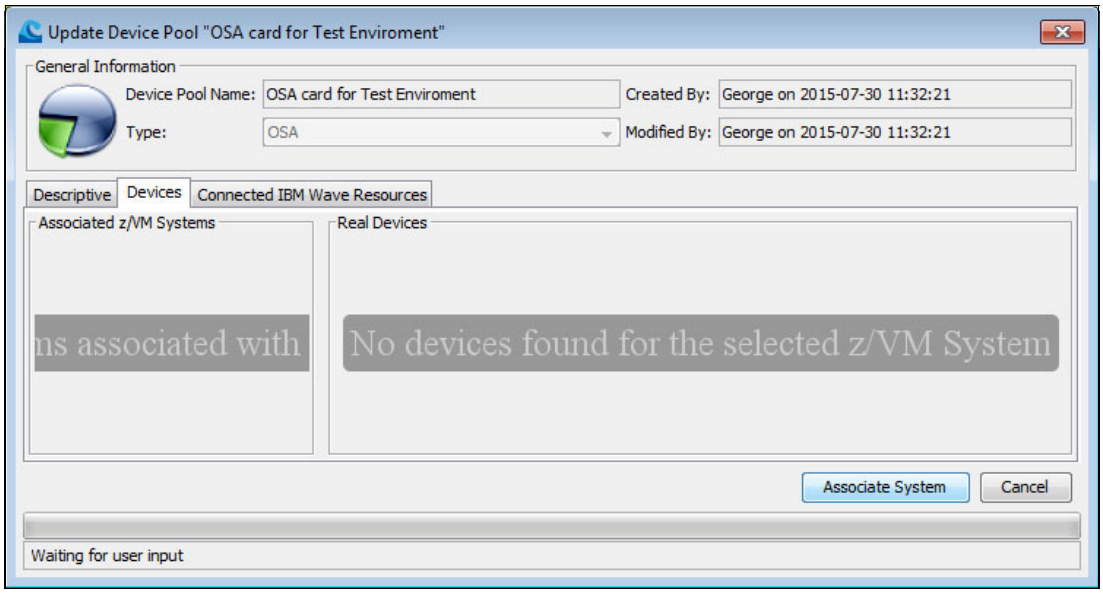


*Figure 6-30 Associate Device Pool to a z/VM system*

7. In the window that opens, select the z/VM system with which you want to associate, as shown in Figure 6-31. Click **Go**.

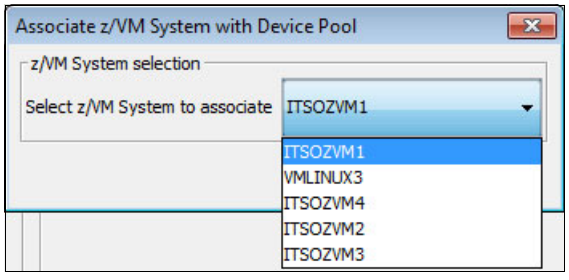


*Figure 6-31 Pick a z/VM system to associate*

8. A window opens that shows that the work unit was dispatched to the BTS. You also see the Associated z/VM systems, as shown in Figure 6-32. In this example, no devices were found yet for the selected z/VM system.

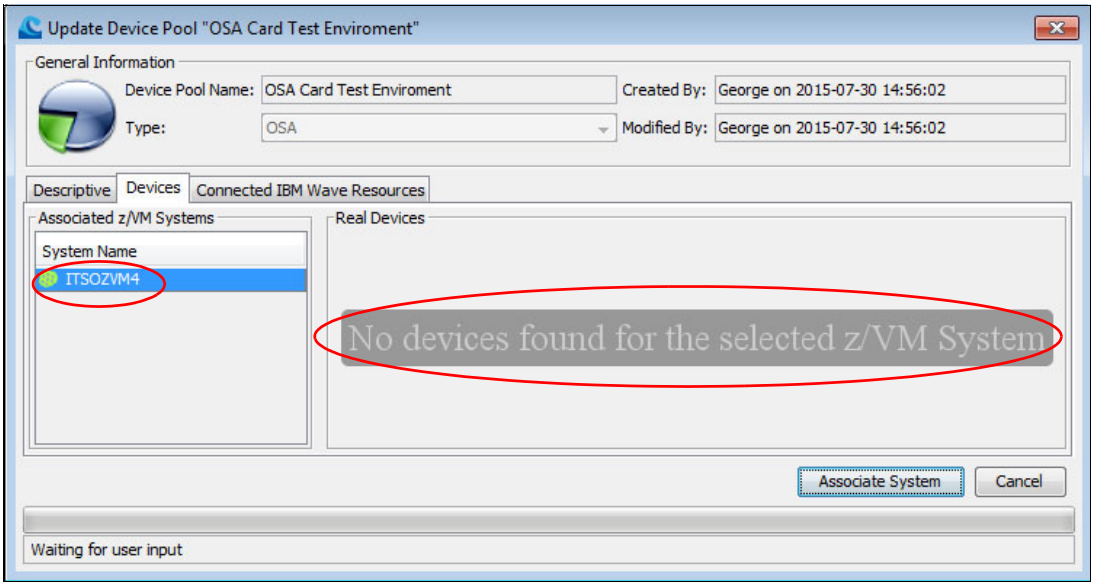


*Figure 6-32 Associated z/VM system*

## 6.2.2 Transferring Real Devices from one pool to another

To transfer Real Devices from one pool to another, the source pool must have some free devices (see Figure 6-33). These free devices must be available on the z/VM system with which the target Device Pool is associated.

To transfer Real Devices from one OSA pool to the newly created pool, complete the following steps:

1. In the Device Pool Manager window, right-click the source pool entry and select **Update Details**, as shown in Figure 6-33.

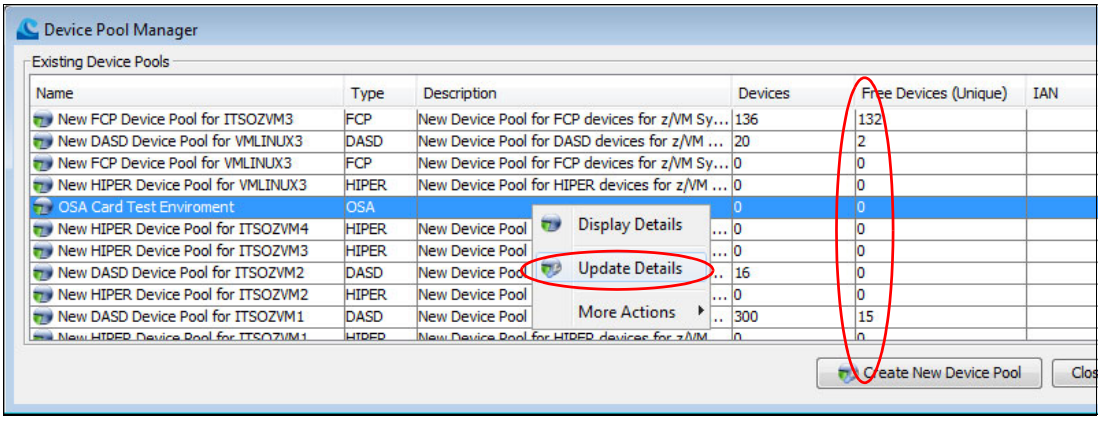


*Figure 6-33 Transfer Real Devices from one pool to another*

2. Select the **Devices** tab, and then select the Real Devices that you want to transfer (in this case, 2009 - 200B) by clicking the devices that you want, holding the Ctrl key, and clicking other devices. Right-click and select **Transfer To** and then the pool that you want to transfer to from the submenu, as shown in Figure 6-34.

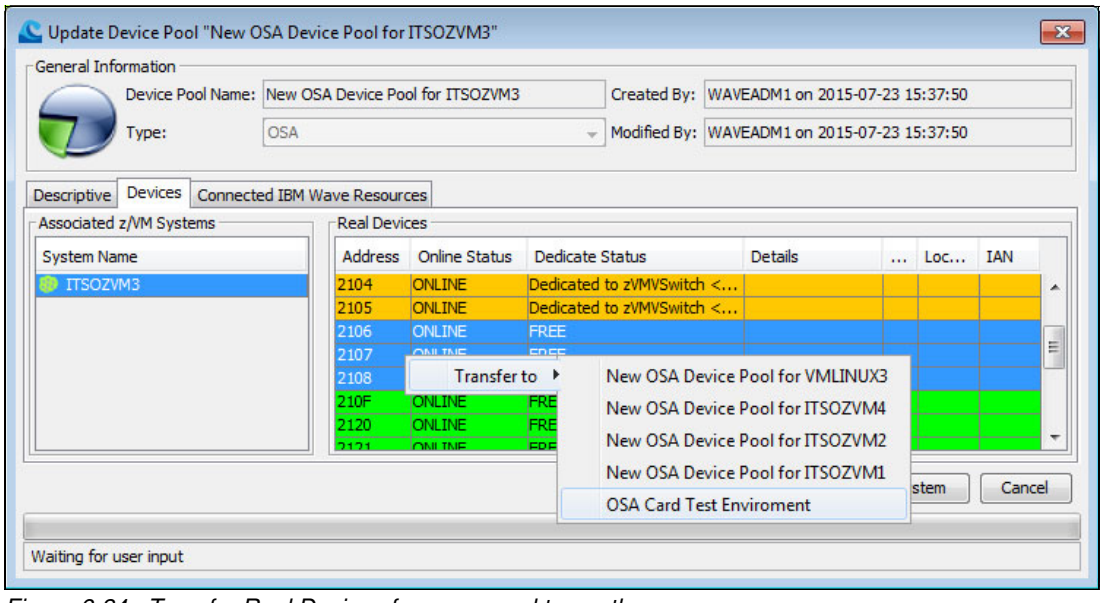


*Figure 6-34 Transfer Real Devices from one pool to another*

3. After clicking the device, you see a window in which it is stated that the work unit was dispatched to the BTS. Close the window and you notice that the Real Devices (2106 - 2109) are removed from the source pool. If you open the target pool, you see that the devices were successfully transferred there.

## 6.2.3 Removing a Device Pool

To remove a Device Pool, transfer all of the devices it includes to another Device Pool by following the process that is described 6.2.2, “Transferring Real Devices from one pool to another” on page 194. Transfer devices back until there are no devices that are found for the selected z/VM system.

If there is a 0 (zero) in the Devices column, there is no Real Device that belongs to this pool. You can delete it by right-clicking the Device Pool’s name and selecting **More Actions** → **Remove Device Pool**, as shown in Figure 6-35.

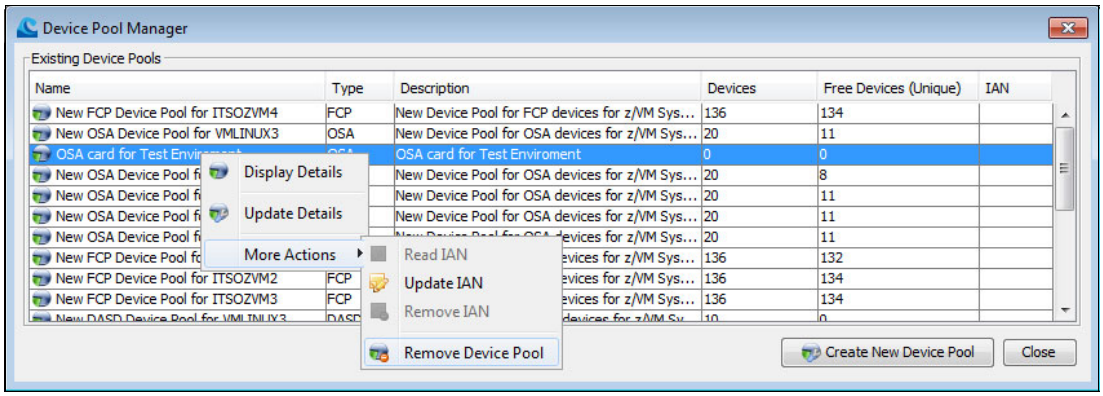


*Figure 6-35 Remove a Device Pool*

You see a window in which you are prompted to confirm that you want to remove the specific Device Pool. Click **Yes** and your request is submitted to the BTS and your Device Pool is removed.

Consider the following points about IBM Wave and Device Pools:

- IBM Wave uses internal z/VM data to retrieve a unique ID for each Real Device. This unique ID is made up of the Logical Control Unit ID with other elements and is a unique identification for each device.
- Attention Required status

  If any real device was dedicated to z/VM guest on two or more z/VM systems, IBM Wave raises an Attention Required status for the real devices. If this situation is acceptable, the Attention Required Entry can be ignored.
- Transfer Real Devices

  Because Device Pools are associated with Real Devices by their Unique ID, all Real Devices with same unique ID are transferred to the target Device Pool.
- It is possible to link a Managed Entity to a Device Pool. Managed entities are pure metadata. For example, it is possible to link an OSA Device Pool to a Router Managed Entity. The Network Viewer displays a link between the Device Pool and the router.

# 6.3 Managing DASD

This section describes how to manage DASD by using IBM Wave.

## 6.3.1 The Storage tab

When you click the **Storage** tab, you see three tabs on the left. The Distribution tab shows the whole space in z/VM systems that is managed by IBM Wave and your distribution. In the Groups tab, you can manage DirMaint groups. Finally, the Volume tab shows the volume disks and the many commands to manage your volumes.

## 6.3.2 Identifying the volume status

There are two DASD types that are usable by z/VM systems: ECKD and emulated device (EDEV). In the Storage tab, there is a specific symbol for each type, and the status of device. To understand what these symbols are, see Figure 6-36.

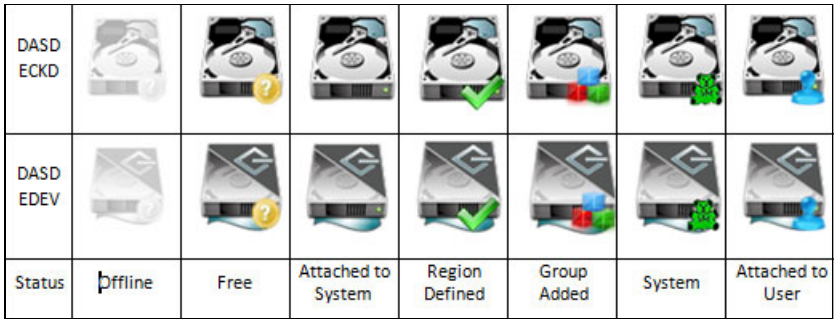


*Figure 6-36 DASD symbols by status*

## 6.3.3 Attaching DASD volumes to a system

To attach DASD volumes to a system, complete the following steps:

1. Open the Hardware Viewer and click the **Storage** tab, as shown in Figure 6-37 on page 197. In the left column of the Current System View, there are three options: Distribution, Groups, and Volumes. Click the **Volumes** tab.

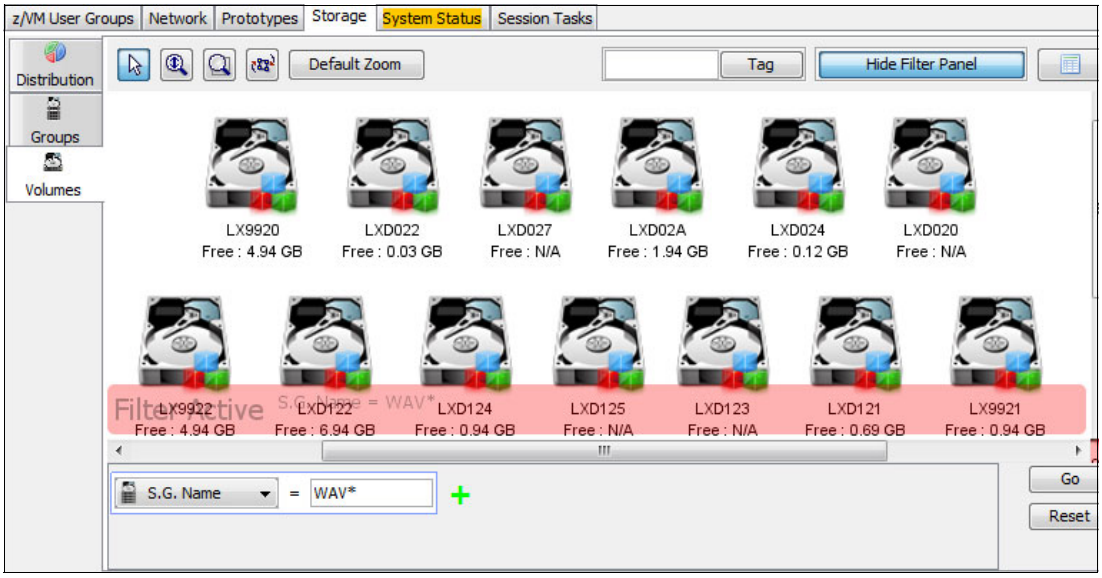


*Figure 6-37 Storage View by Volumes*

2. Depending on the number of DASD volumes that are accessible by your system, you might need to apply a filter to find the icons that represent the volumes that you want to add. In this case, it is easiest to sort the volumes by providing a range of real addresses, so click **Real Address**, as shown in Figure 6-38.

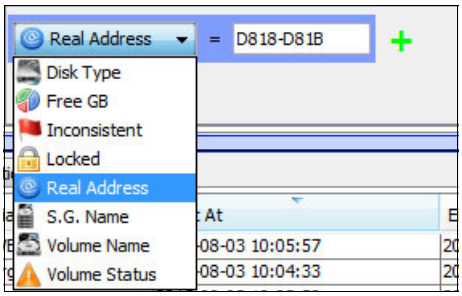


*Figure 6-38 Select a filter criteria*

After the filter criteria is applied, the number of volumes that is displayed decreases and the icons resize to fill the display, as shown in Figure 6-39.

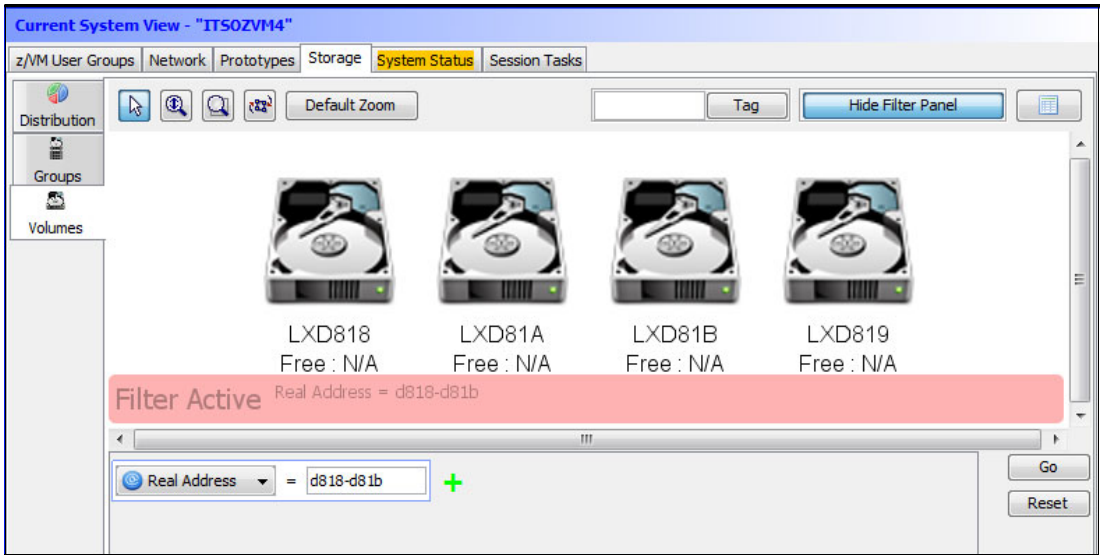


*Figure 6-39 Storage volumes that are filtered by real address*

3. After you display the specific volumes that you want to attach to the system, select one or more volumes, right-click them, and select **Attach to System**, as shown in Figure 6-40.

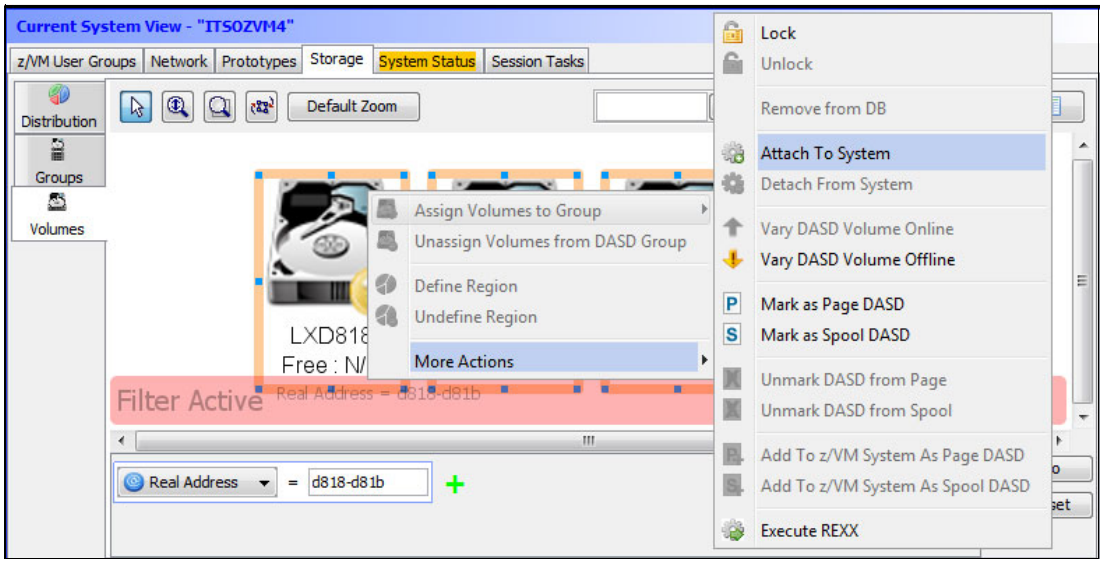


*Figure 6-40 Highlight storage volumes and select the Attach to System option*

4. A window that is called Attach DASD Volumes to z/VM system opens, as shown in Figure 6-41 on page 199. The volumes that were highlighted when you right-clicked and selected the option are listed in this window. By default, each of these volumes is highlighted. Click **Go** to start the task.

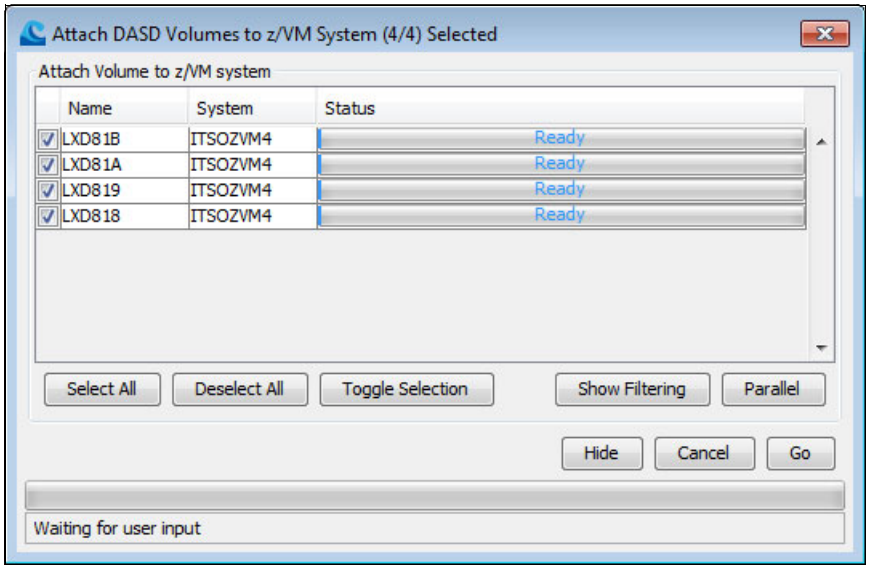


*Figure 6-41 Attach DASD Volumes to z/VM system window*

5. After clicking **Go**, you see a window that states that the work unit was dispatched to the BTS. You can then close the Attach DASD Volumes to z/VM system window.

**Note:** IBM Wave does not issue CP commands to any z/VM at once. Therefore, all CP commands must be performed at each z/VM if needed. If you want to attach a DASD for all z/VM systems in an SSI environment, you must run the command on each z/VM system and member of the SSI.

## 6.3.4 Including a volume for DirMaint

Now that the DASD volumes are attached to the system, you must define the volume in DirMaint extent control as a whole-volume region by completing the following steps:

1. In IBM Wave, select one or more volumes that are attached to the system, right-click the volume or volumes, and select **Define Region**, as shown in Figure 6-42.

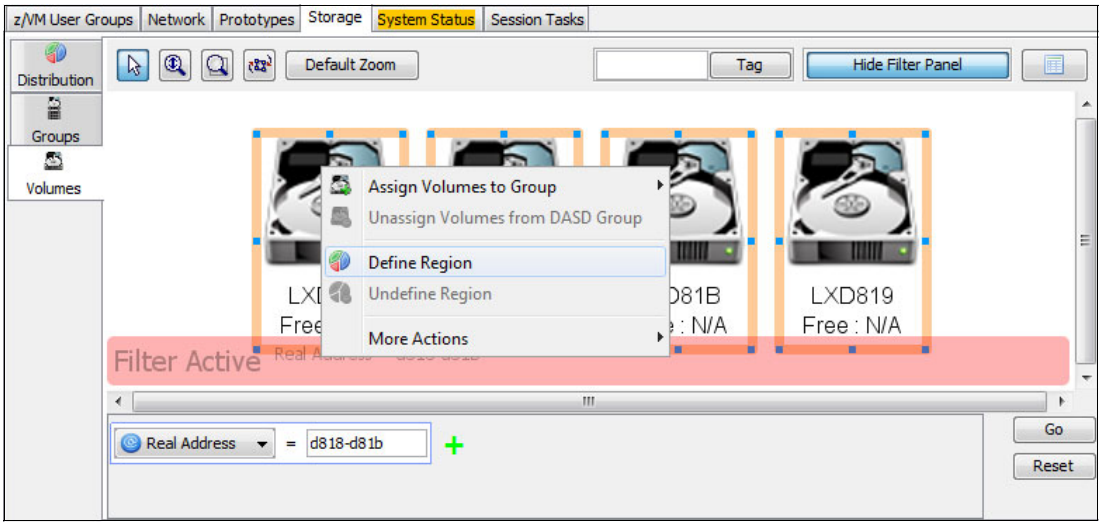


*Figure 6-42 Selecting Define Region from multiple attached storage volumes*

2. Click **Define Region**. The Define Regions to DASD Volumes window opens, which includes one or more attached volumes. By default, all listed volumes are selected, as shown in Figure 6-43. To define all of the listed volumes to DirMaint extent control as regions, click **Go**.

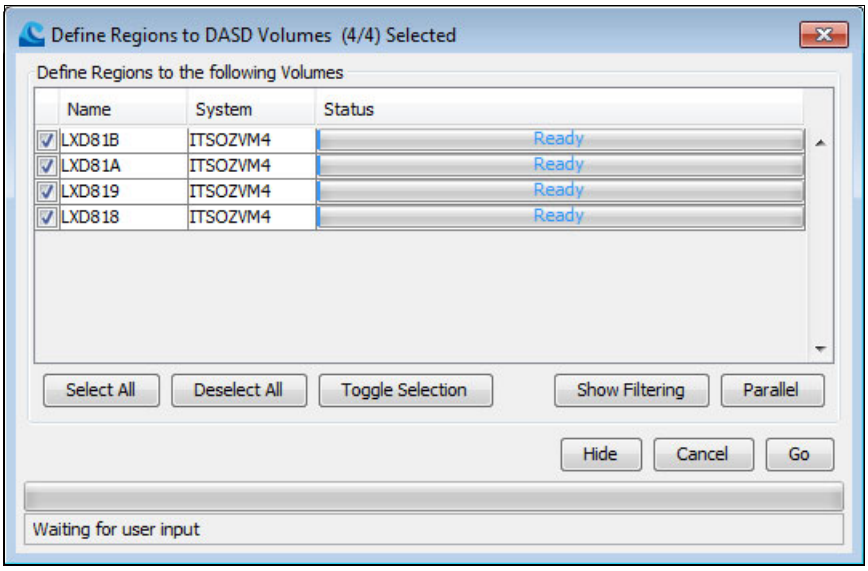


*Figure 6-43 Define Regions to DASD Volumes window*

If you are defining EDEV volumes, another window opens and prompts you for the type of device, as shown in Figure 6-44 on page 201.

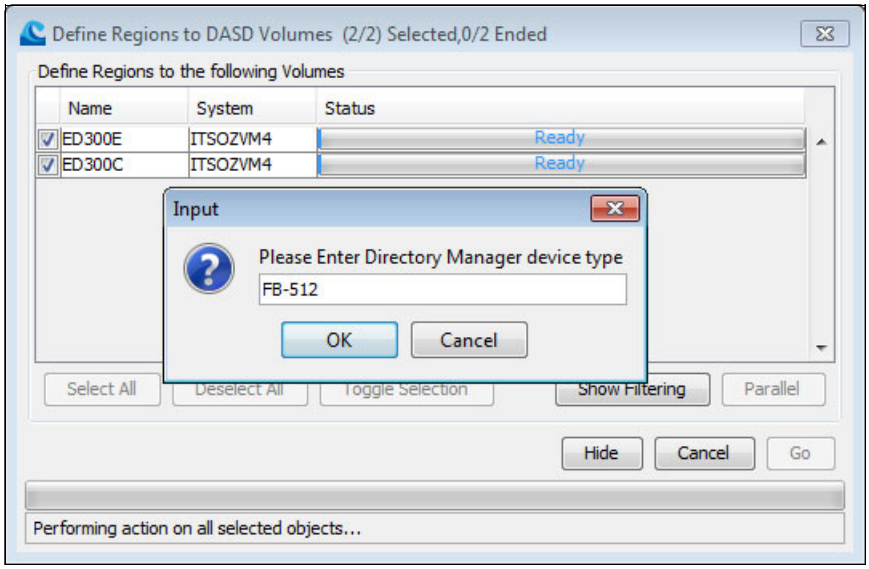


*Figure 6-44 Define Regions for EDEV Volumes window*

3. After the work unit is dispatched to the BTS, you can close the window.
4. After the attached volumes are defined as whole-volume regions to DirMaint extent control, they must be added to a DirMaint group. Select one or more volumes, hover your cursor over the Assign Volume to Group menu option, and then click one of the groups that appears in the menu, as shown in Figure 6-45.

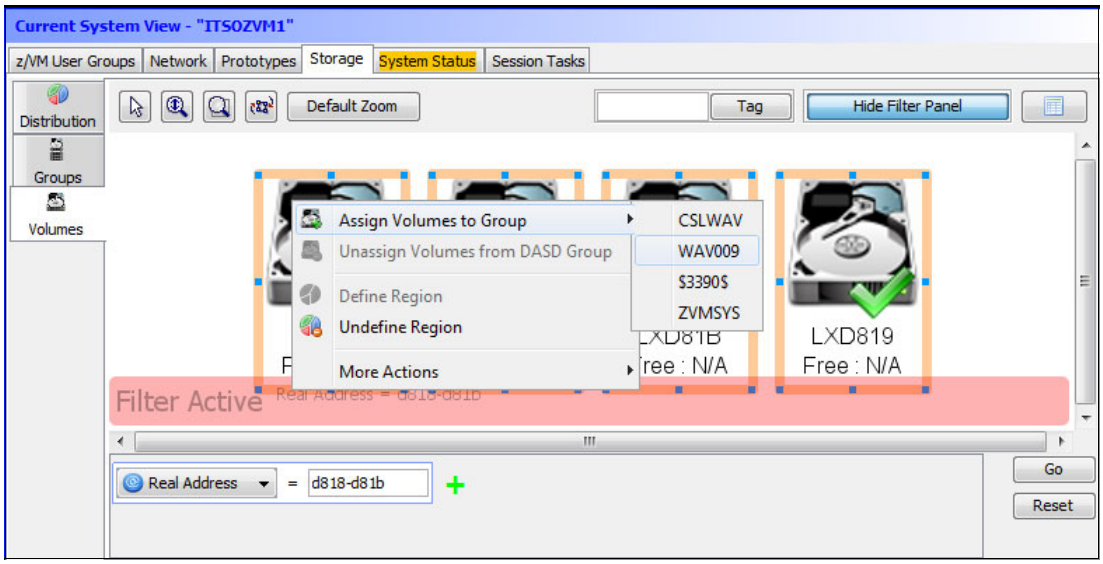


*Figure 6-45 Assign the volumes to a group*

5. Clicking a group name in the Assign Volumes to Group menu opens the Assign DASD Volumes to DASD Group window, which includes one or more attached volumes. By default, all of the listed volumes are selected, as shown in Figure 6-46. To assign all of the listed volumes to the selected DirMaint group, click **Go**.

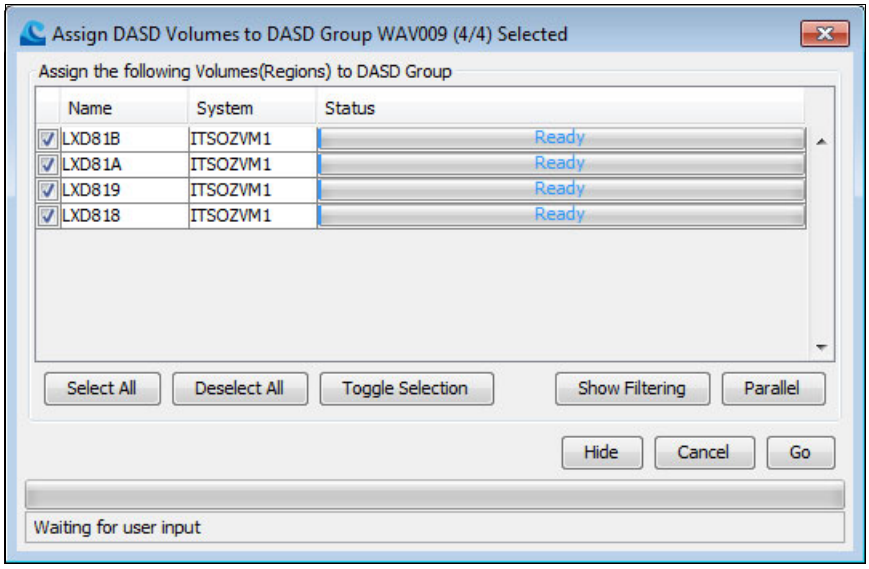


*Figure 6-46 Assign DASD volumes to a DASD group*

6. The work unit is dispatched to the BTS. Close the window. Each volume is fully configured and ready for use by IBM Wave, as shown in Figure 6-47.

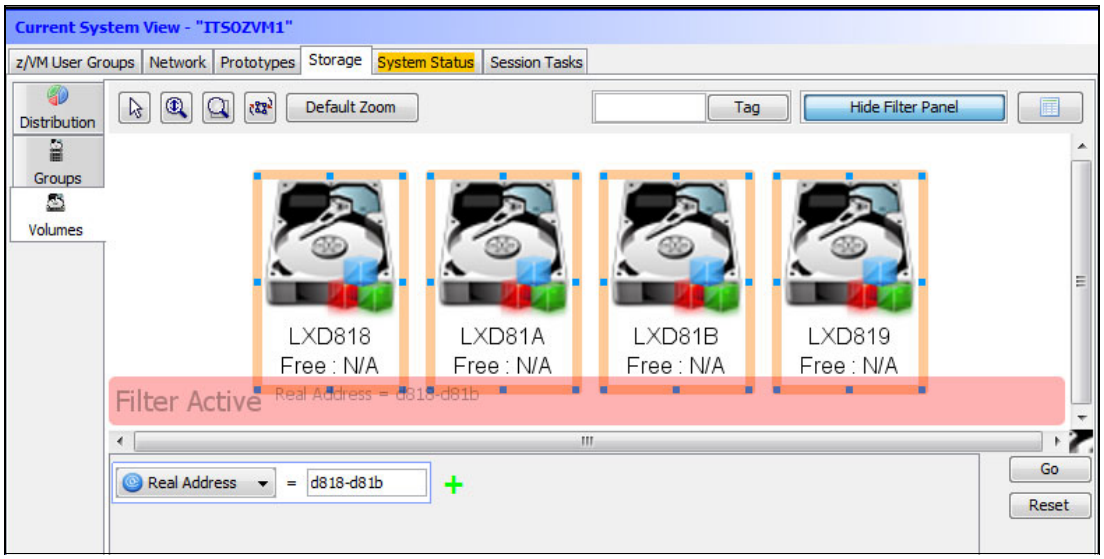


*Figure 6-47 New volumes are fully configured*

## 6.3.5 Creating a group

In the Groups tab, you can add a group to your SSI or stand-alone z/VM system. This action changes the `EXTENT CONTROL` file and the new group is available to insert new volumes.

To create a group, complete the following steps:

1. Click the **Groups** tab and right-click any empty space in the window, as shown in Figure 6-48.

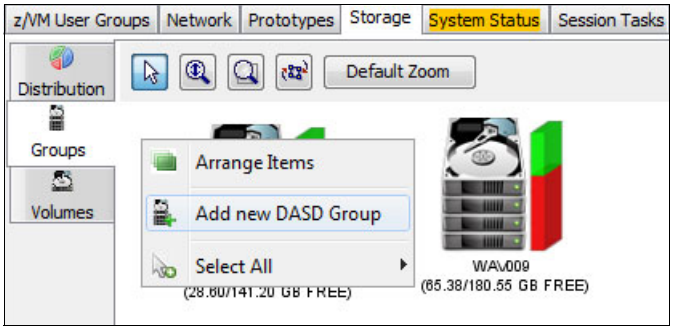


*Figure 6-48 Add a new DASD Group*

2. Complete the fields and click **Create**, as shown in Figure 6-49.

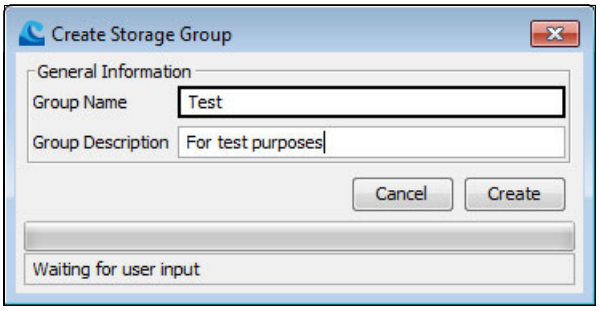


*Figure 6-49 Create Storage Group*

A work unit is submitted to BTS. After that process, completes, the new group is ready and you can add volumes there.

## 6.4 Managing the network

This section describes how to create a vSwitch, and after the process is complete, how to create a virtual network segment.

To see the network configuration, go to the Hardware Viewer display in the left pane of the IBM Wave window and highlight the z/VM LPAR with which you want to work. In the right pane, click the **Network** tab to see the layout of device pools, vSwitches, virtual network segments, and connected users.

## 6.4.1 Creating a vSwitch

To create a vSwitch, complete the following steps:

1. In the network topology window, right-click any white (unused) space and then click **Add New VSwitch**, as shown in Figure 6-50.

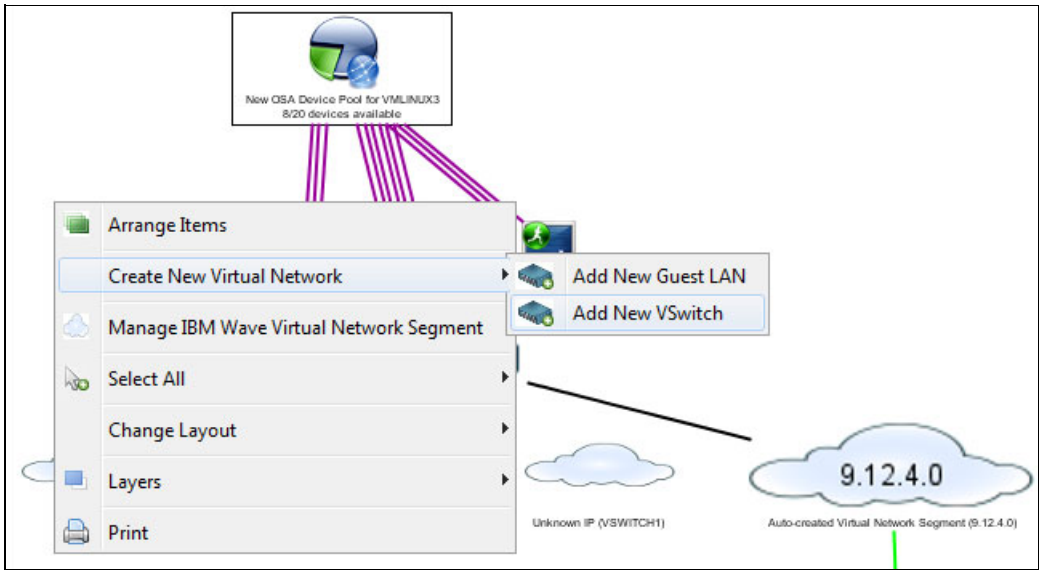


*Figure 6-50 Create a vSwitch*

2. In the General Information section (Figure 6-51), there are two required fields to complete: Name and Default NIC. In the Descriptive tab, you can enter comments about the virtual switch. In the Connected Segments tab, you do not need to complete the fields now. In the Vswitch Specific tab, you must select the Device Pool Name that has the three OSA devices to use in the virtual switch. You can optionally select the **VLAN Aware** option and enter a Default VLAN ID. The Layer type field default is IP (Layer 3), and you can also choose to use MAC (Layer 2).

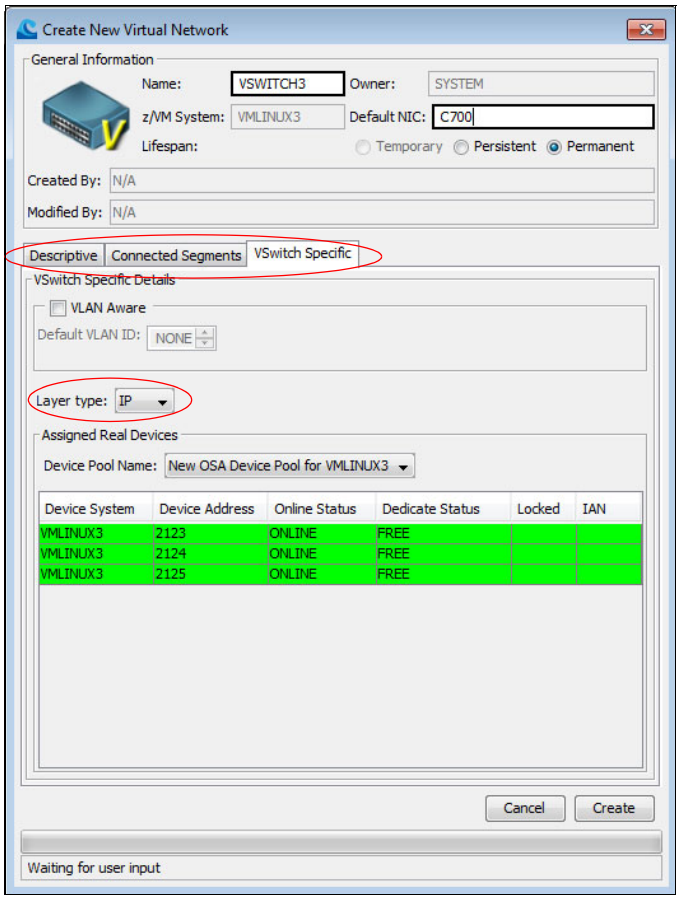


*Figure 6-51 Create New Virtual Network options*

3. After you click **Create**, a request is sent to the BTS and the new vSwitch is added to the network topology, in this case Vswitch3, as shown in Figure 6-52.

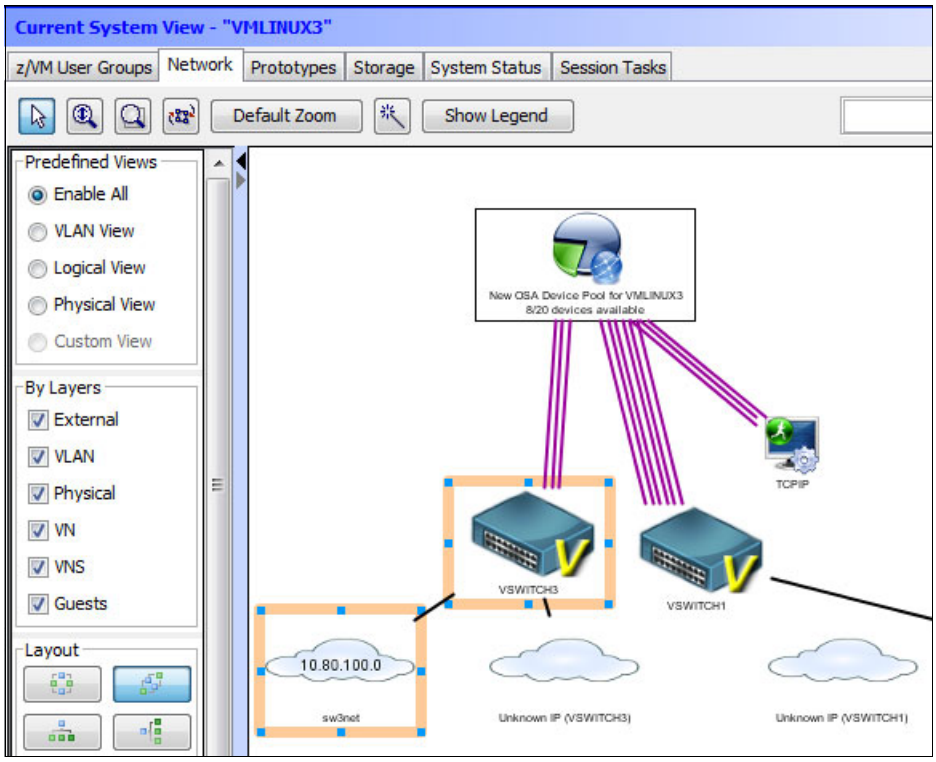


*Figure 6-52 Creating a vSwitch*

## 6.4.2 Creating a virtual network segment

Now that the vSwitch is created, you must “fill in the cloud” that is associated with it by adding a virtual network segment. Complete the following steps:

1. From the Administrative menu that is at the top of the window, click **Site Management** → **Manage IBM Wave for z/VM Virtual Network Segment**. Click **Create New Virtual Network Segment**, as shown in Figure 6-53 on page 207.

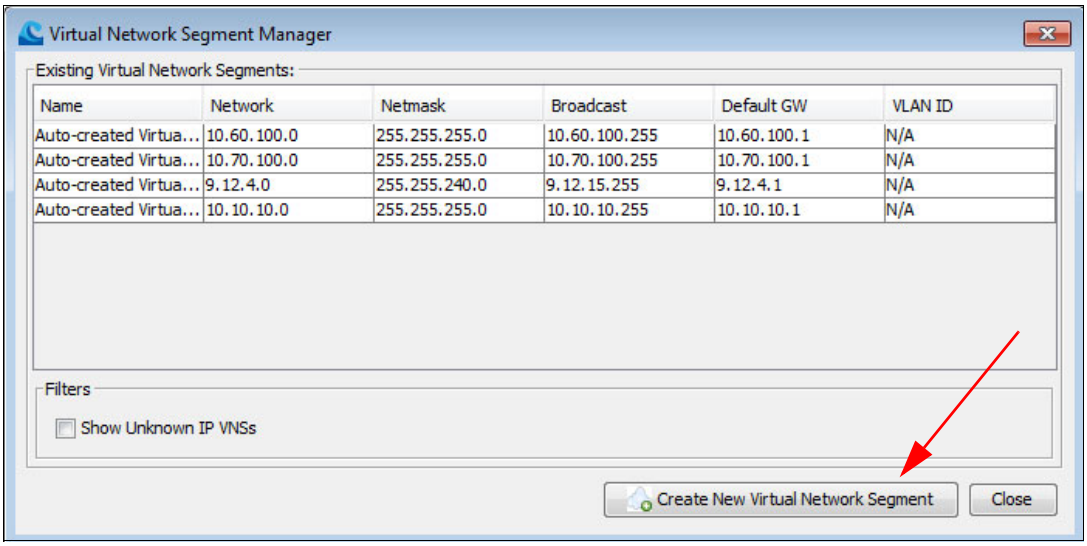


*Figure 6-53 Virtual Network Segment Manager*

A new window opens and you see information similar to what is shown in Figure 6-54.

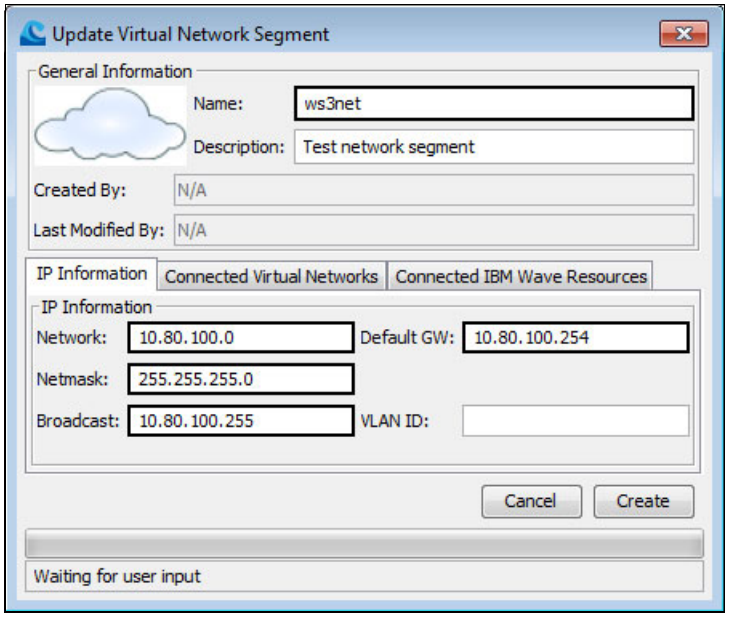


*Figure 6-54 Create a virtual network segment*

2. Enter the name of the Virtual Network Segment (VNS) in the Name field and the IP Information section.
3. Click **Create**.
4. After the segment is created, go to the Auto Detect menu that is at the top of the window and click **Refresh** → **Schedule z/VM Network Update**. A request is sent to the BTS.
5. When this segment is created, you connect it to a virtual network. Click **Administrative** → **Site Management** → **Manage IBM Wave for z/VM Virtual Network Segment**. Highlight and right-click the newly created network segment and select **Update Information**.

6. Click the **Connected Virtual Networks** tab and then click **Connect to Virtual Network**, as shown in Figure 6-55.

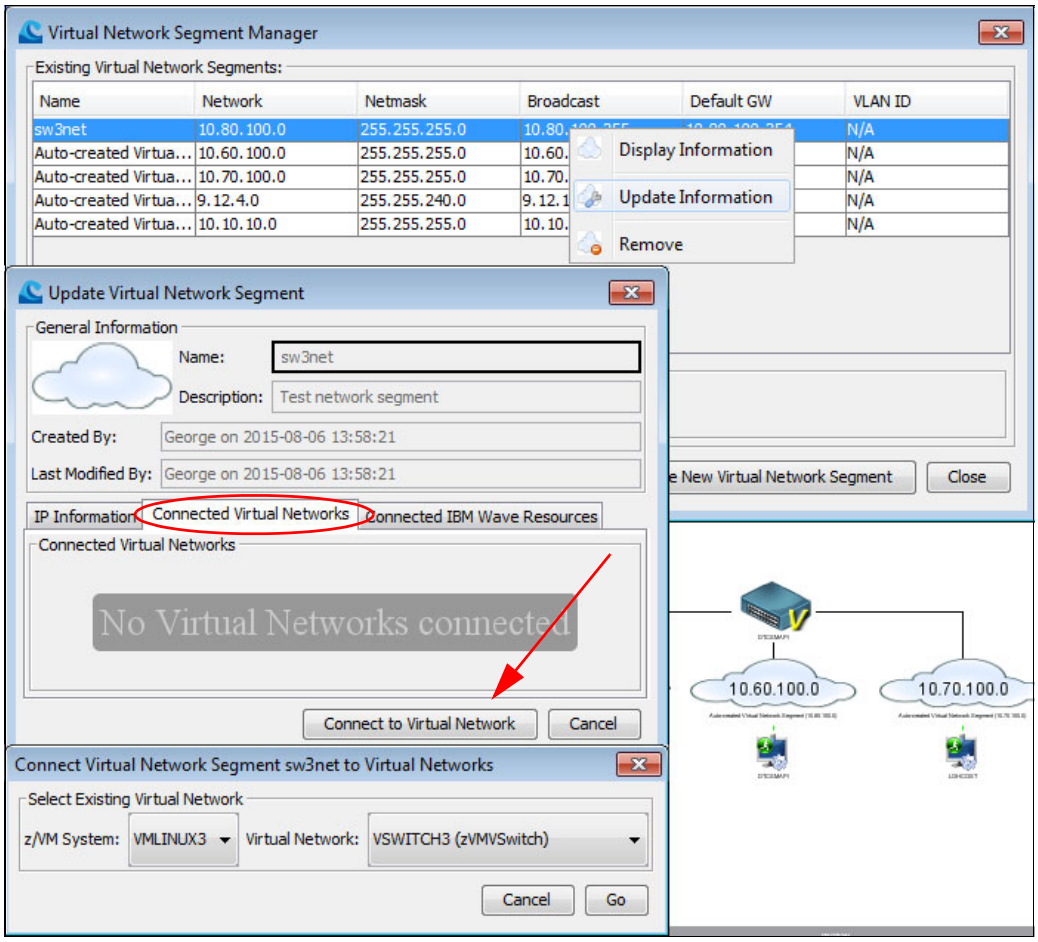


*Figure 6-55 Connect the virtual network segment to the virtual network*

7. Click **Go**. A work unit is sent to BTS and then you see that the network IP address that you supplied in the Network field in the Network Segment Manager window now appears in the VNS cloud, as shown in Figure 6-56.

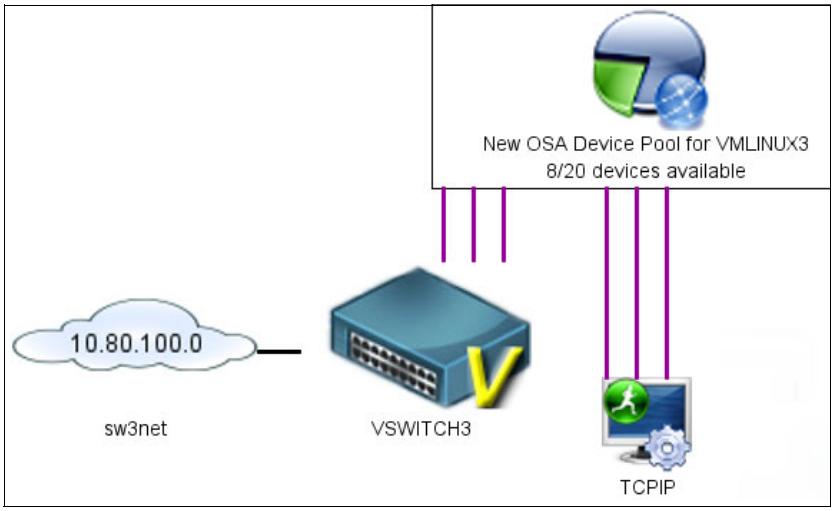


*Figure 6-56 VNS designation*

# 6.5 Managing z/VM SSI clusters

Managing z/VM in SSI clusters requires several additional steps in planning and installation. The additional considerations and definitions for SSI can be found in the following sections:

- “Preparation for SSI Installations” on page 16
- 4.1.7, “Single-system image definitions” on page 39

## 6.5.1 Live Guest Relocation

LGR is a feature in z/VM V6.3 that you can use to move z/VM Linux guests from one system in an SSI cluster to another system while the guest is active.

If you want to relocate one or many z/VM guests, you can use IBM Wave to do so by using dragging or by running a command.

### The dragging version

Here is an example of two Linux guests being moved from one system to another system by using the dragging method. Complete the following steps:

1. Switch to the Enterprise Viewer and group by z/VM systems, as shown in Figure 6-57.

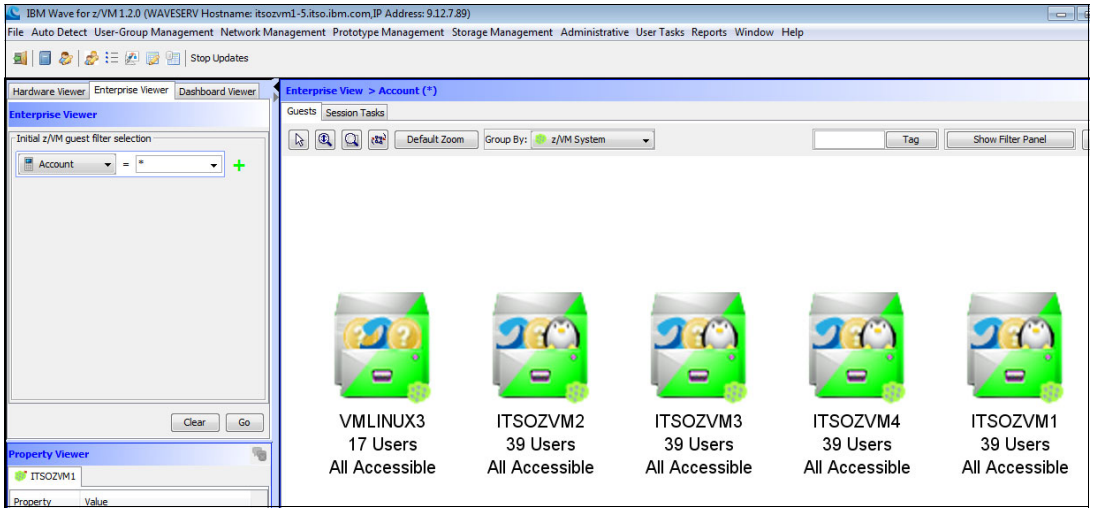


*Figure 6-57 Enterprise View - group by z/VM systems*

2. Set some filters (Figure 6-58) for the guests that you want to move. In a production environment, you might be dealing with several hundreds guests. Even in the test environment, you have 39 users, so set the filter to active users, which brings the number of users down to nine, which can be easily managed.

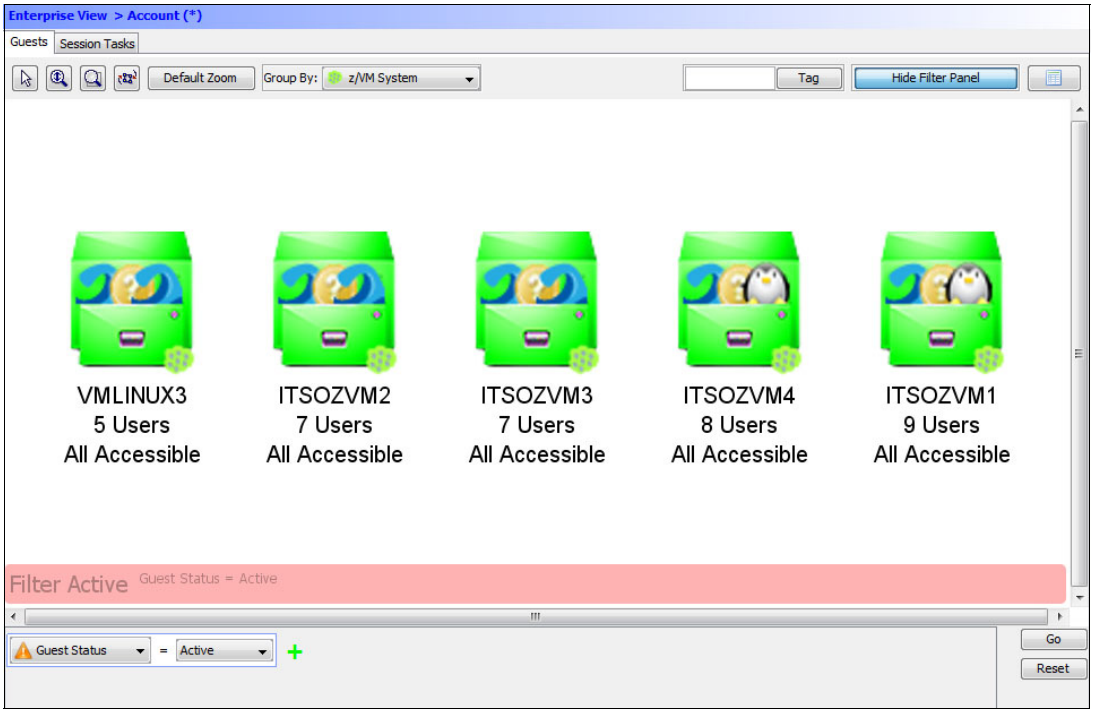


*Figure 6-58 Enterprise View - set filters*

3. Open the system from which you want to relocate the guests (Figure 6-59).

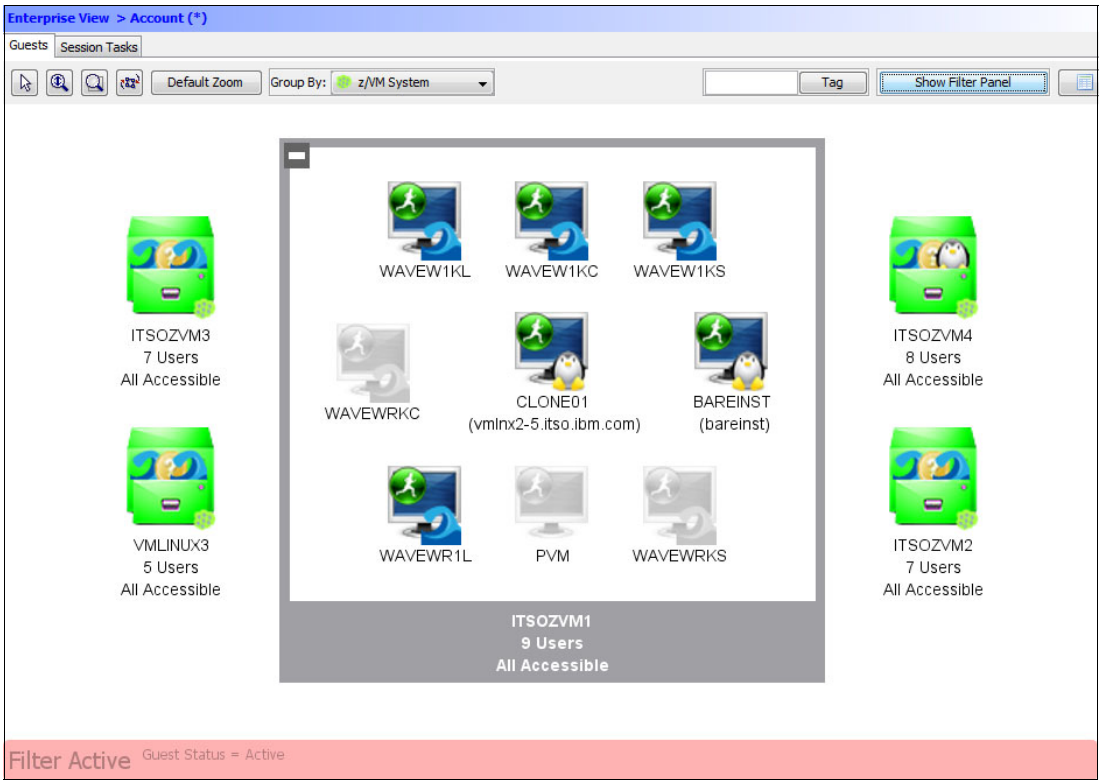


*Figure 6-59 Enterprise View - open system 1*

In this example, you want to move BAREINST and CLONE01 from ITSOZVM1 to ITSOZVM2. To select the guests, press the Shift key and click. After you select one or many guests to relocate, press the Ctrl key and keep holding it, then drag the selected guests to the system to which you want them to be relocated. In Figure 6-60, you first try to relocate to VMLINUX3, which is not a member of the SSI cluster. A red rectangle shows that this action is impossible and IBM Wave cannot do it.

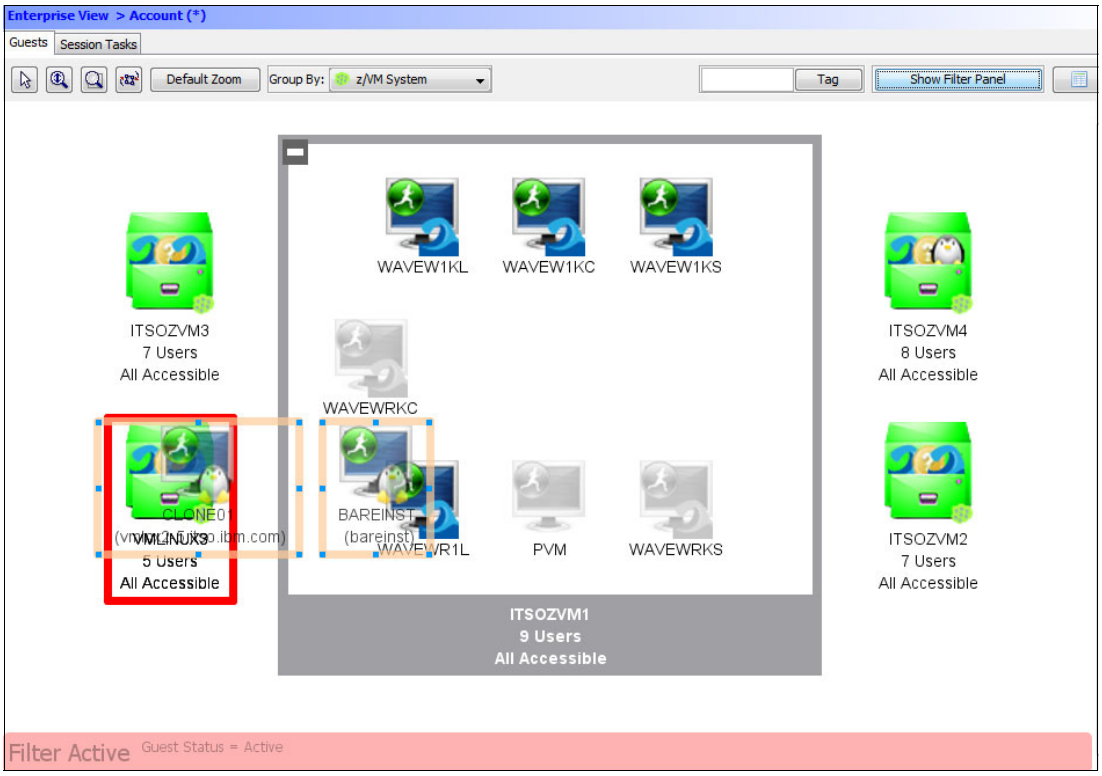


*Figure 6-60 Live Guest Relocation - attempt to move to a system that is not a member of the SSI*

4. When you select a valid member of the SSI cluster, the target system shows a green rectangle. Drop the guests onto the target system, as shown in Figure 6-61.

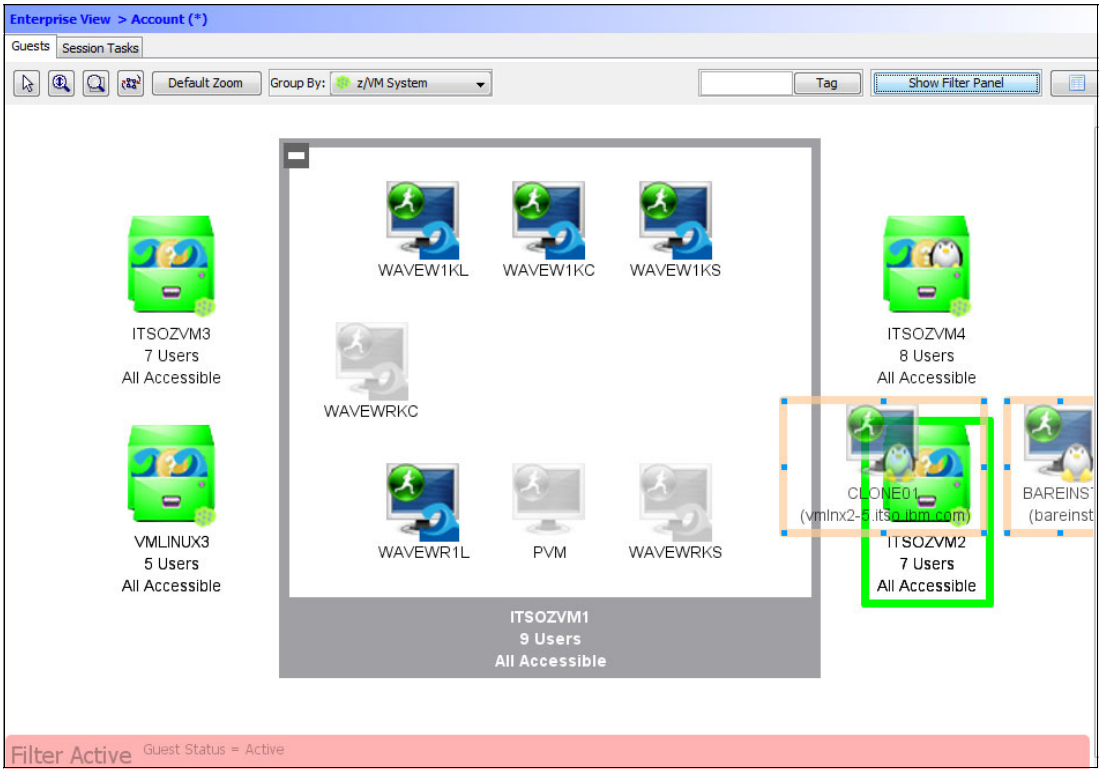


*Figure 6-61 Live Guest Relocation - attempt to move to a system that is a member of the SSI*

5. The z/VM LGR window opens, as shown in Figure 6-62. Click Go, which opens the window that is shown in Figure 6-63.

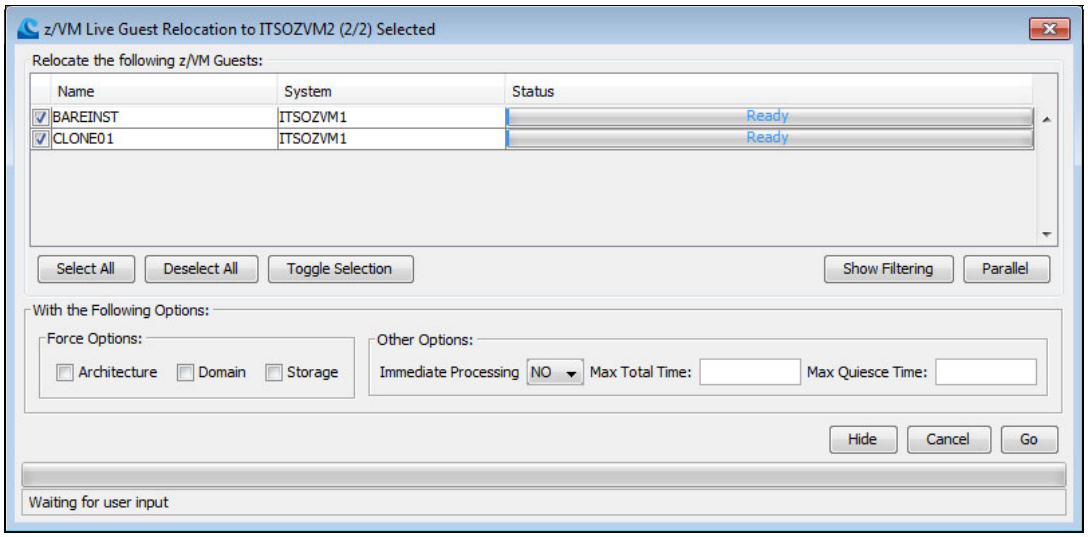


*Figure 6-62 Live Guest Relocation start*

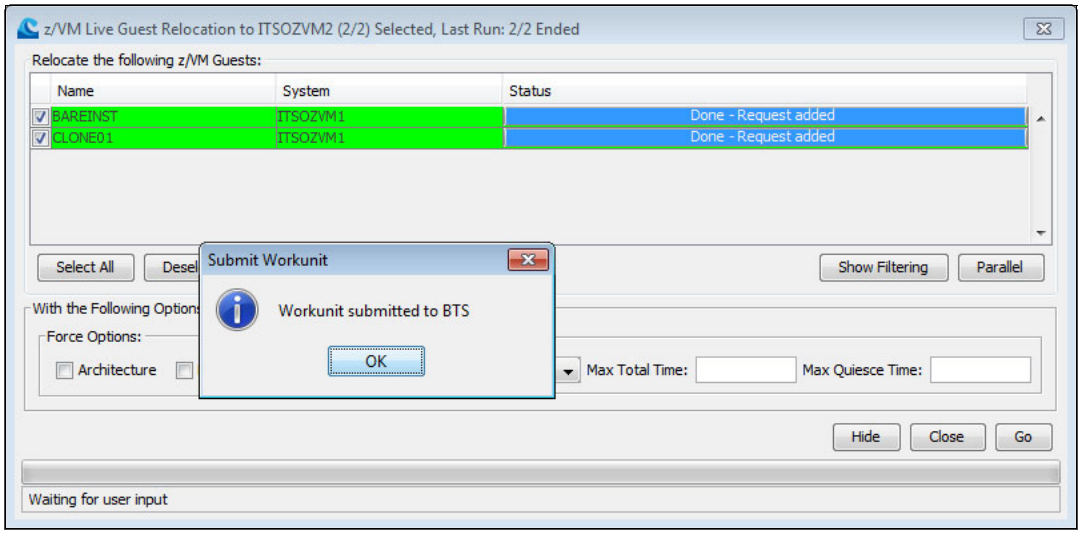


*Figure 6-63 Live Guest Relocation work that is submitted to BTS*

The work unit is submitted to BTS. IBM Wave starts to relocate the selected guests, as shown in Figure 6-64 on page 215.

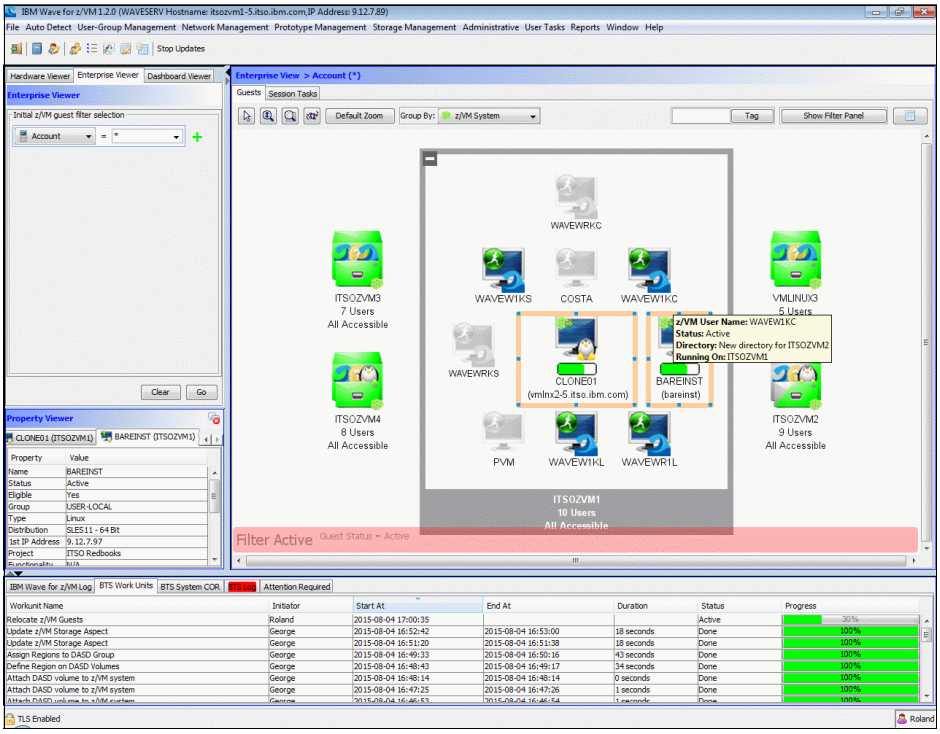


*Figure 6-64 Live Guest Relocation in progress*

You see a progress bar below the selected guests, as shown in Figure 6-65. You can also check the status in the BTS Workunits window, which shows the progress of “Relocate z/VM guests”. The two guests also appear on the target system with progress bars.

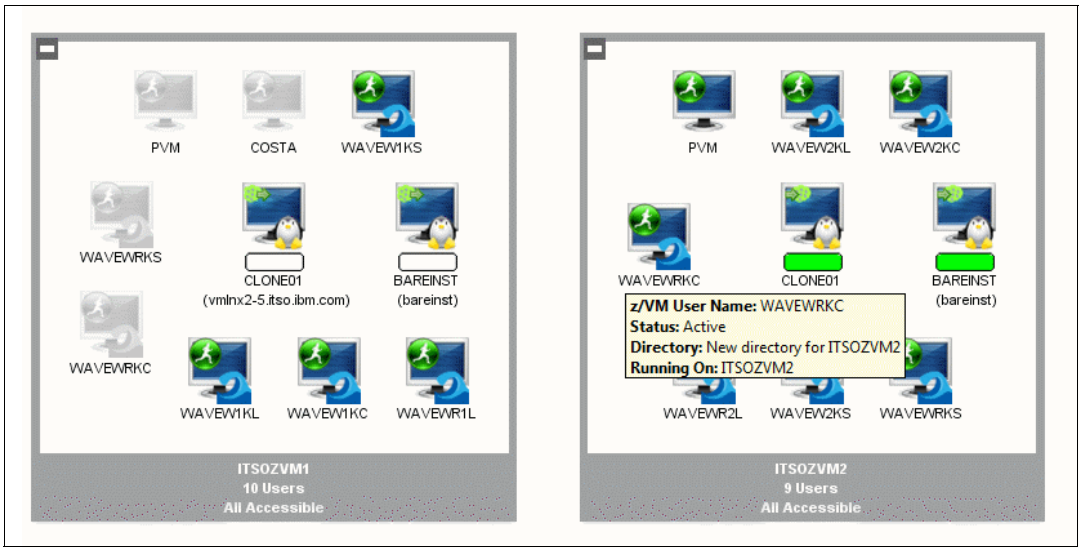


*Figure 6-65 Live Guest Relocation progress*

When the relocation is complete, the guests switch to Inactive on the source system. On the target system, IBM Wave might need a little time to detect the right network resources (red icon), as shown in Figure 6-66. The guests continue to work the whole time. In this example, the guests disappear on the source system because you set the filter to “active systems”.

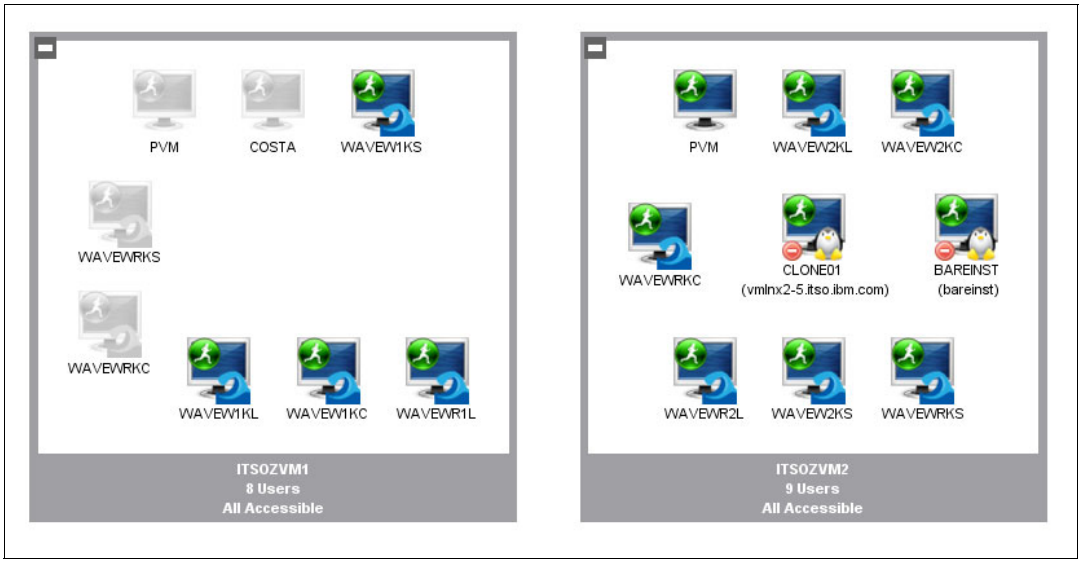


*Figure 6-66 Live Guest Relocation complete - IBM Wave network detection pending*

Finally, IBM Wave detects the network connections, as shown in Figure 6-67.

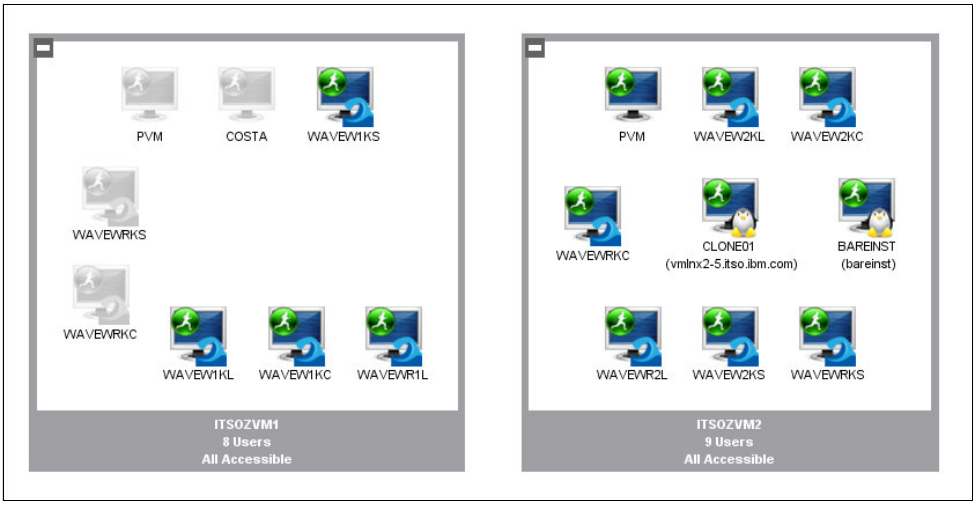


*Figure 6-67 Live Guest Relocation complete*

## The command version

Here is an example of two Linux guests being moved from one system to another system by running an IBM Wave option / command. Complete the following steps:

1. Select the guest that you want to relocate, right-click it, and select **More Actions** → **Relocate to**.

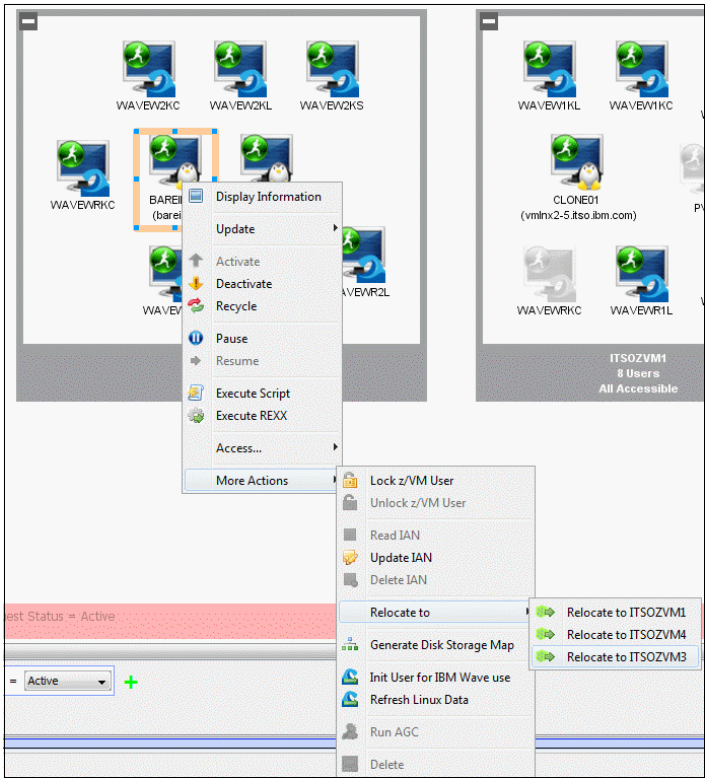


*Figure 6-68 Relocate to*

2. When the next window opens, complete step 5 on page 214.

7

# Managing Linux virtual machines

A golden master is a template for a virtual machine (VM). It can also be referred to as a *golden image*, a *master image*, or a *base image*. Usually, a golden image is manually installed and clones are created with minimal modifications. By using IBM Wave for z/VM (IBM Wave), a prototype can be created that automates cloning golden images.

This chapter describes how to use the IBM Wave interface to install Linux guests from scratch by using installation media that is mounted on a file server, convert Linux guests into golden master prototypes, and clone guests from a prototype.

This chapter builds on some of the core concepts that were described in Chapter 6, “Managing IBM z/VM” on page 169. Traditionally, these tasks were performed by a z/VM systems programmer. However, the bare metal installation and cloning features simplify z/VM administration to enable Linux administrators to provision and deprovision independently Linux guests on z/VM.

This chapter includes the following topics:

- Preparing to install Linux on the virtual machine
- Creating a z/VM guest
- Initializing guests by using IBM Wave
- Creating a golden master
- Cloning virtual machines from a golden master

# 7.1 Preparing to install Linux on the virtual machine

IBM Wave has two features that you can use with the GUI to create Linux guests from scratch. This function offers the possibility of building Linux golden images without the use of 3270 interfaces.

There are times when the system programmer uses traditional 3270 methods to build a golden image and then uses IBM Wave to turn this golden image into associated prototypes. This practice is acceptable.

There also are times when a bare metal installation is necessary. A bare metal installation can serve as a suitable means for building golden images when a z/VM systems programmer is not available. This section describes using a bare metal installation to build a golden image.

## 7.1.1 Installing an X server on your workstation

VNC is not supported by IBM Wave during a bare metal installation; therefore, you must install an X server on your workstation to tunnel X Window System through an SSH connection. This configuration enables the use of the fully graphical Linux installation wizard. There are many options available, including XMing, Cygwin with XFree86/X.org, and Hummingbird Exceed.

If you cannot install X11, you must install the pseudo-graphical Ncurses from YaST, as shown in Figure 7-1.

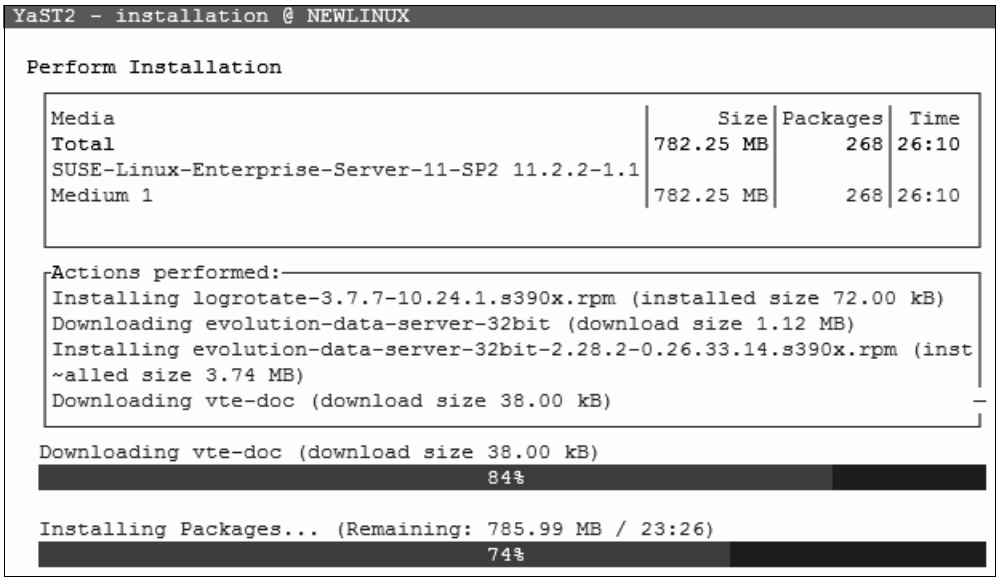


*Figure 7-1 Ncurses pseudo-graphical interface*

After accepting all of the defaults during the installation wizard, start your X server. With most X11 applications, you know that the X server started when the X icon in your system tray changes color from red to green or an X icon appears in your system tray.

While your X server is running, other applications, such as PuTTY, can access the X Window System services.

**Note:** If you are a Macintosh user, you can use XQuartz, which is available at the following website:

`http://xquartz.macosforge.org/landing/`

Run the `*.dmg` file and copy the application icon to your Application folder.

If you are a Linux user, you most likely already have an X server that is preinstalled on your system.

## 7.1.2 Logging in as the IBM Wave administrator

Log in as the IBM Wave administrator. If you are already logged in to IBM Wave, log out and log back in as the administrator.

## 7.1.3 Configuring an external SSH program to start sessions with X services

PuTTY is a no-charge implementation of Telnet and SSH for Windows and UNIX operating systems, along with an xterm terminal emulator. You can use any SSH terminal, but this section describes how to configure your SSH terminal by using PuTTY in IBM Wave.

To configure an external SSH program, complete the following steps:

1. In IBM Wave, open User Preferences by clicking the small taskbar icon in the upper left corner of your IBM Wave GUI that shows a person with a pencil, as shown in Figure 7-2.

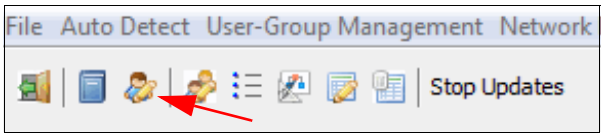


*Figure 7-2 User preferences*

A Change Preferences window opens in which you can set preferences that are specific to your IBM Wave user ID, as shown in Figure 7-3.

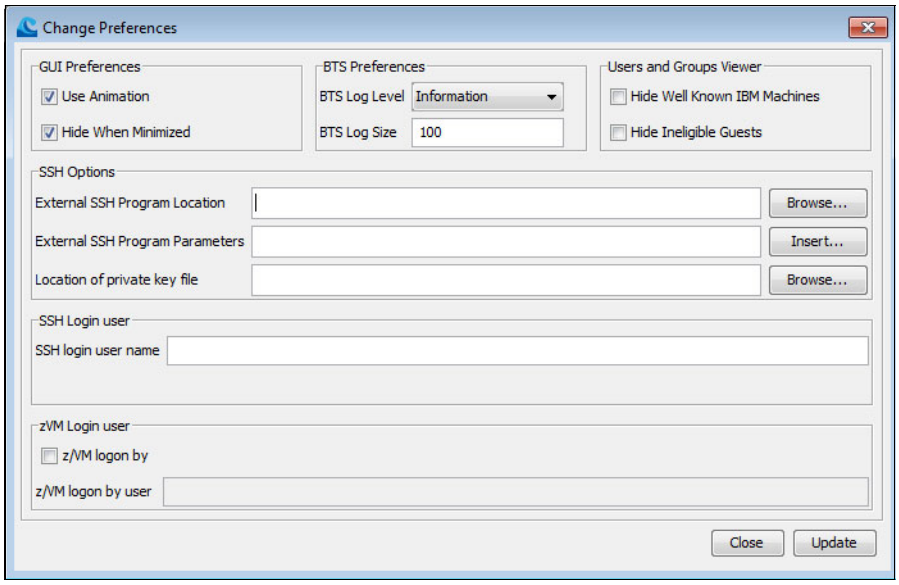


*Figure 7-3 Change Preferences window*

2. Click **Browse** that is next to the External SSH Program Location field. A window opens. Browse to the location of your `putty.exe` file and click **Select**, as shown in Figure 7-4.

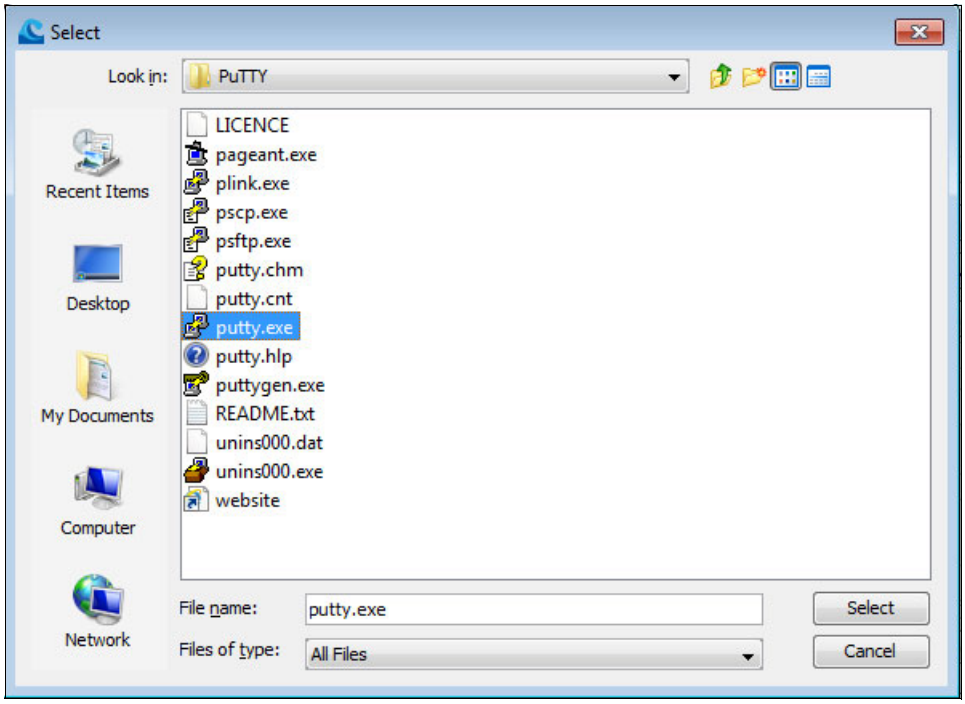


*Figure 7-4 Browse to the External SSH program location*

Your external SSH program is placed into the External SSH Program Location field, as shown in Figure 7-5.

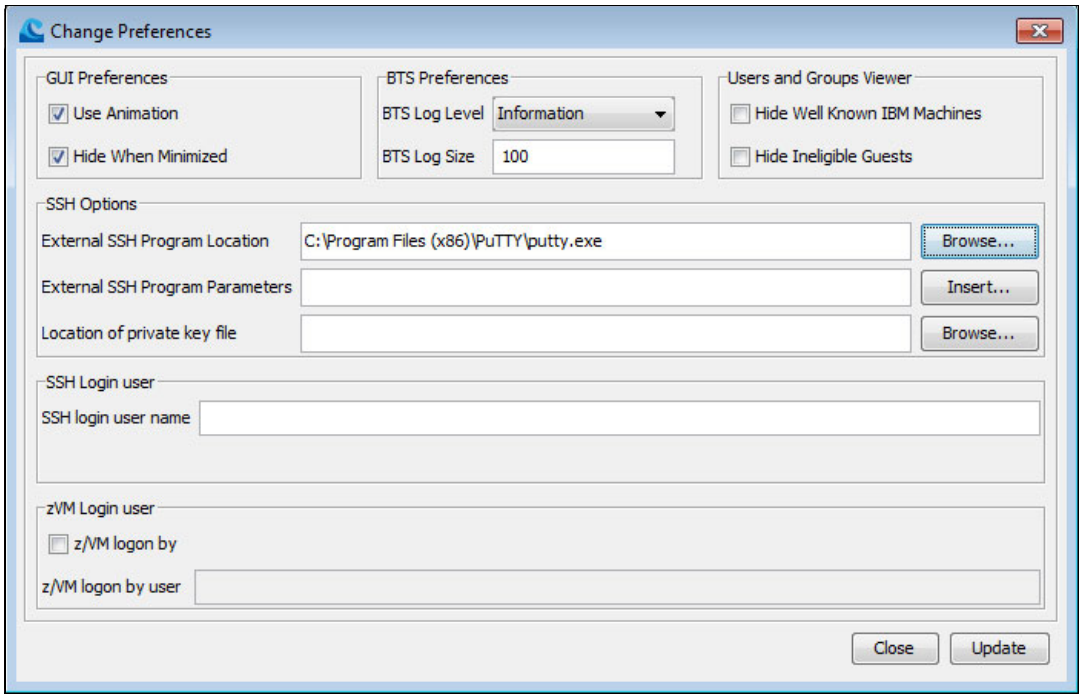


*Figure 7-5 External SSH Program Location field completed*

IBM Wave must understand the correct format for passing in the IP address, user name, and password to the SSH program as arguments. In the case of PuTTY, this is done by clicking **Insert** that is next to the External SSH Program Parameters field. A menu appears similar to that shown in Figure 7-6. Through this menu, you can format the field with the symbolic notation that enables IBM Wave to insert dynamically arguments at run time.

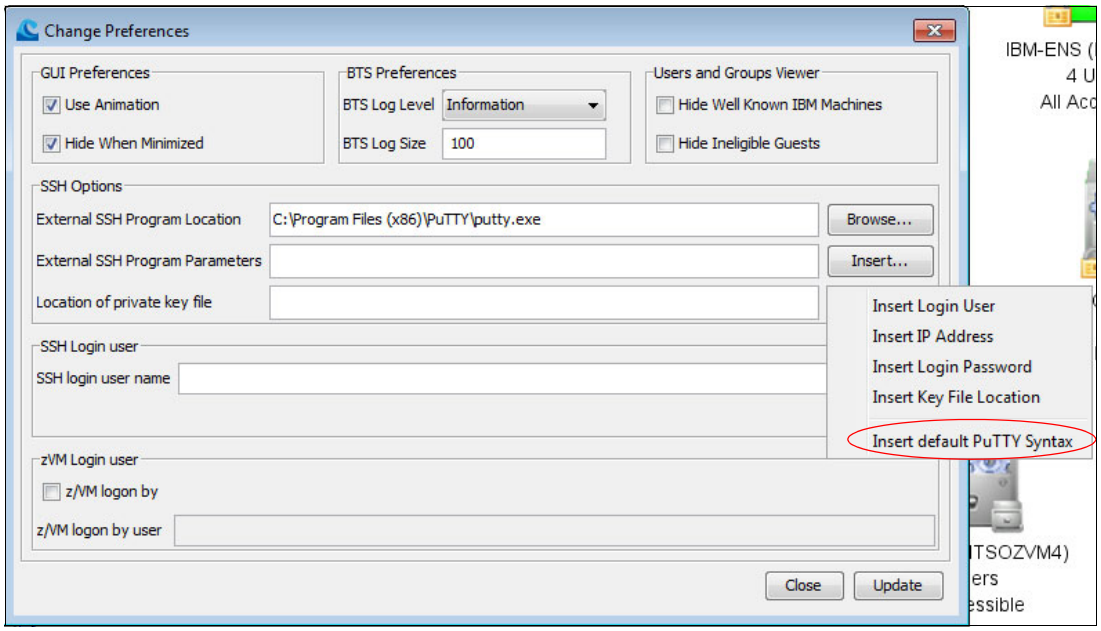


*Figure 7-6 Insert default PuTTY Syntax option*

3. Click **Update**. The External SSH Program Location field now should appear similar to the example in Figure 7-7.

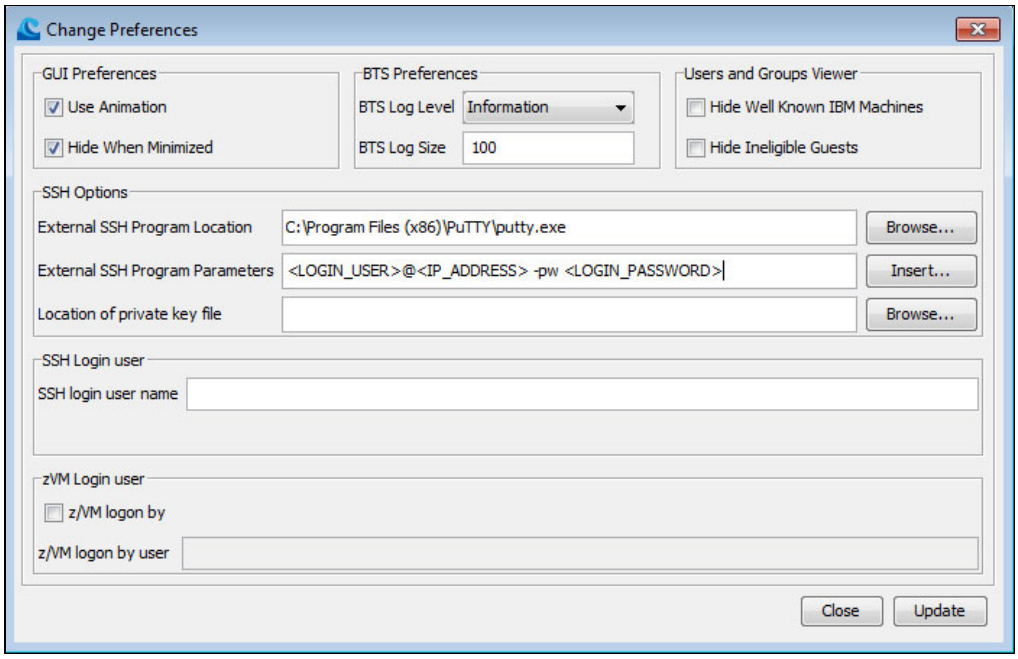


*Figure 7-7 External Program Location with Default PuTTY syntax*

To start your SSH program sessions with X services, you must modify the statement in the External SSH Program Parameters field to include the `-X` parameter. This parameter indicates to the SSH program that an X server is locally available from your workstation for this SSH session. An X server makes it possible to start graphical windows, such as yast2, from your SSH session.

4. Add the `-X` parameter to the External SSH Program Parameters, as shown in Example 7-1.

*Example 7-1 Include the -X parameter*

```
-X <LOGIN_USER>@<IP_ADDRESS> -pw <LOGIN_PASSWORD>
```

5. Click **Update**. Your SSH sessions now use X Window System when possible.

**Note:** If you use the same IBM Wave user ID on multiple computers, you must ensure that the path to `putty.exe` is identical on each computer to avoid having to reconfigure this setting.

## 7.1.4 Starting an SSH program from IBM Wave

Now that your clones are active and your external SSH program is configured, you can start SSH sessions and access any of your clones that are running the Linux operating system.

Complete the following steps:

1. Right-click the active Linux system that you want to access. As shown in Figure 7-8 on page 225, right-click TESTSUSE and select **Access** → **SSH Access**.

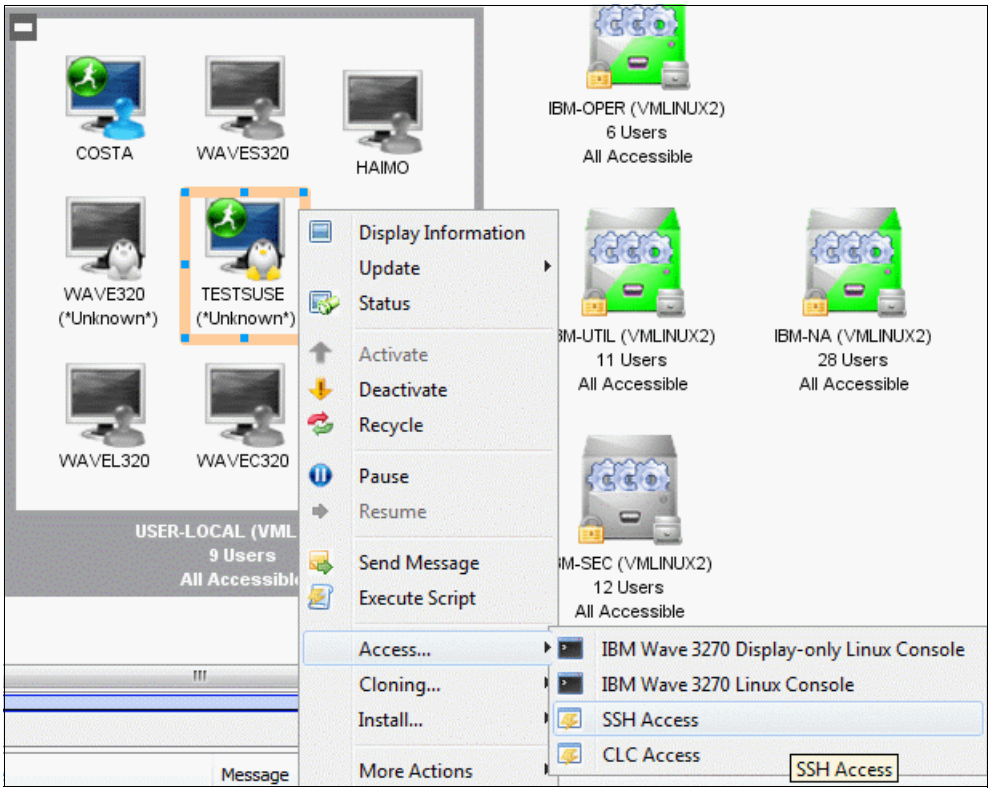


*Figure 7-8 Access an external SSH program*

2. Enter the user name and password for the system in the window, as shown in Figure 7-9.

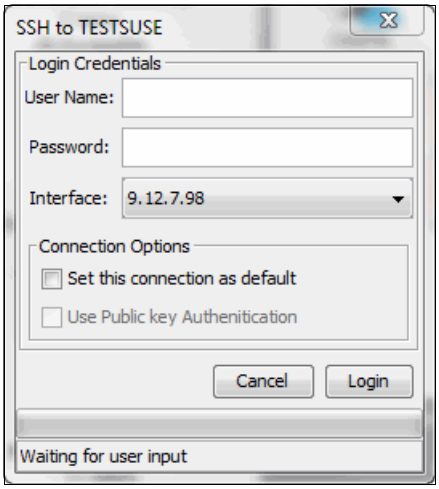


*Figure 7-9 SSH credentials*

You are now connected to your Linux guest by using the external SSH application (in this case, PuTTY). This example uses XMing as the X Window System provider.

## 7.2 Creating a z/VM guest

After the user preference is updated, a new z/VM guest must be created. Use the ties of IBM Wave to z/VM to use the graphical environment to define a guest VM. During this process, IBM Wave presents a series of windows with which you set the characteristics of the guest that you want to define.

In the background, IBM Wave works through a directory manager to define a new directory entry for your guest and dynamically places this directory entry into production. Finally, IBM Wave refreshes its view of the underlying z/VM environment to graphically indicate this new guest.

The example view is set at Site Defined Groups. Ensure that the correct group is selected when you are creating your z/VM user. In this example, use the Group by: Project group to create the user, as shown in Figure 7-10.

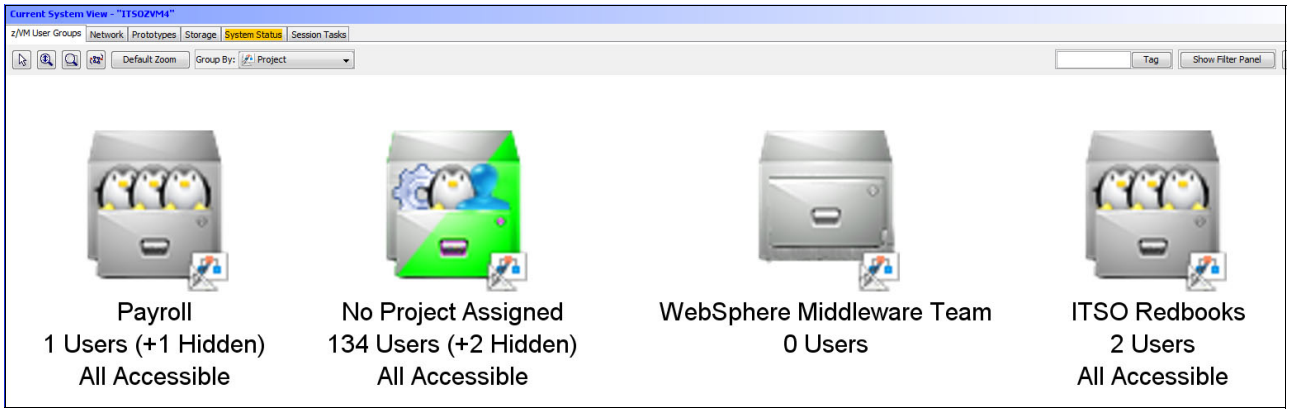


*Figure 7-10 Group by: Project*

Complete the following steps:

1. Define a new guest in the IBM Wave project that is named ITSO Redbooks. Right-click anywhere in the white space within the project to open a menu and select **z/VM User Actions** → **Create New z/VM User**, as shown in Figure 7-11.

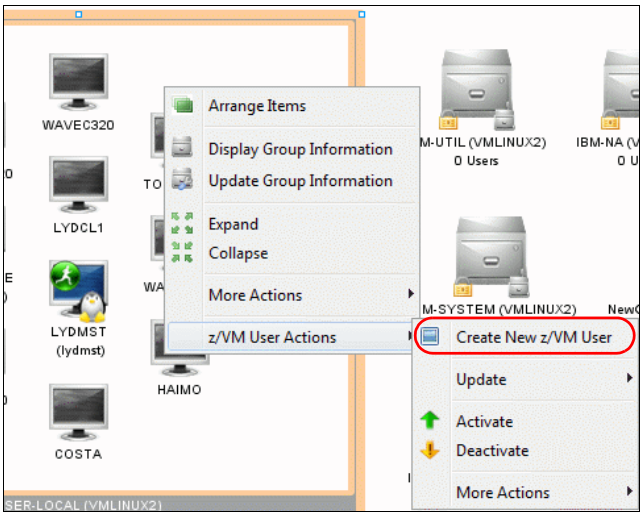


*Figure 7-11 Create New z/VM User option*

A window that is called Create New z/VM guest with the fields that are required to create a z/VM user opens. An example is shown in Figure 7-12.

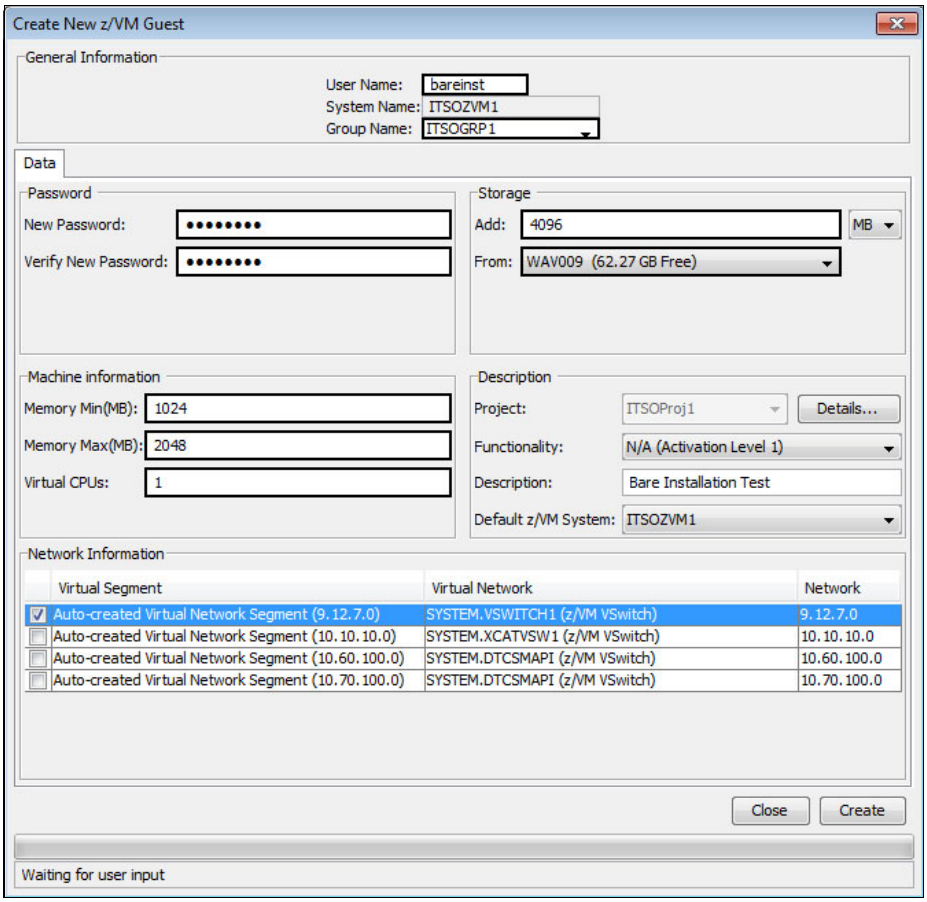


*Figure 7-12 Parameters to create z/VM guest*

The window has the following fields:

- User Name: The z/VM User ID of your new guest. This name is limited to up to eight characters. This example uses the value BAREINST.
- Password: The password that is used for logging in to z/VM as your new guest.
- Storage: The number and menu identify the size of the storage that you want to define and allocate to your new guest. The menu to the right of the “from” label is used to identify the storage group from which you intend to allocate your storage (this example uses WAV009). The default value is 2048 MB (this example uses 4096 MB).
- Memory Min/Max (MB): In these fields, you identify the minimum and maximum amount of memory that your system can be allocated by the z/VM Control Program. IBM Wave does resize these memory values. Size the memory values for the memory requirements of your workload. By using IBM Wave, you cannot perform a bare metal installation if your memory minimum is smaller than 1024 MB.
- Virtual CPUs: In this field, you set the number of virtual processors that are available to the guest. The default value is 1.

- Project: By using this menu, you can assign the guest to a predefined project. When this task is run from within the context of an IBM Wave project, the name of the project is shown here and the menu is disabled. If you run this task outside of the context of a project, you must select a project here. You can view the more information about your selected project by clicking **Details**.
- Functionality: This menu is used to express the interdependencies between Linux guests. For example, if an application server depends on a database server, the database server has a higher activation level to ensure that it is available when the application server starts. Because we are not using this function in this example, N/A (Application Level 1) is the only option, which means that all Linux guests have the same priority for startup.
- Description: This is a field that is used to insert a description of the function that is performed by the guest. This information is useful because of the z/VM limitation that a user name can be no longer than eight characters, which prevents the user name from effectively conveying the purpose of the guest. This field is stored as metadata in the IBM Wave database. Enter a string of your choosing in this field.
- Default z/VM system: In this field, you define the default member of an single-system image (SSI) cluster that the guest loads if the entire cluster is restarted. You can also select a stand-alone z/VM system here.
- Network Information: This area displays a list of available virtual network segments for your guest. The number of available segments can be zero to many, depending on the network configuration of your z/VM environment. If you define a virtual switch, you might see at least one segment that is automatically generated to connect your guest to your virtual switch. Select VSWITCH or NIC, whichever you want your z/VM user to use.

2. When you complete the window, click **Create**. If there is an invalid entry in any of the required fields, you receive an error message that returns you to the window to correct the field. If there is a missing optional field (such as the Description field), you see a window in which you are prompted to confirm that you want to create the guest without the optional field.

   After all of the fields are accepted, a progress bar appears at the bottom of the window next to the area that explains the z/VM tasks that IBM Wave is automatically performing.

To see your new z/VM user ID, double-click the project group where it was created, in this case, the ITSO Redbooks project. You see BAREINST is listed and disabled, which means that it is inactive, as shown in Figure 7-13.


*Figure 7-13 New z/VM guest BAREINST in the ITSO Redbooks project group*

Double-click the new user ID to see the general information about your user ID. The first tab, Data, shows on which member of your cluster the new user ID can run (if you are using an SSI cluster), as shown in Figure 7-14.

**Note:** If the cluster is recycled, your user ID loads on the z/VM system in which it was created.

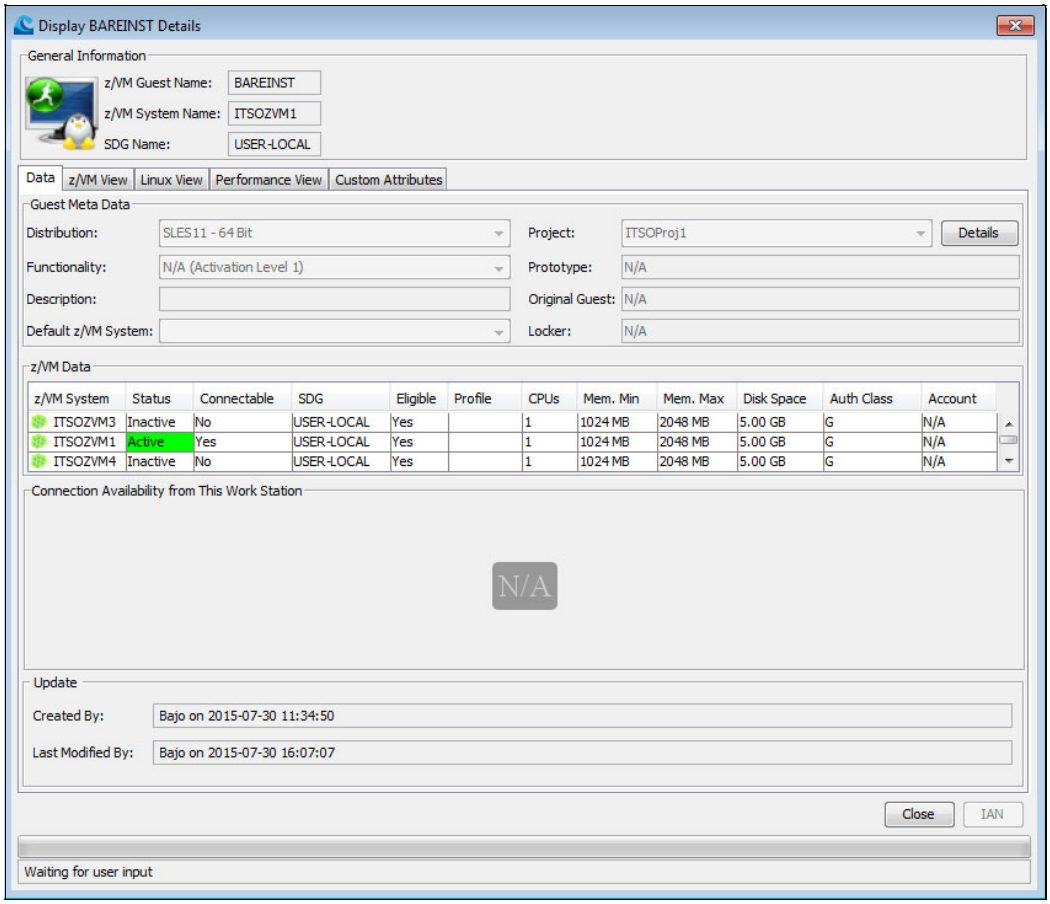


*Figure 7-14 New user ID general information: - Data tab*

You can examine your z/VM directory information by clicking the **z/VM view** tab. Here, you see that your user ID was created with the name BAREINST, with one virtual CPU, a NIC that is connected to a vSwitch that is named VSWITCH1, and a minidisk that is 5 GB, as shown in Figure 7-15.

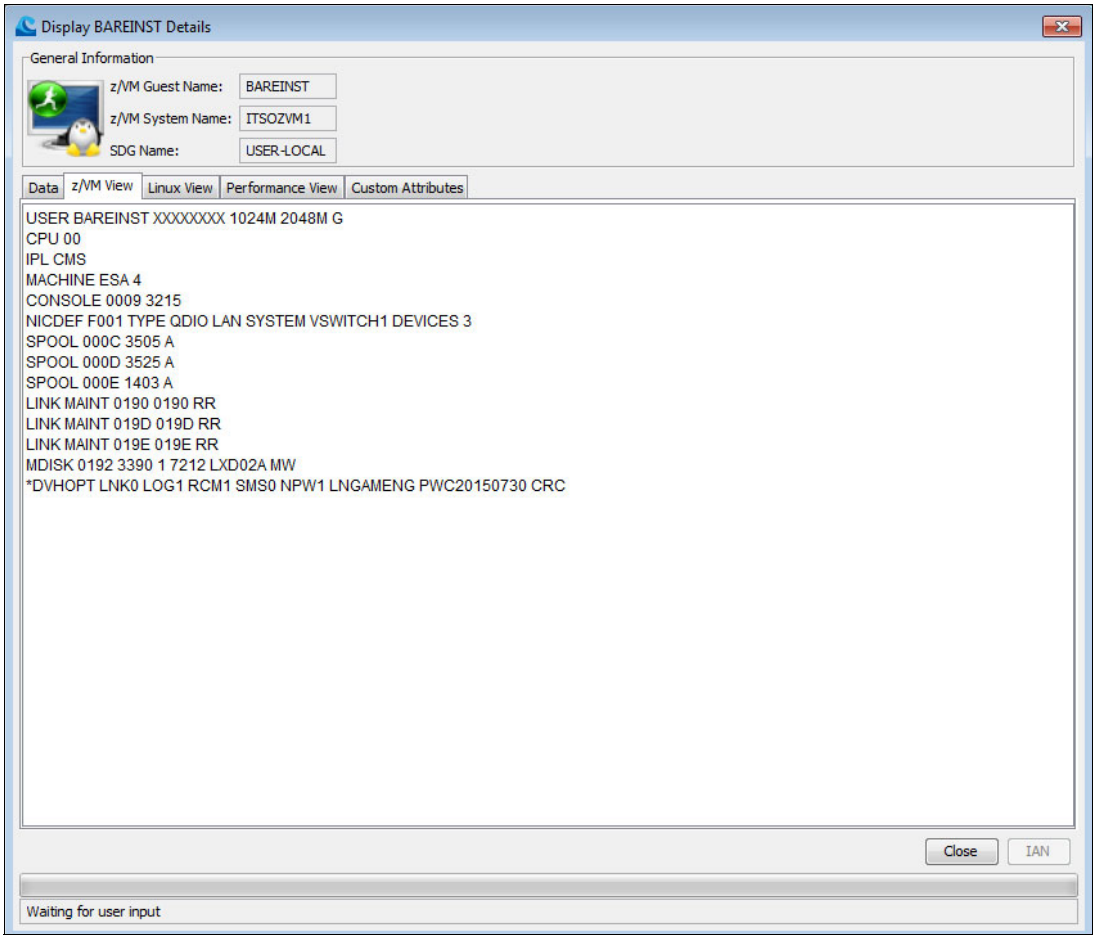


*Figure 7-15 New user ID general information: - z/VM tab*

The z/VM tab features the following information:

- USER: BAREINST is user ID name, in which XXXXXXX represents the password and 1024 MB and 2048 MB are the minimum and maximum storage values. The last parameter, G, is the privilege class (in this case, General).
- CPU: In this case, 00 represents a single virtual process that is reserved for this guest.
- NICDEF: This is the Network Interface Card at virtual address F001 that connects this user ID to vSwitch VSWITCH1.
- MDISK: 0192 is the IBM Wave default minidisk address that is always used for guests that were created by using the Create New z/VM guest window. The size of this minidisk is 7212 cylinders on volume LXD02A.

By opening the Network View, you see a graphical representation of the new user ID, in this case, BAREINST. Figure 7-16 shows the connection to Unknown IP(VSWITCH1). An IP address is not assigned to the BAREINST user because it is still an empty shell that does not have an operating system with which it is associated.

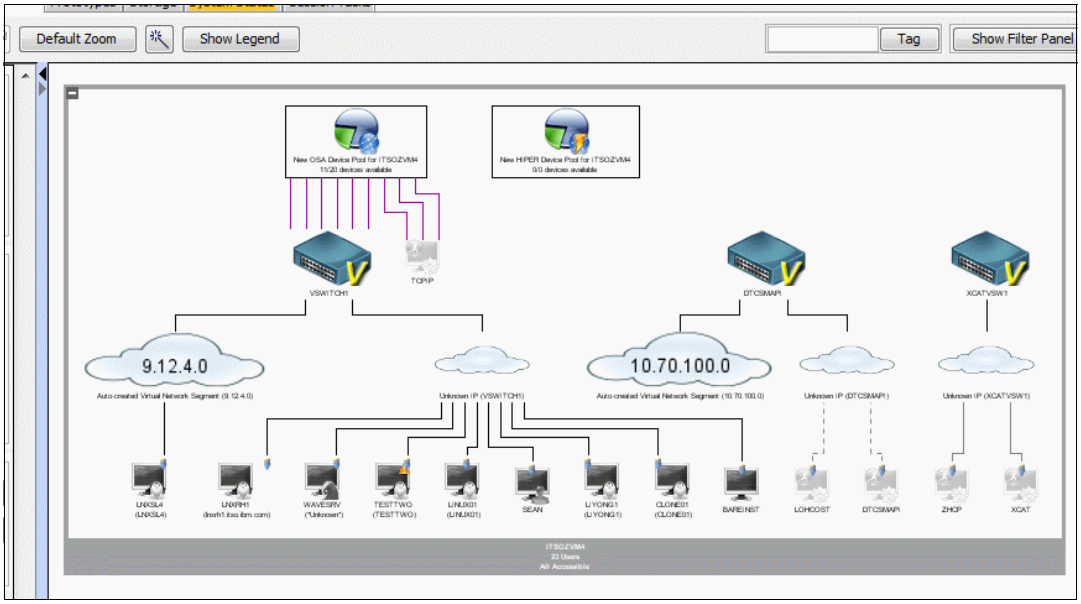


*Figure 7-16 Network view of new guest virtual machine BAREINST*

## 7.2.1 Activating, logging in to, and deactivating a new guest virtual machine

Your new guest is an empty shell that lacks a bootable disk or operating system. However, all guest VMs that are running under z/VM can access the traditional Conversational Monitoring System (CMS) operating system that is stored in a common, read-only disk that is controlled by the z/VM Control Program. This configuration enables guests that lack their own operating systems to perform z/VM systems management and programming tasks by using the traditional 3270-base CMS interface.

Complete the following steps:

1. Before you log in to the new guest VM, BAREINST must be activated. Return to the z/VM User Groups tab, click **Group By: Project**, right-click the BAREINST icon, and then click **Activate**, as shown in Figure 7-17.

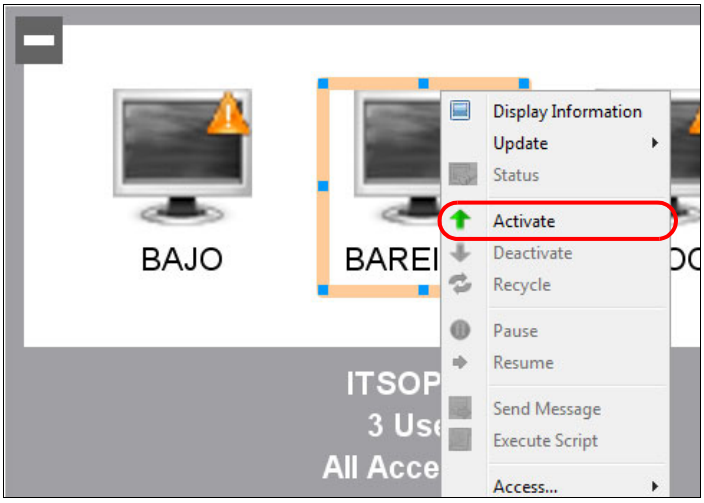


*Figure 7-17 Activate the new BAREINST guest ID*

2. A confirmation window opens, as shown in Figure 7-18. Ensure that the guest that you want activated is selected and click **Go**.

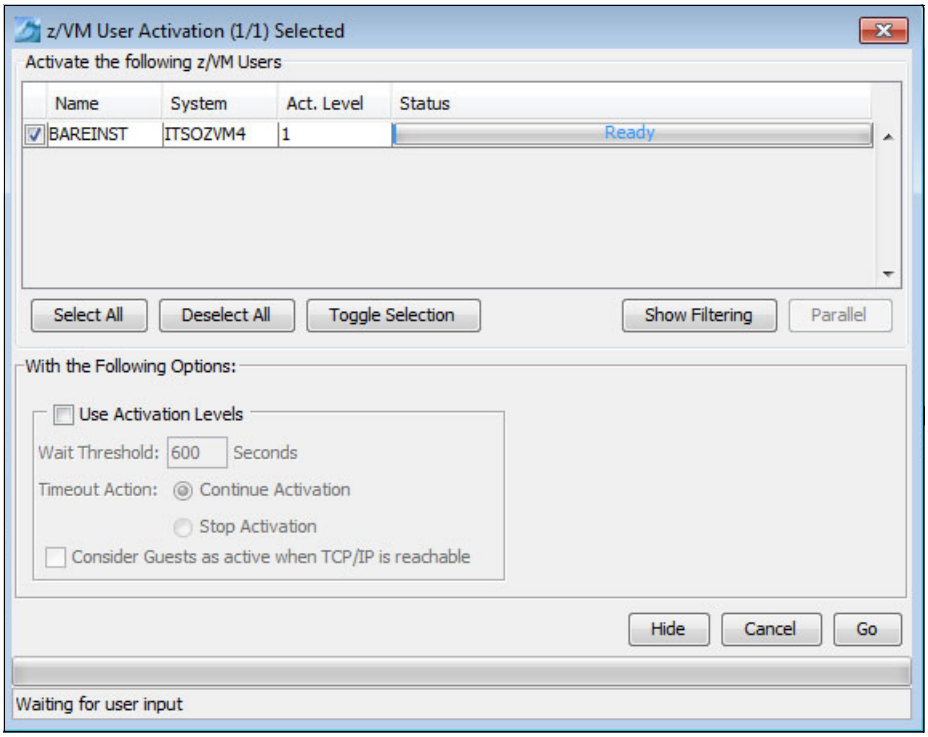


*Figure 7-18 Select a single guest ID for activation*

A request is sent to the Background Task Scheduler (BTS). Upon completion, the BAREINST icon in the Project Group features a green running man symbol, which indicates that this guest is now active, as shown in Figure 7-19 on page 233.


*Figure 7-19 Activation completion for new guest ID*

3. To access the BAREINST 3270 console, right-click the BAREINST icon, select **Access**, and then click **CSL-Wave 3270 Console**, as shown in Figure 7-20.

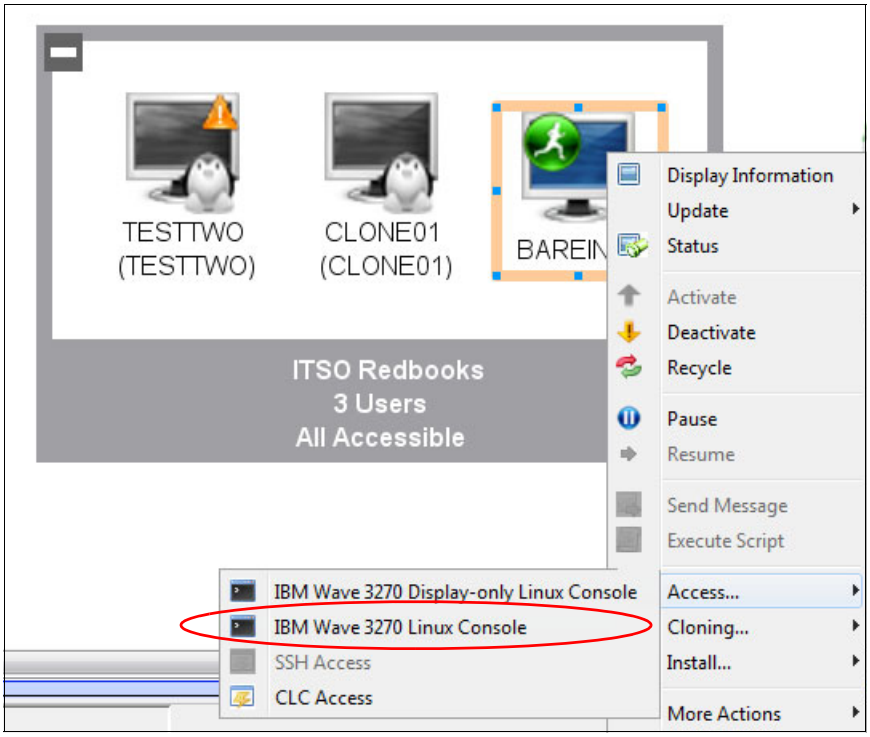


*Figure 7-20 Access the 3270 console*

4. You are prompted for the BAREINST password, as shown in Figure 7-21.


*Figure 7-21 Password prompt*

The BAREINST console opens by using the IBM Wave integrated 3270 terminal emulator, as shown in Figure 7-22.

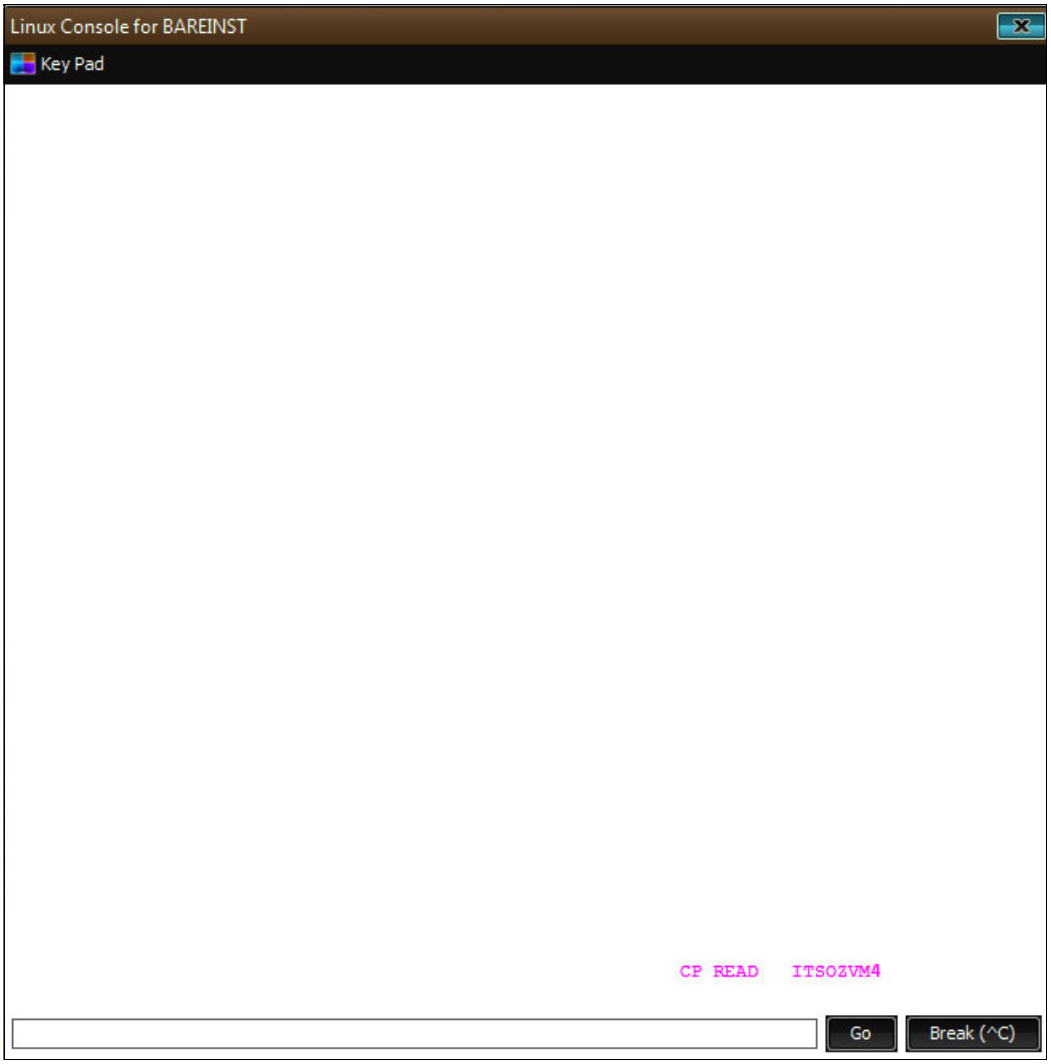


*Figure 7-22 3270 console display*

5. Run the **Q DASD** command to see the direct access storage device (DASD) that was assigned to you by IBM Wave and the directory manager, as shown in Figure 7-23.

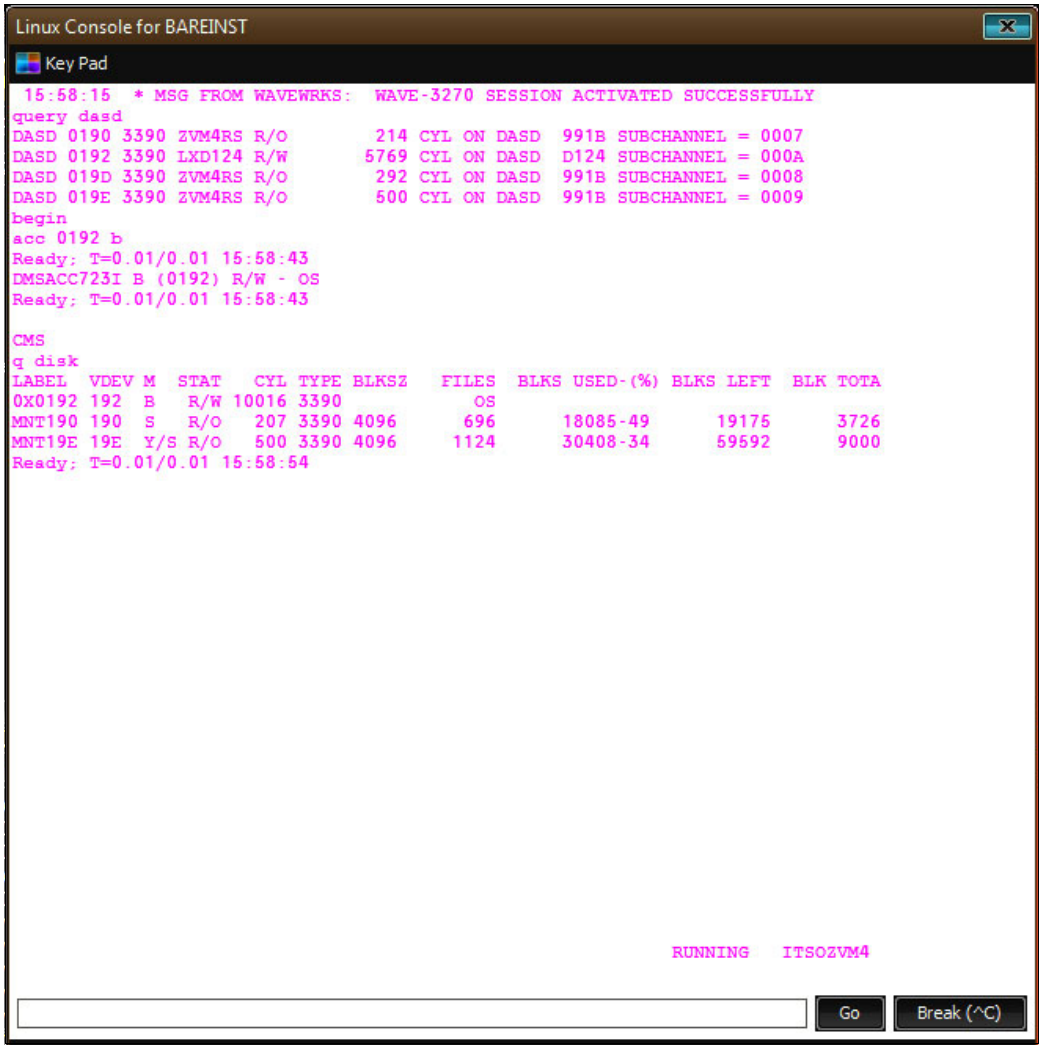


*Figure 7-23 z/VM console commands*

6. The 192 disk that has 5769 cylinders on volume LXD124 is shown in the output of the command. Enter **b** on the command line and click **Go** to enter the CMS environment. You then can run other CMS commands. Run **access 192 b** to assign this disk to file mode B. Run **q disk** to view the various allocated and linked disks for this user.
7. Click the red **X** in the upper right corner to close BAREINST's console and return to IBM Wave.

8. To proceed with the installation of Linux onto the new BAREINST guest, it must be deactivated. Right-click the BAREINST icon and select **Deactivate**, as shown in Figure 7-24.

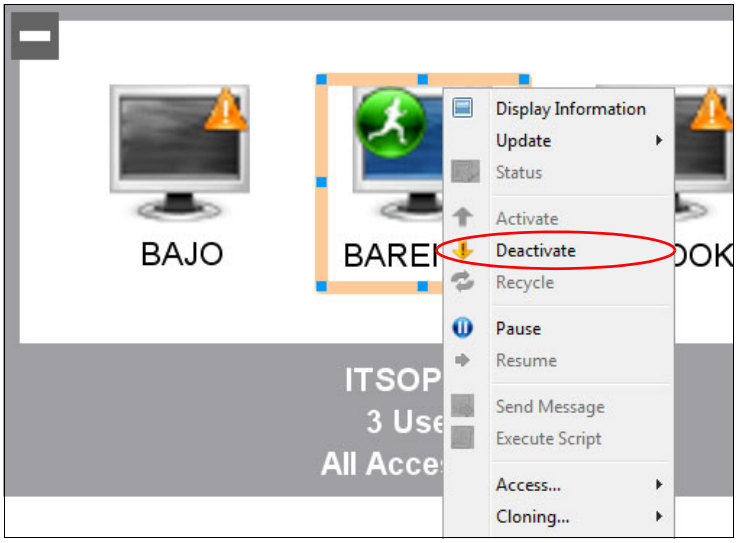


*Figure 7-24 Deactivate a guest ID*

When the deactivation process is complete, the BAREINST icon is disabled. You are now ready to install Linux.

**Note:** To make the guest created eligible for LGR in installations with SSI, some settings must be changed. If necessary, remove all local resources.

## 7.2.2 Installing Linux on your guest virtual machine

You now start the SUSE Linux Enterprise Server installation through the IBM Wave Bare Metal Installation wizard feature. This feature automates the process of building bootable Linux media that can log in to a Linux repository through TCP/IP, transferring this media to BAREINST, and starting BAREINST by using this installation media. This process eliminates the need to manipulate manually SPOOL devices to get to a Linux Installation wizard.

Complete the following steps:

1. To start a bare metal installation, right-click **BAREINST** and select **Install** → **Launch Linux Installation**. You see a window that is similar to the window that is shown in Figure 7-25 on page 237.

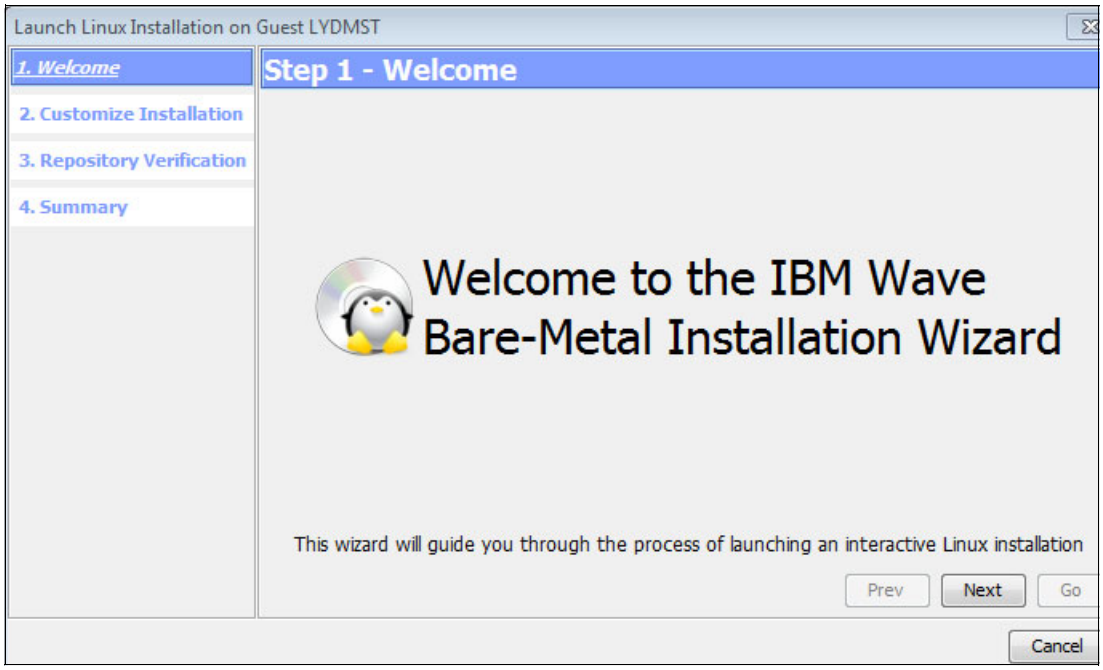


*Figure 7-25 Wizard Welcome window*

2. Click **Next**, which opens another window, as shown in Figure 7-26. Ensure that the Hostname and Networking Interface reflect your installation. Change the Temporary Installation Password to the permanent value that you want.

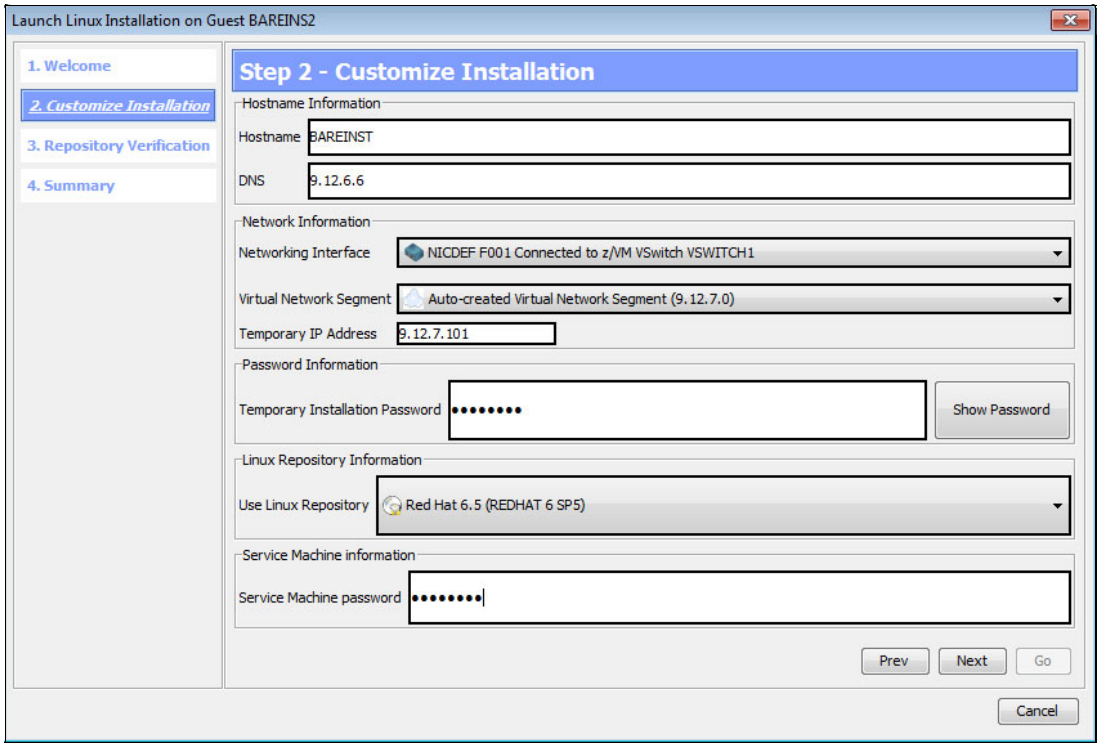


*Figure 7-26 Customize the installation*

3. Click **Show Password** so that you can see, in clear text, the value that you entered, as shown in Figure 7-27.


*Figure 7-27 Password display*

4. Click **Next** and you see a flashing window while the repository is verified. After a few seconds, you see a message that reads, “Repository Verified”, as shown in Figure 7-28.

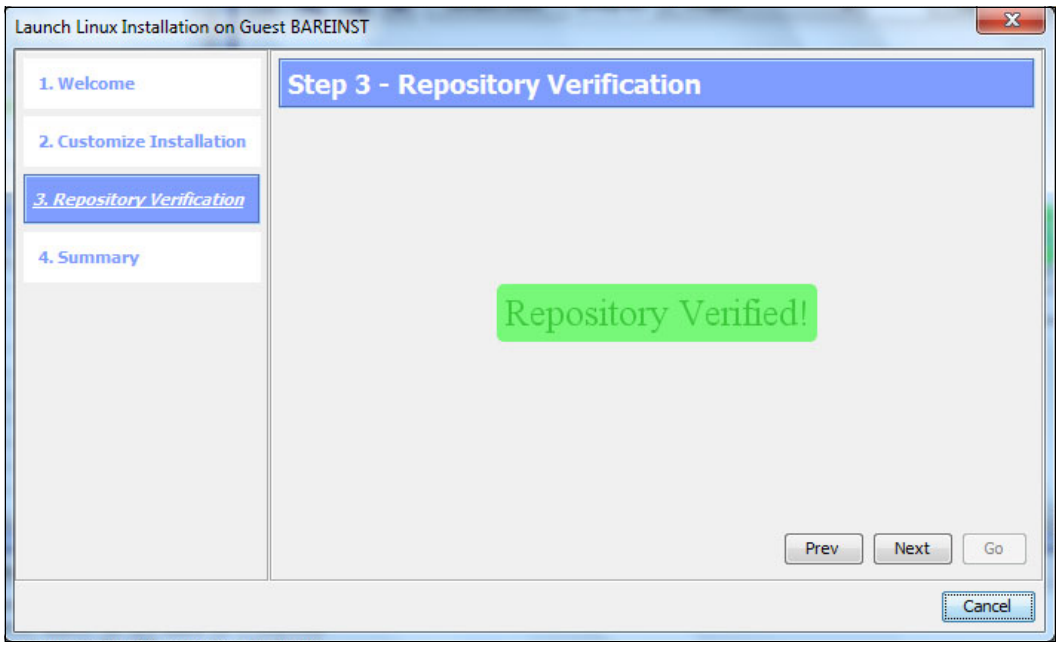


*Figure 7-28 Linux repository verification*

5. Click **Next**. Another window in the Repository Verification process opens. On the Select Storage Group, select your storage group, in this case, it is $3390$, as shown in Figure 7-29 on page 239.

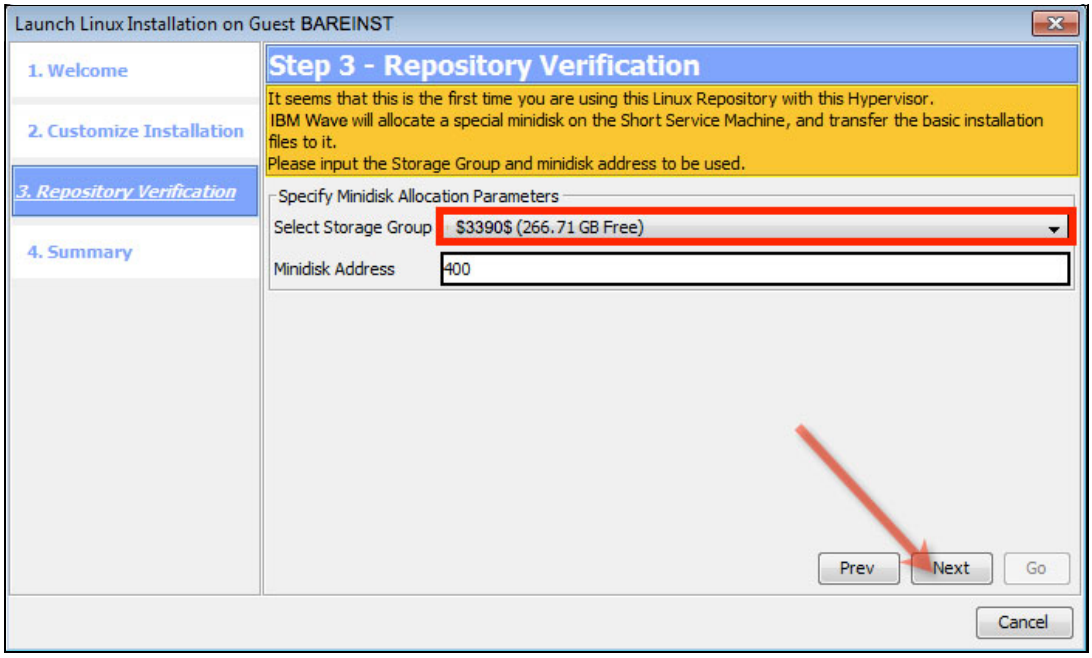


*Figure 7-29 Repository Verification Storage Group selection*

6. After the repository is verified, click **Next** to view the Summary, as shown in Figure 7-30.

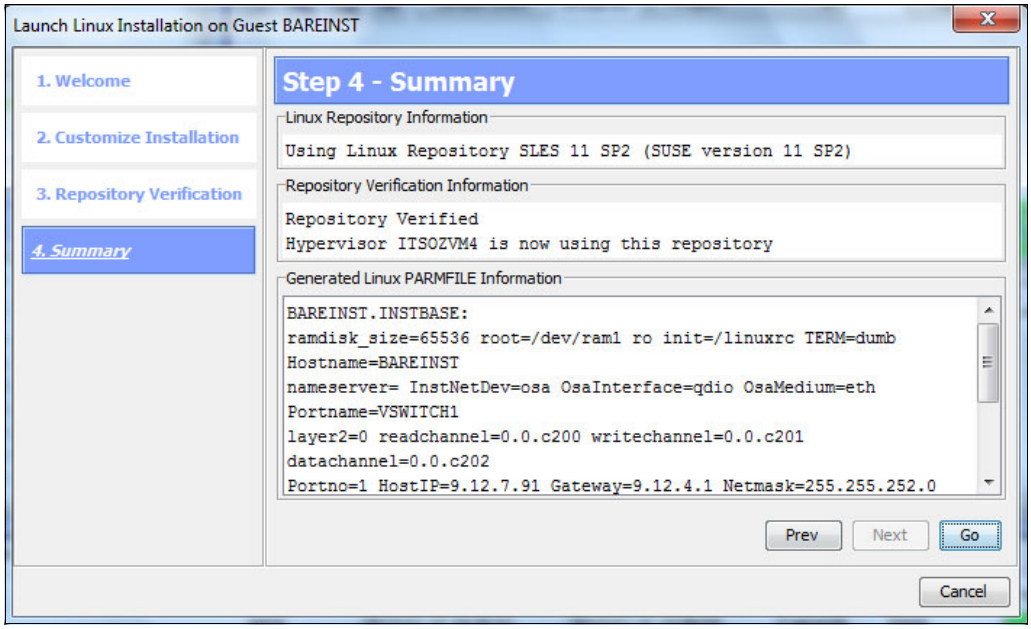


*Figure 7-30 Repository Information Summary*

You should verify the following information in the summary:

- You are installing the Linux system that you expect (in this example, SUSE Linux Enterprise Server 11 SP 2).
- Virtual NIC addresses are assigned to the vSwitch (in this example, c200-c202 a).

- TCP/IP information regarding the IP address for the Host, Gateway, Netmask, and Broadcast are correct.
- Installation repository location is correct.

7. Click **Go**. The Linux Installation work unit is started, which submits a request to the BTS, as shown in Figure 7-31.

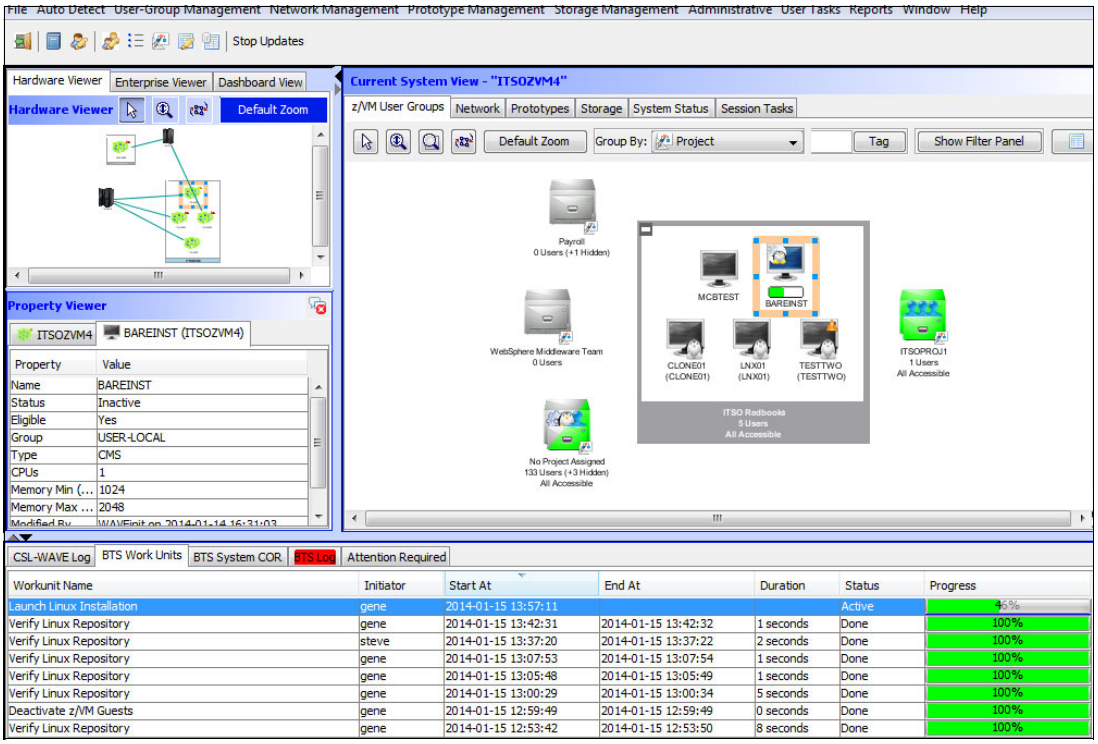


*Figure 7-31 Work unit progress display*

8. Double-click the BTS Work Units tab to view more detailed information, as shown in Figure 7-32.

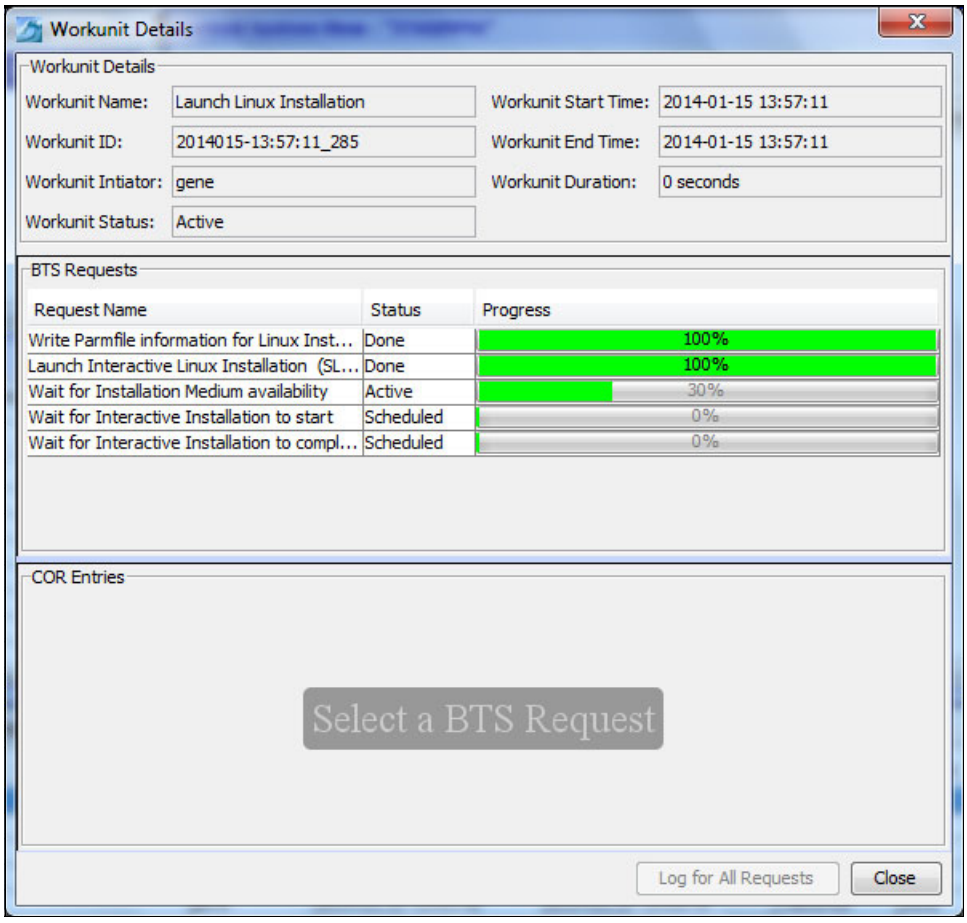


*Figure 7-32 BTS work unit details*

Upon completion, a window opens that indicates that the Interactive Installation is ready, as shown in Figure 7-33.

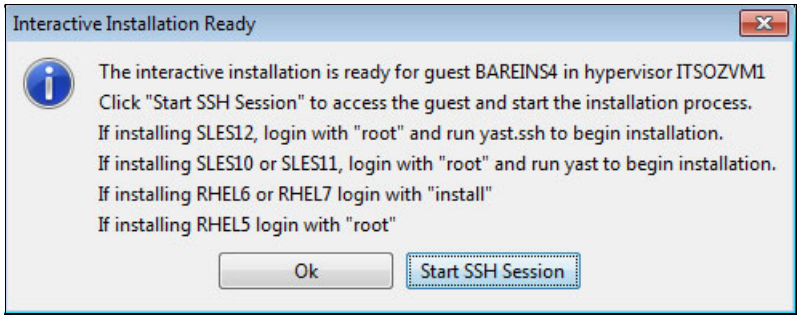


*Figure 7-33 Interactive Installation message*

From here, you must continue manually by clicking **Start SSH Session** only. For more information about a manual Linux installation, see *The Virtualization Cookbook for IBM z Systems Volume 2: Red Hat Enterprise Linux 7.1 Servers*, SG24-8303 or *The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3*, SG24-8147.

9. After the installation is complete, you must right-click the BAREINST icon and select **Reset Installation Status**, as shown in Figure 7-34.

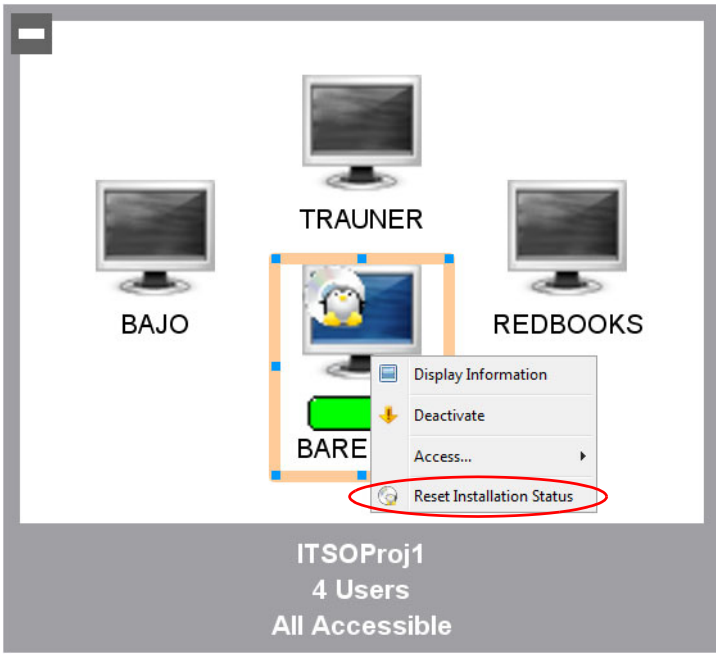


*Figure 7-34 Reset Installation Status*

You see that a penguin icon was added to the BAREINST icon, as shown in Figure 7-35.

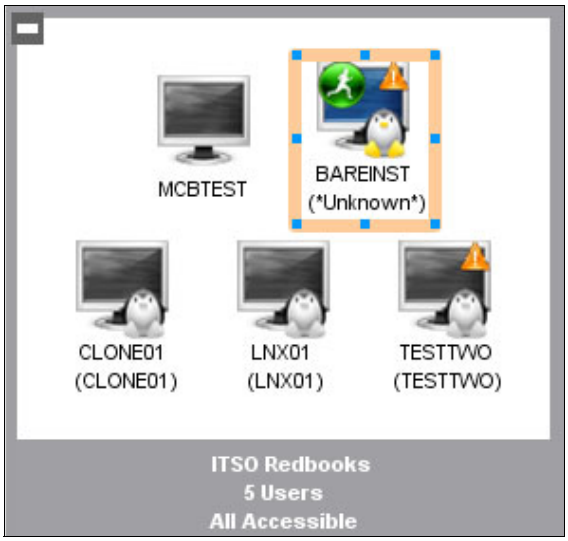


*Figure 7-35 Linux installation completion*

## 7.3 Initializing guests by using IBM Wave

After a new Linux guest is installed or imported, the guest must be initialized by IBM Wave to be managed in the GUI environment.

Running the Init User for IBM Wave use task on a Linux guest has the following results:

- Opens an SSH tunnel into the Linux guest.
- Verifies that a version of IBM Wave is not installed on the guest.
- Verifies root privileges.
- Verifies the following Linux guest prerequisites:
  - The CMSFS package that contains the driver and utility that are required to allow Linux to access CMS files and minidisks
  - The VMCP package that contains the driver and utility that are needed to allow Linux to issue commands to the z/VM Control Program
- Verifies that the guest has the **LOGOFF ON SHUTDOWN** flag set in the `zipl.conf` file.
- Verifies that both the `zipl.conf` and `fstab` files adhere to the supported formats of IBM Wave.
- Links the WAVEWRKS 399 minidisk to the z/VM guest that is running Linux in the directory entry with read-only access.
- Adds a series of scripts to `/usr/wave`.
- Creates the Linux user ID `waveuser`, adds it to the `/etc/sudoers` file, and sets its password.

The process of initializing guests for IBM Wave management has a number of prerequisites. For both SUSE Linux Enterprise Server and Red Hat Enterprise Linux installations, a package that is called `cmsfs` is required, which allows Linux to read from CMS minidisks.

To initialize BAREINST for IBM Wave management, you must use an SSH session to access BAREINST and manually install this package.

Complete the following steps:

1. Right-click your guest and select **Access** → **SSH Access**, as shown in Figure 7-36**.**

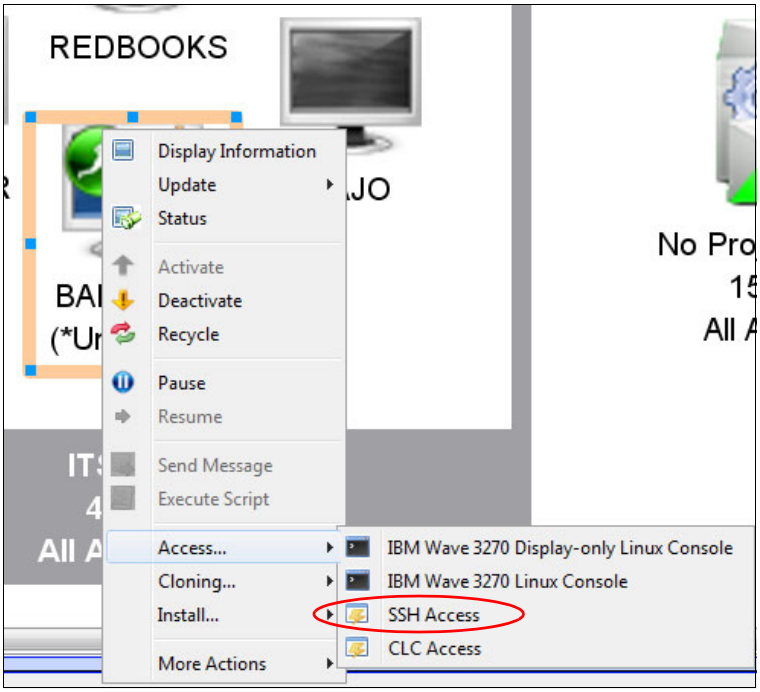


*Figure 7-36 SSH Access*

After you select **SSH Access**, you are prompted for a user name and password, as shown in Figure 7-37.

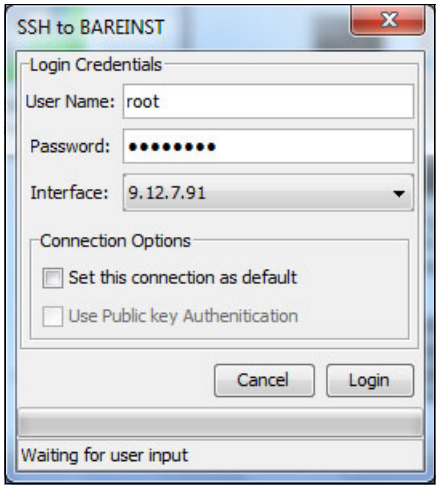


*Figure 7-37 SSH password prompt*

2. Enter the user name and password and click **Login**. An SSH session opens, as shown in Figure 7-38.

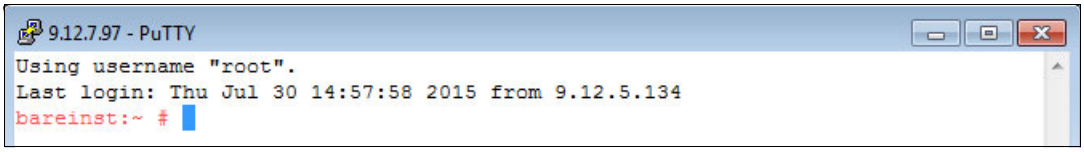


*Figure 7-38 SSH session*

3. It is possible to install the package by using a CLI, as shown in Example 7-2. The display of CLI option on SUSE Linux Enterprise Server 11 SP4 is shown in Figure 7-39.

*Example 7-2 Installation command for cmsfs*

```
# zypper install cmsfs (On SUSE distributions)
# yum install cmsfs (On Red Hat distributions)
```

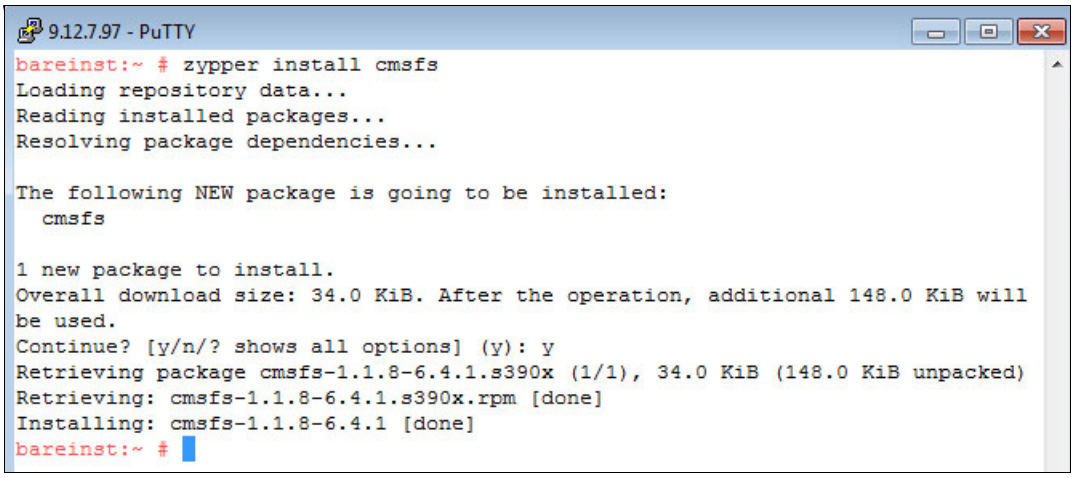


*Figure 7-39 Command-line interface installation option*

To do a graphical installation, ensure that you have an X11 terminal emulator installed. In this case, use MobaXterm. If you do not have an emulator initialized, you see a “cannot open display” error message.

4. Run the **yast2** command to open the YaST2 Control Center in an X Window. Scroll down to the Software section and click **Software Management**, as shown in Figure 7-40.

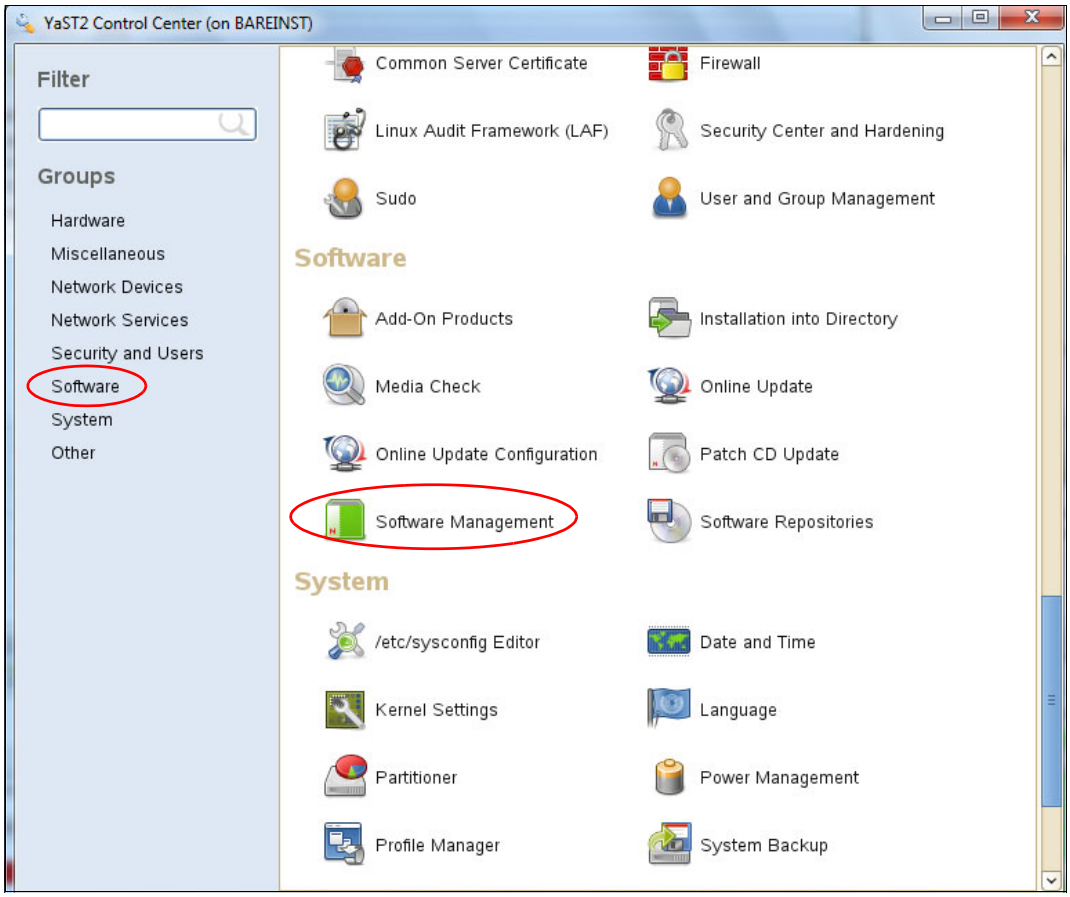


*Figure 7-40 YaST2 Control Center*

5. When the next window opens, enter `cmsfs` into the field that is next to the Search option, and then click **Search**.
6. When `cmsfs` is found, select the package that you want to install and click **Accept,** as shown in Figure 7-41 on page 247.

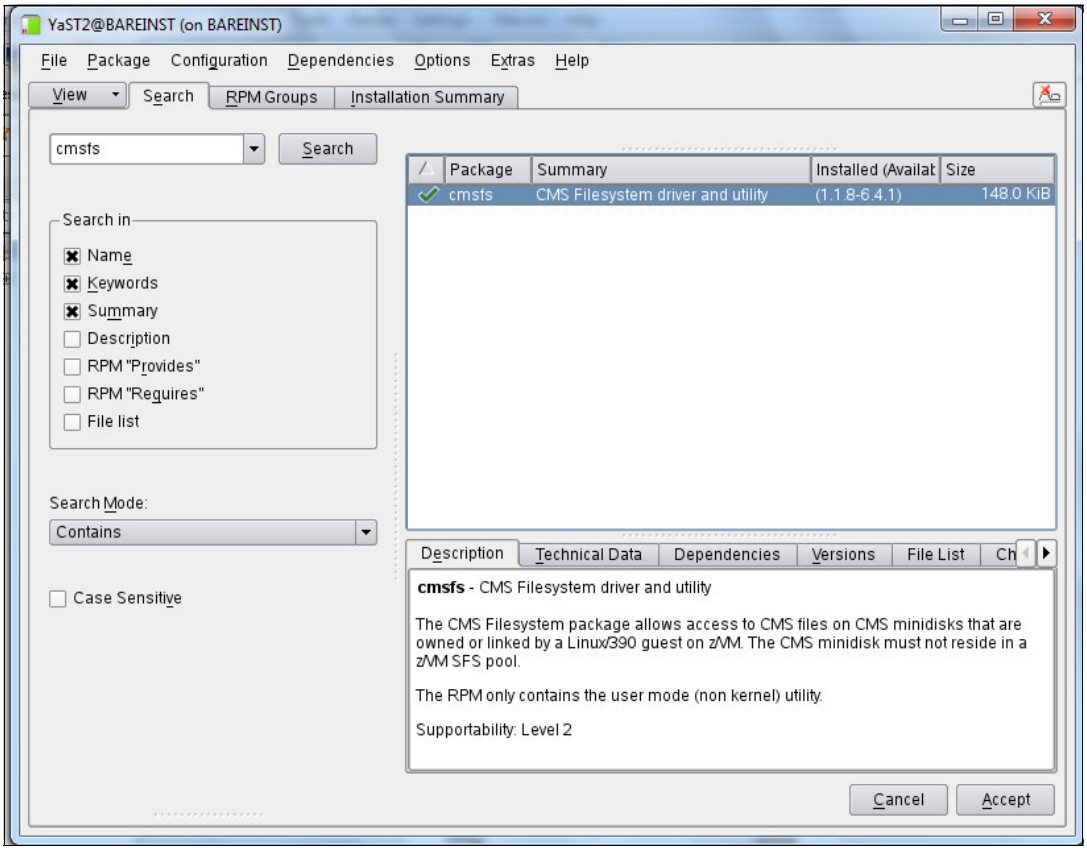


*Figure 7-41 The cmsfs package search*

7. After `cmsfs` is installed, you are returned to the YaST2 Control Center. Click the red **X** in the upper right corner to close the window. You are now back in your SSH session where you can verify that `cmsfs` was installed by running the following command:

   ```
   rpm -qa | grep 'cmsfs'
   ```

   If the `cmsfs` package is displayed, you successfully installed the package. Run the **`exit`** command and close your SSH program, as shown in Figure 7-42.

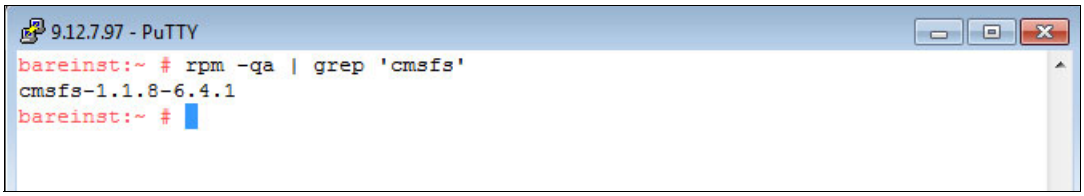


*Figure 7-42 The cmsfs installation confirmation*

You can now successfully initialize BAREINST for IBM Wave management.

8. Right-click the BAREINST icon and select **More Actions** → **Init User for IBM Wave Use**, as shown in Figure 7-43.

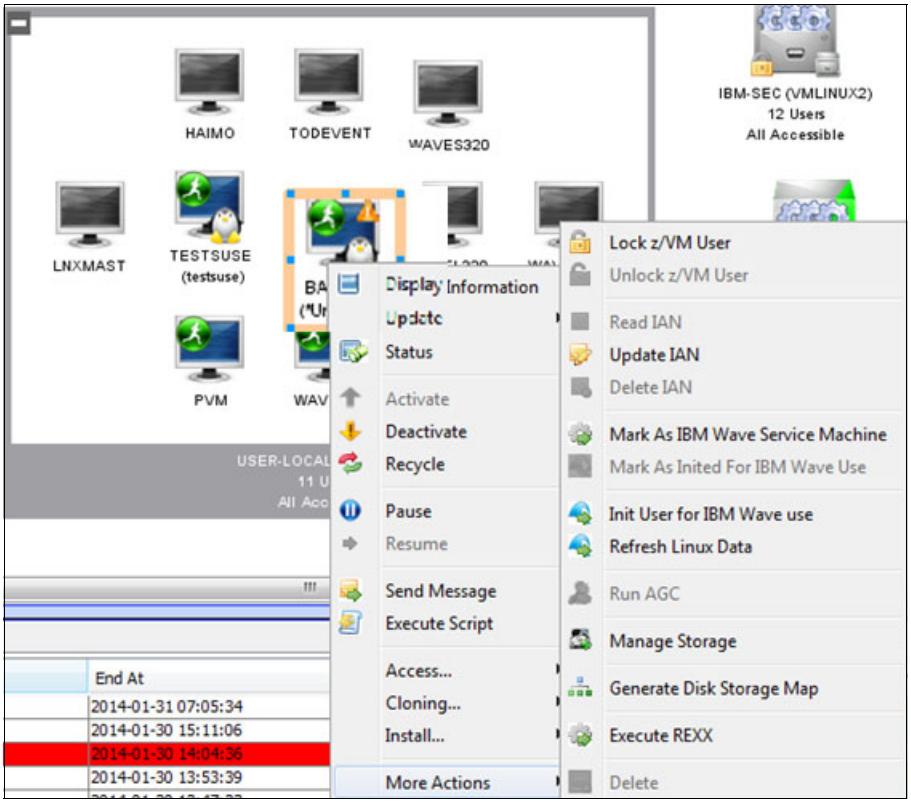


*Figure 7-43 Init User for IBM Wave use*

9. In the Initialize User for IBM Wave Use window, ensure that BAREINST is selected, enter the password, and then click **Go**, as shown in Figure 7-44 on page 249. A work unit is submitted to the IBM Wave Background Task Scheduler.

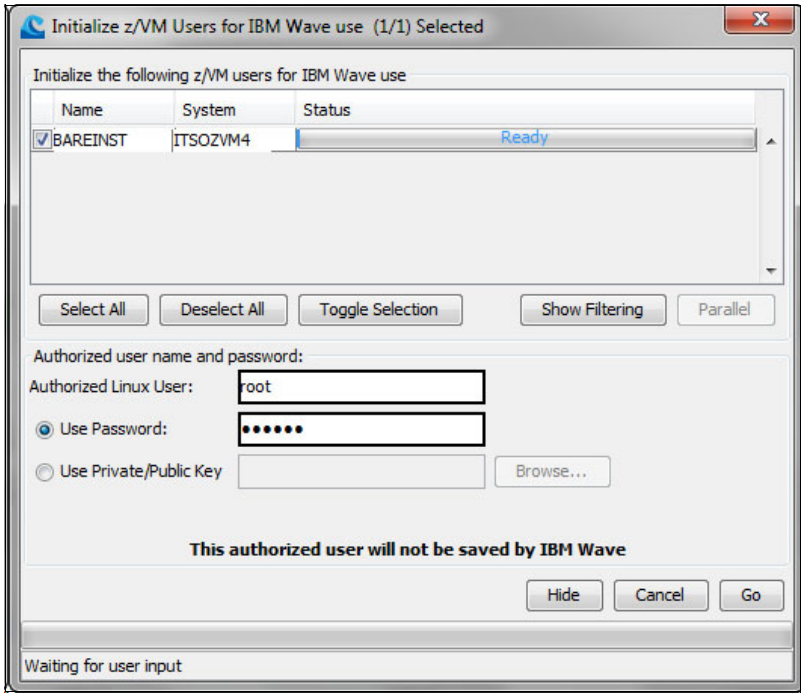


*Figure 7-44 Selection of BAREINST with password prompt for use by IBM Wave*

After a few moments, the caution symbol on the icon and the *Unknown* text beneath the user ID field disappear and BAREINST reappears, as shown in Figure 7-45.


*Figure 7-45 Completed BMI guest icon*

You now successfully created a SUSE Linux Enterprise Server guest from scratch by using the IBM Wave Add New z/VM User and bare metal installation function. In a similar way, you can create an RHEL guest.

## 7.4 Creating a golden master

A prototype is a complete z/VM User Directory entry that is used by the directory manager to create and add dynamically a guest VM to the z/VM user directory. By using this addition, a z/VM systems programmer can add a guest with the same directory attributes as the prototype by using the **LIKE** parameter.

A golden master is a template for a VM. It can also be referred to as a *golden image*, a *master image*, or a *base image*. Usually, a golden image is manually installed and clones are created with minimal modifications.

By using IBM Wave, a prototype can be created that automates the cloning of a golden image. To do this task, the following components should be in place before a clone can be created:

- Sufficient storage space (DASD)
- Sufficient memory (RAM)
- A vSwitch
- A pool of IP addresses that are connected to the vSwitch
- A fully installed Linux server that is deactivated

This section describes changing one of the Linux guest servers to become an IBM Wave prototype. You can call this prototype a golden master or template of future clones. After you create a prototype, this section describes the steps that are necessary to clone by using this prototype.

**Important:** After you change a Linux server to become an IBM Wave prototype, the Linux server you choose no longer can be logged in to; therefore, it is no longer bootable.

In IBM Wave, the concept of a golden master is implemented by using a feature that is called an *associated prototype*. An associated prototype is composed of a traditional prototype directory entry and an associated z/VM guest with assigned minidisks and an installed guest operating system. This feature is synonymous with the term golden master.

The associated guest is deactivated and has a password set to **NOLOG** to prevent it from being actively used. Within the IBM Wave GUI, the guest is filtered from the z/VM User Groups tab and visible in the Prototypes tab only.

By using associated prototypes, the IBM Wave administrator can clone a commonly used pattern. When it is combined with the functions of the Script Manager, the IBM Wave Administrator can automate much of the process of dynamically provisioning guest Linux VMs.

Associated prototypes allow for *physical cloning*, which is the process of creating a z/VM guest definition according to the definition of the prototype, or *full cloning*, which is the process of creating a z/VM guest definition and copying the data from the assigned user to the newly created user.

A disassociated prototype does not have a z/VM guest that is associated with it. When a user is cloning from a disassociated prototype, the user can choose physical cloning only. The disassociated prototype is visible only to the IBM Wave user that has the Site Level administrator role.

To create and convert a z/VM guest into an associated prototype, the guest must be inactive. For more information about how to deactivate a z/VM guest, see Chapter 6, “Managing IBM z/VM” on page 169.

If you are running in a non-SSI environment, you can easily see whether a member is inactive because the z/VM guest icon is disabled. If you are running in an SSI environment, it is impossible to tell easily if a z/VM guest is inactive on all members of the cluster because the views are generated on a per-member basis. For this reason, you must open the detailed information window to determine whether this guest is active on any members in the SSI. This window is accessible by double-clicking the guest icon or right-clicking the guest icon in any view and selecting **Display Information**. The middle of the data tab displays a table that shows the state of the guest on each member of the SSI cluster. In Figure 7-46 on page 251, guest LNXSL1 is Inactive in all members of the cluster.

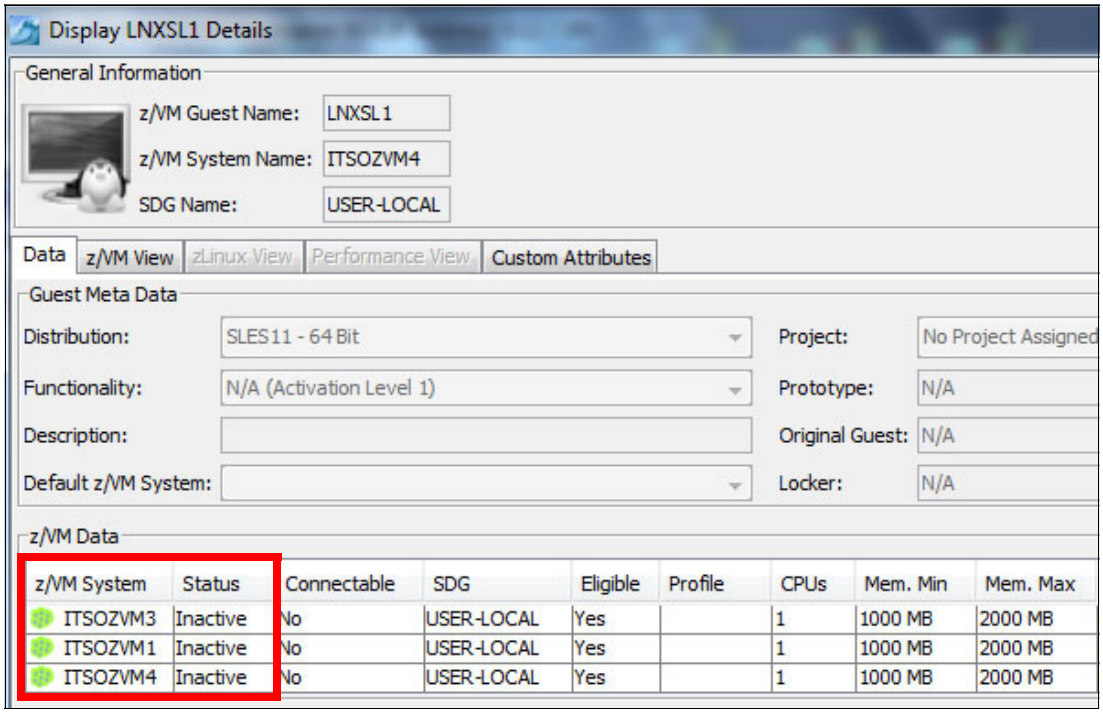


*Figure 7-46 Status of a guest in members of a cluster*

After the z/VM guest is inactive, right-click the icon and select **Cloning** → **Convert VM User to Prototype**, as shown in Figure 7-47.

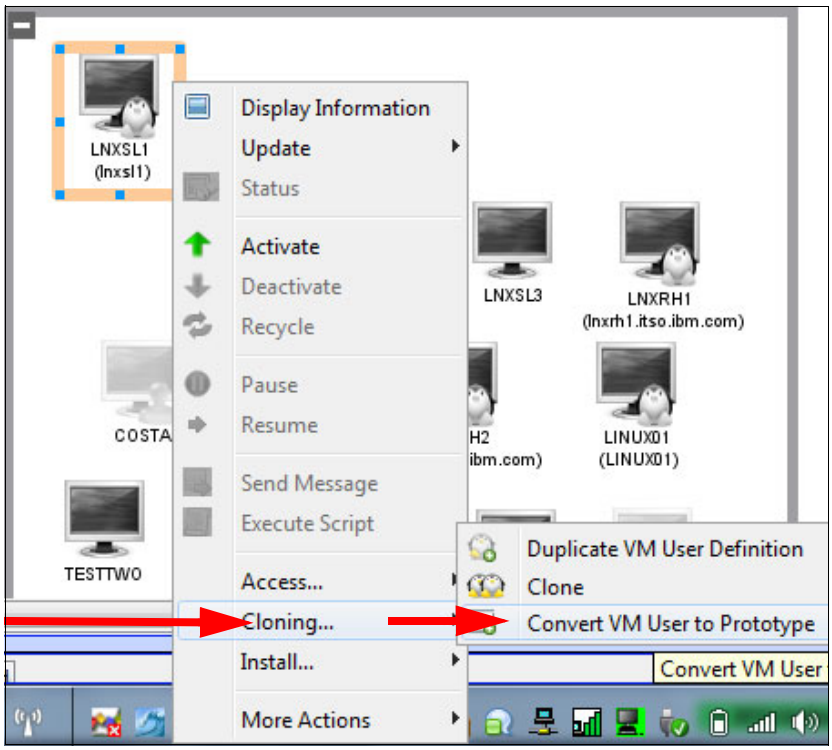


*Figure 7-47 Select "Convert VM User to Prototype" from an Inactive z/VM guest*

A Create Prototype From z/VM User window opens, as shown in Figure 7-48.

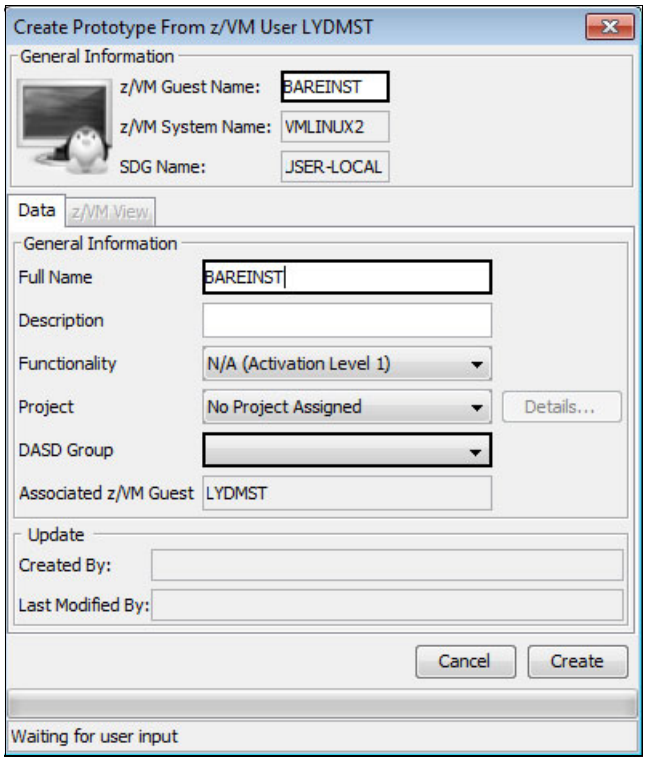


*Figure 7-48 Create Prototype from z/VM User*

The window includes the following fields:

- z/VM guest Name: This field has nothing to do with a z/VM guest. This value sets the name of the new prototype directory entry that IBM Wave instructs your directory manager to create. This name is restricted to eight characters.
- z/VM Full Name: Because the name that is passed to the directory manager is limited to eight characters, IBM Wave provides another field to allow for a more descriptive name. This field is stored in the IBM Wave knowledge base as metadata.
- Description: This field allows for more information about the prototype and might be used to describe software packages and versions that are associated with this prototype. Because clones from an IBM Wave prototype duplicate the Linux user names and passwords of the original prototype, this field also might list the user names and passwords on the machine.
- Project: This field allows for the prototype to be associated with a default project. Although this field is optional, IBM Wave presents an error message when a user attempts to clone from a prototype that lacks a default project.
- DASD Group: This field selects the storage group that is to be used for creating minidisks during cloning operations from this prototype.

When the fields are complete, click **Create**. A work unit is dispatched to the BTS, as shown in Figure 7-49 on page 253.

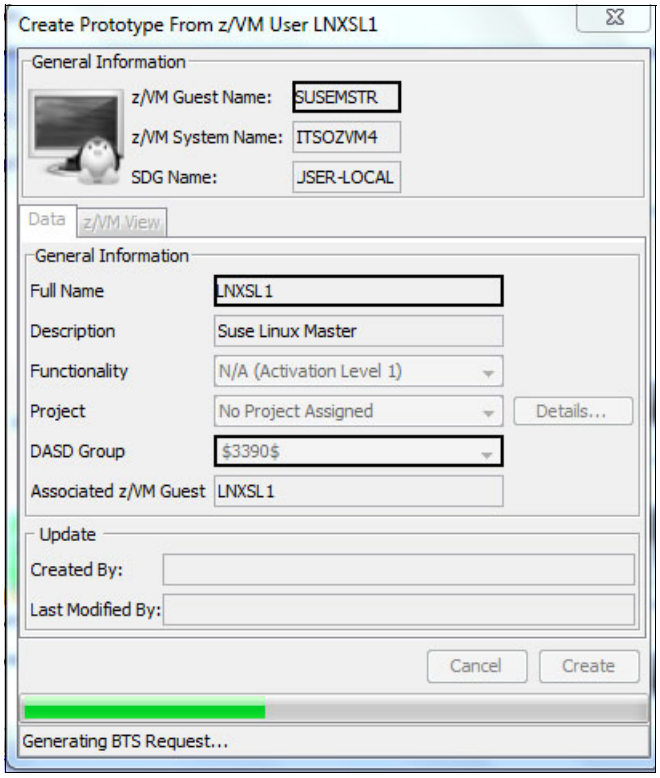


*Figure 7-49 The Create Prototype from z/VM user window*

After the prototype is created, the z/VM guest that was used to create the prototype disappears from the z/VM User Groups tab and the associated prototype appears in the Prototypes tab, as shown in Figure 7-50.

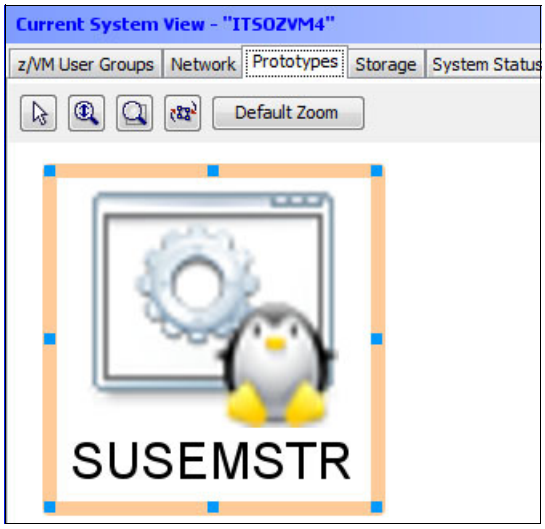


*Figure 7-50 New IBM Wave associated prototype*

Double-click the prototype to display information about the prototype on the Data tab, as shown in Figure 7-51.

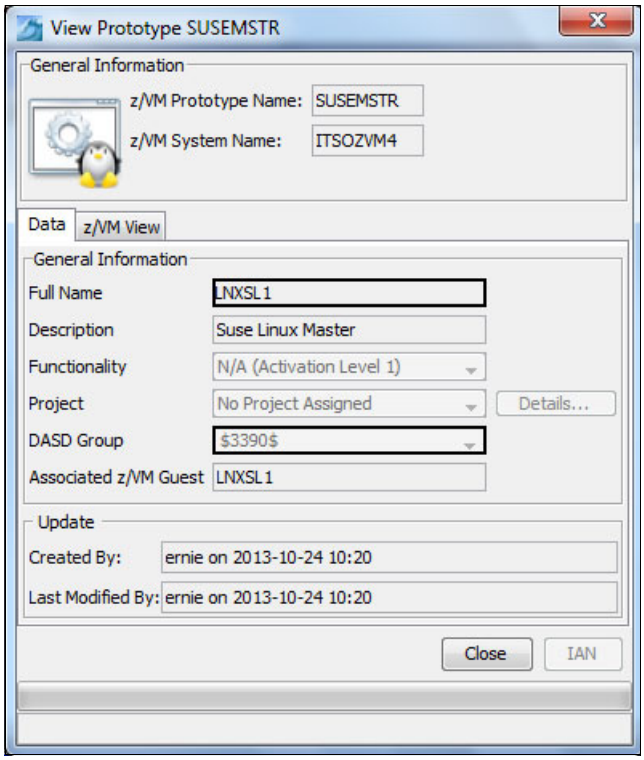


*Figure 7-51 VIew Prototype window*

The information about the prototype is nearly identical to the directory entry of the original guest. The following differences are featured:

- The directory entry is renamed from the name of the original guest to the one to eight character string that is specified in the z/VM field from the Create Prototype From z/VM User window.
- The MDisk statements are changed into directory manager commands to Add Minidisk (AMD) with **AUTOG** keywords that request that the directory manager dynamically create the 3390 DASD from the storage group that was specified in the Create Prototype From z/VM User window.

The z/VM View tab of the View Prototype window displays the prototype directory entry that is associated with the IBM Wave prototype, as shown in Figure 7-52 on page 255.

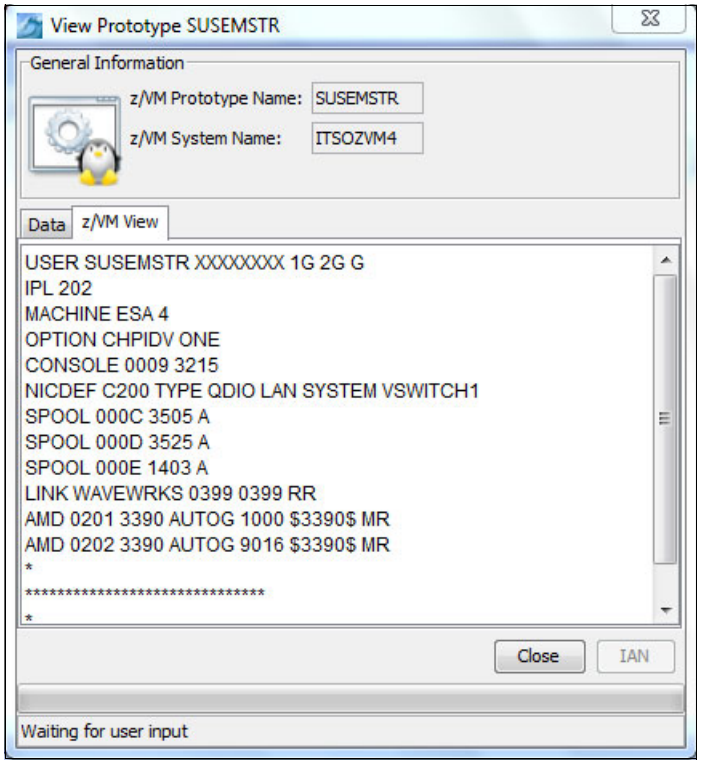


*Figure 7-52 Prototype directory entry*

## 7.5 Cloning virtual machines from a golden master

To clone a VM from a golden master, complete the following steps:

1. Right-click and select **Clone From This Prototype**, as shown in Figure 7-53.

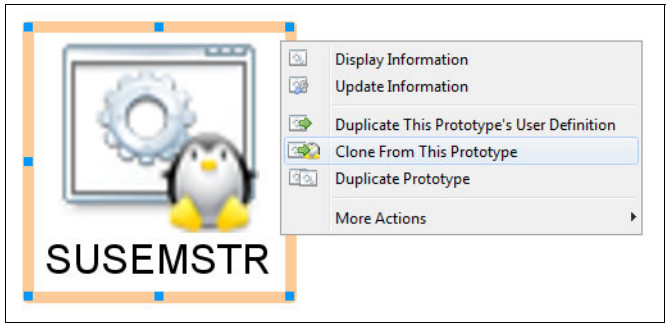


*Figure 7-53 Clone from this prototype*

If your prototype does not have a Default Project that is assigned, you receive an error message, as shown in Figure 7-54.

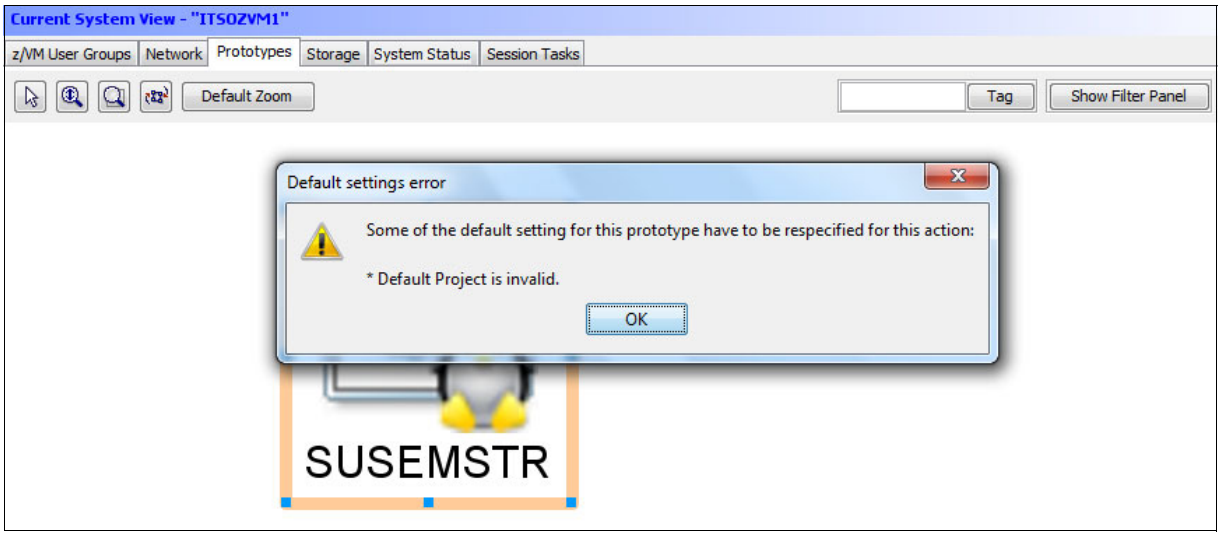


*Figure 7-54 Error message*

The Cloning from z/VM Prototype window opens, as shown in Figure 7-55 on page 257.

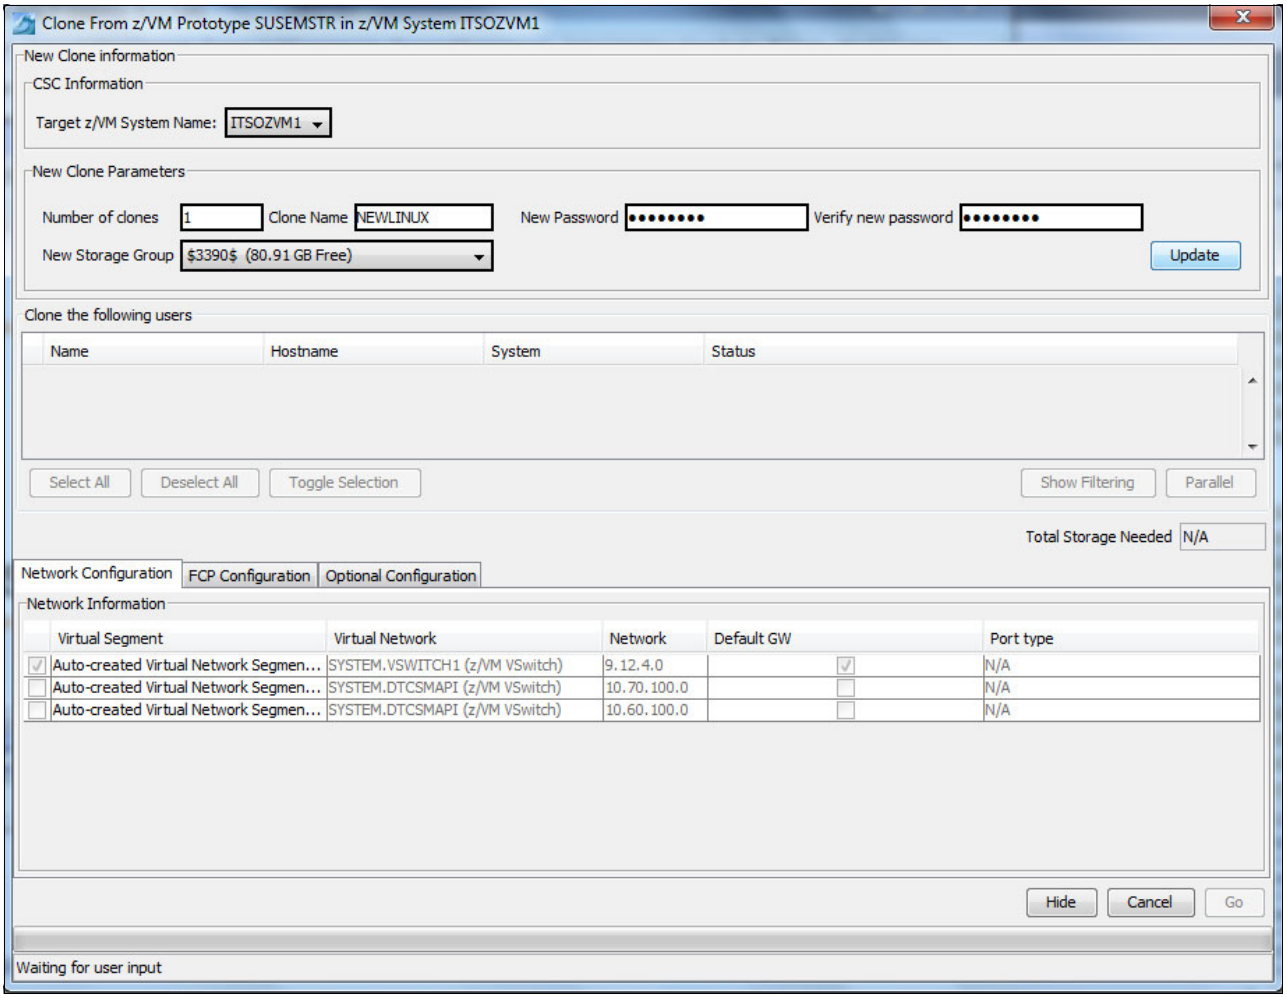


*Figure 7-55 Cloning from z/VM Prototype window*

Complete the following steps in the Cloning from z/VM Prototype window:

a. Select the Target z/VM system Name.

b. Enter the number of clones that you want.

c. Enter the base name for the clones. If multiple clones are to be created, this name is the base name of the clone and each clone has the base and a number.

d. Enter a password and verify the password.

e. Click **Update**. The Clone the following users table in the middle of the window is completed, as shown in Figure 7-56. An IP address is suggested. If you have multiple clones, the Virtual Network Segment (VNS) column is completed with suggested IP addresses on that VNS's IP segment.

   The IP address can be changed by double-clicking it. A window opens, as shown in Figure 7-56. Double-click the IP address, edit the field by typing your change, and then click **IP** to change it in the table.

f. Click **Go** to submit the clone request to the BTS.

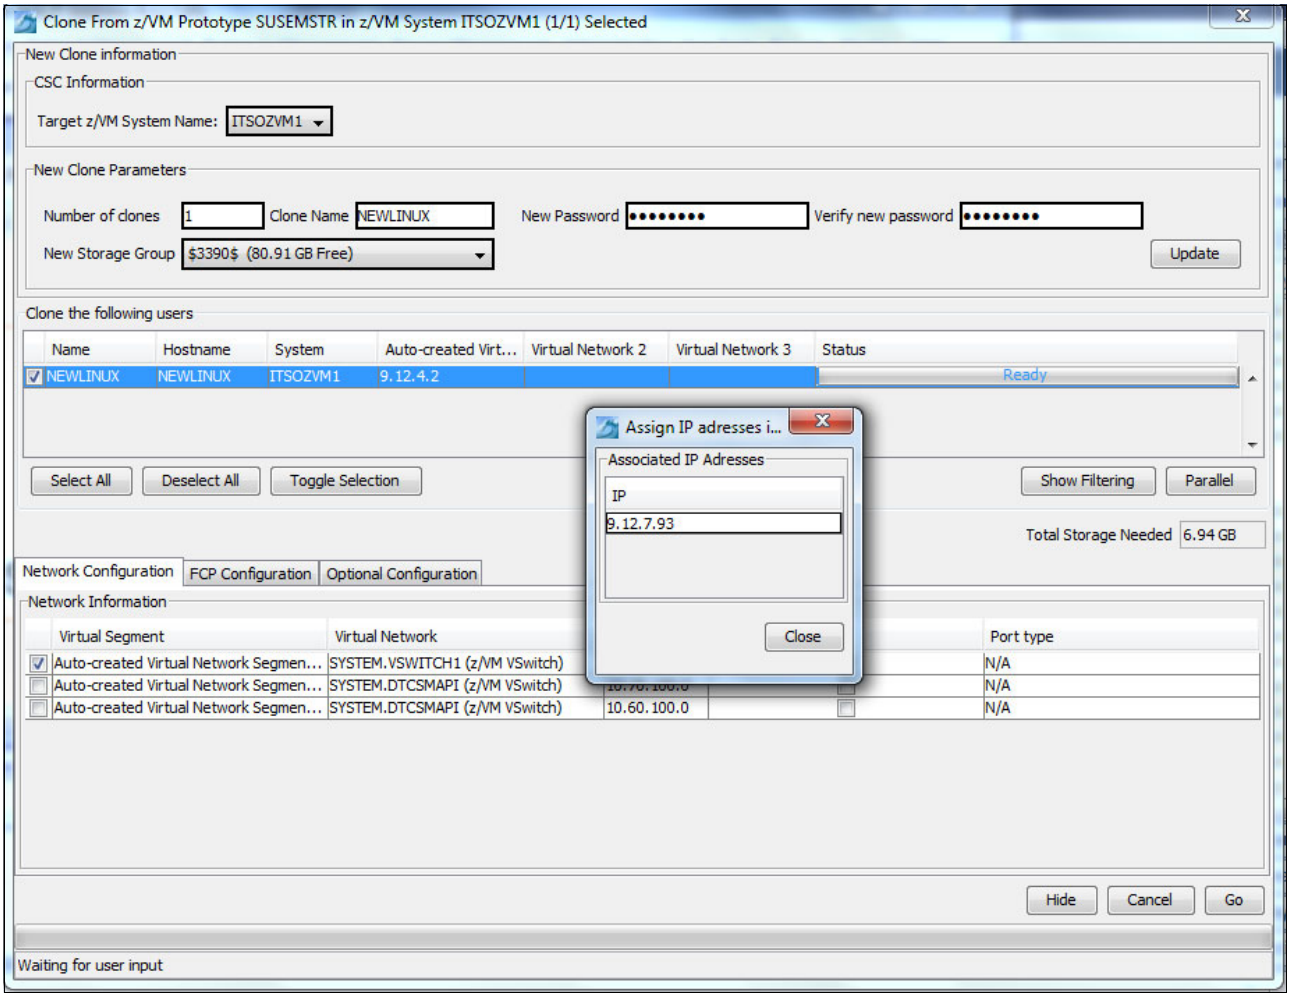


*Figure 7-56 Clone from z/VM prototype*

The clones appear in the z/VM User Group window under USER-LOCAL, as shown in Figure 7-57 on page 259.

The icons have a red circle with a white slash over them, which indicates that the WAVESRV Linux server does not have IP connectivity to the z/VM guest (Figure 7-57 on page 259).

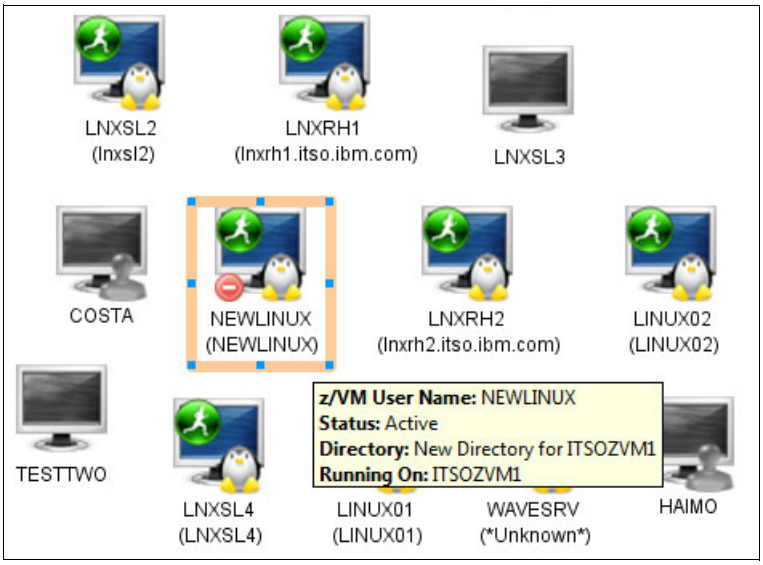


*Figure 7-57 No connectivity indicator on new Linux guest*

Connectivity is live (Figure 7-58) and you can connect through an external ASCII terminal emulator.

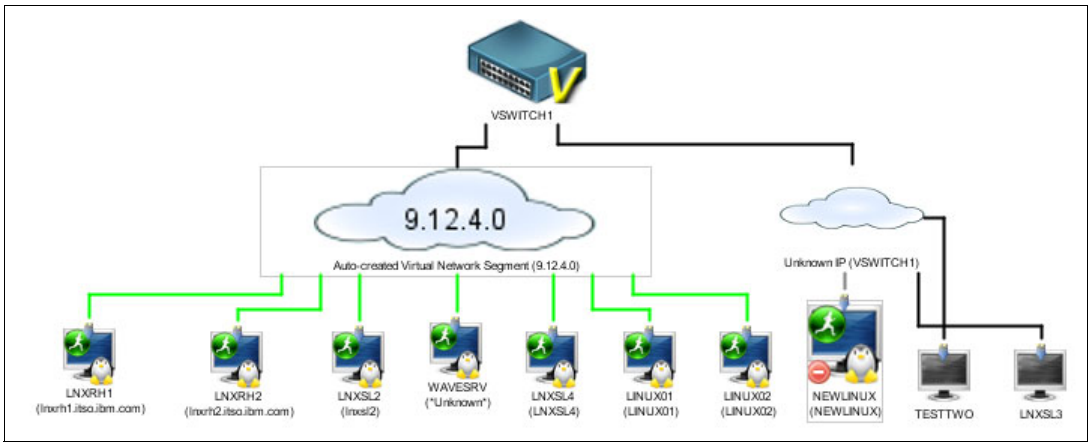


*Figure 7-58 Verify the network connectivity*

2. The red circle with a white bar on the guest is a No Connectivity icon that indicates that the IBM Wave metadata must be refreshed. Click **Auto Detect** → **Refresh** → **Schedule z/VM Network Update**, as shown in Figure 7-59.

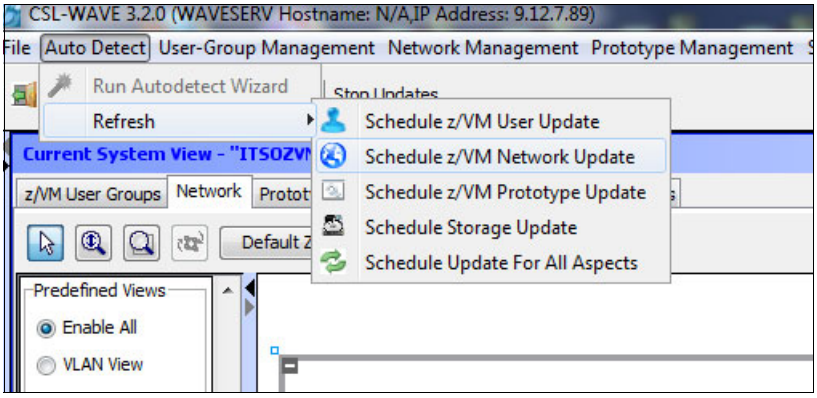


*Figure 7-59 Update the z/VM network metadata*

After the metadata is refreshed, you can see that the red circle with the white slash is gone (Figure 7-60), which indicates that the network is now connectable.

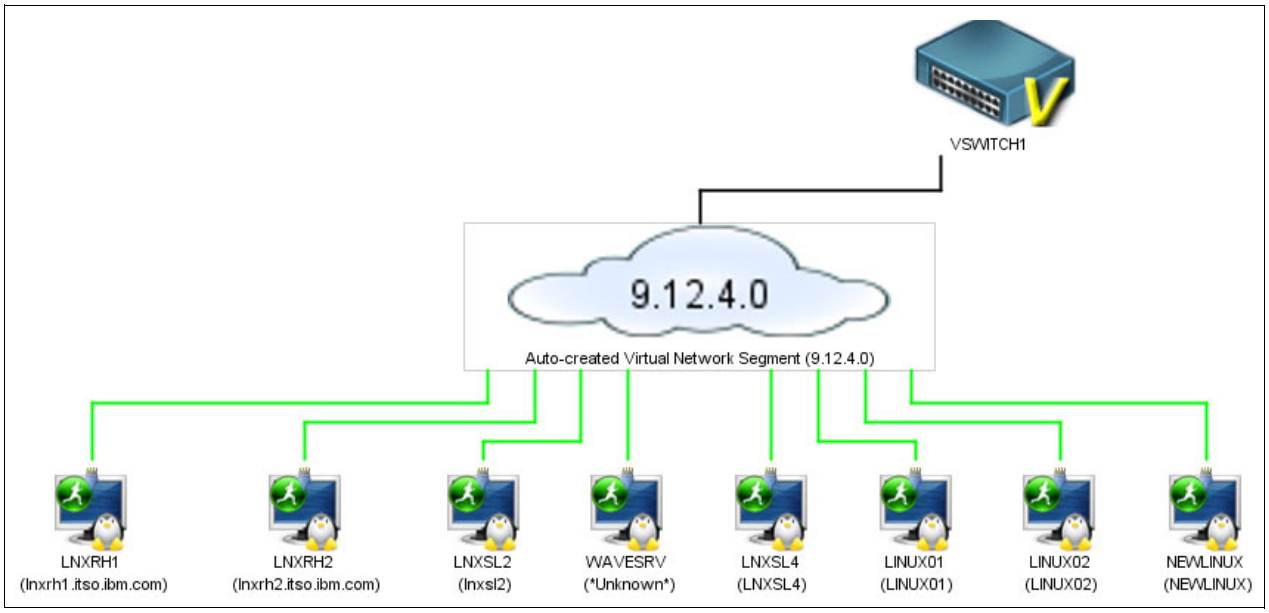


*Figure 7-60 Network connectable*

8

# IBM Wave for z/VM parameters

This chapter provides information about various Background Task Scheduler (BTS) parameters in IBM Wave that influence performance and resource consumption.

The BTS does all the work in IBM Wave for z/VM (IBM Wave). It schedules commands and runs regular background checks. In this chapter, those parameters might be set to tune performance for IBM Wave, but also reduce resource consumption. IBM Wave does not use many resources. Most resources are used by administrator actions, and those resources are used in a shell / command-line interface (CLI) driven environment. IBM Wave does not set any hooks anywhere, but polls periodically to get information about the environment. The administrator also can do this updating. This task is similar to a series of commands a systems programmer might run to check the health of the system.

This chapter includes the following topics:

- IBM Wave Parameters window
- General parameters
- BTS Manager window
- Restarting the BTS
- Dumping the BTS

# 8.1 IBM Wave Parameters window

The Manage Parameters action is used to configure a wide variety of system-wide IBM Wave settings. To use action, click **Administrative** → **Manage Parameters** (Figure 8-1).

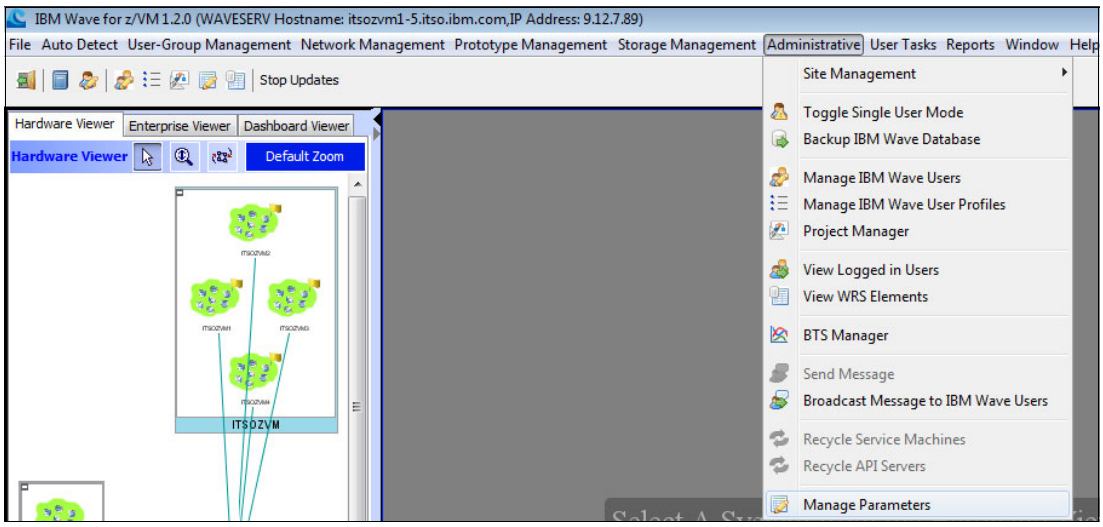


*Figure 8-1 Administrative - Manage Parameters*

The IBM Wave Parameters window opens (Figure 8-2). Click the **BTS** tab.

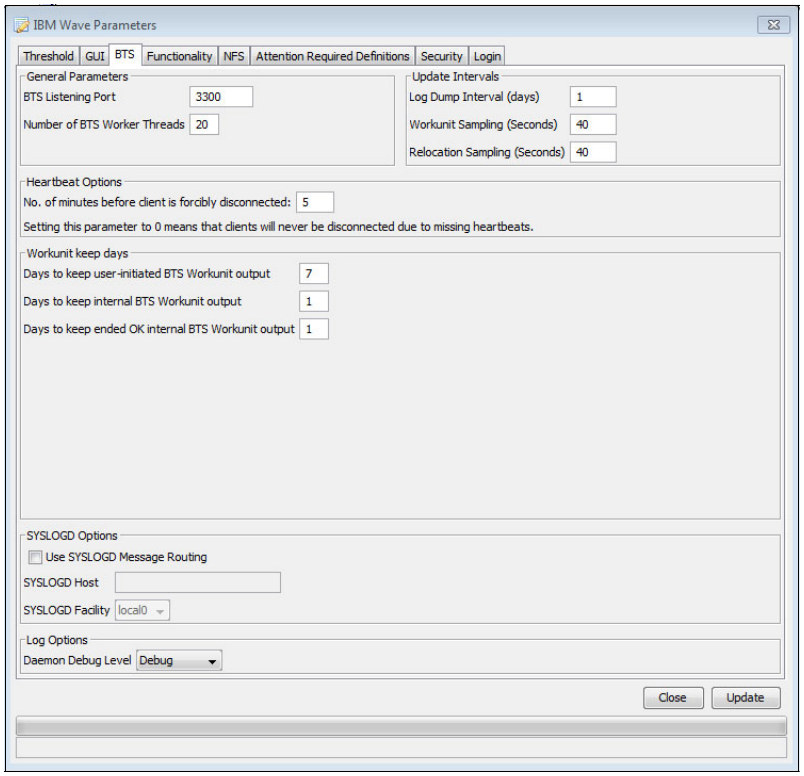


*Figure 8-2 IBM Wave Parameters window - BTS tab*

## 8.2 General parameters

Here are some of the general parameters:

- BTS Listening Port is a setting that you do not change unless your network administrator instructs you otherwise.
- Number of BTS Worker Threads defines the amount of parallel work that can be done by the BTS. A good starting value is 20. You can adjust the amount, but it is better to overprovision than underprovision for this setting. If you have many concurrent active IBM Wave users and do many concurrent cloning actions, consider doubling the value.

  To determine the number of user worker threads, use the following formula:

  <Number of concurrent IBM Wave Users>*2 + <Number of concurrent clone operations>*<maximum of minidisks in a source guest for cloning>

  There are also internal worker threads that run the internal IBM Wave work units and requests, such as the periodic tasks. These threads are not customizable and are determined by the number of z/VM systems that are managed by the BTS.

### 8.2.1 Update intervals

The Workunit Sampling setting by default has a value of 30. Work unit sampling occurs only when there are active work units that run in the z/VM system. Long running work units are created by actions such as copying or deleting z/VM guests. Keep this parameter around 30 or 40 seconds.

Also, keep the Relocation Sampling setting around 30 to 40 seconds, which tracks the long running relocation actions.

### 8.2.2 BTS work units

Every request that is running in the BTS runs under a BTS work unit. BTS work units can be viewed from the BTS Work Units viewer in the GUI client, as shown in Figure 8-3.

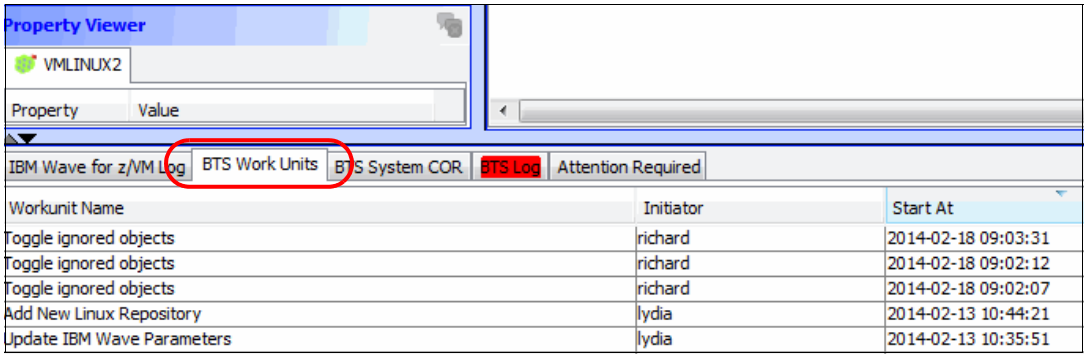


*Figure 8-3 BTS Work Unit viewer*

A BTS work unit is composed of one or more BTS requests. All output that is generated by the BTS and all BTS requests are stored in the IBM Wave common output repository (COR). The COR is made up of the log COR entry, which is used to log all activity of a specific BTS request. Each BTS request is allocated at least one COR entry (a Log COR Entry) upon initiation, which logs all the activity of the BTS request. The Log COR Entry for a specific request can be viewed by accessing the BTS work unit to which this BTS request belongs and selecting the specific request. The log COR entry also has System COR entries that can be viewed from the BTS System COR entry viewer in the GUI, as shown in Figure 8-4.

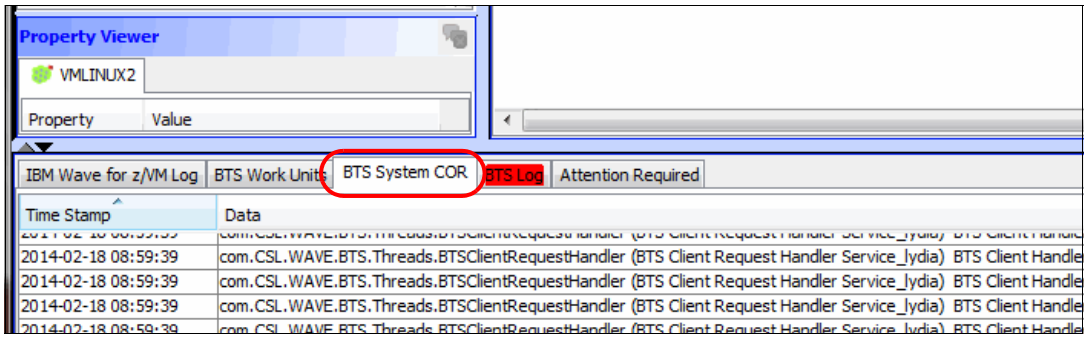


*Figure 8-4 BTS System COR entry viewer*

Other Log COR Entries can be viewed from the BTS work units Details window, which is accessed by double-clicking a specific BTS work unit and then selecting a specific BTS request. You can manually delete COR entries from the BTS work unit Details tab. When a BTS work unit is deleted from IBM Wave, COR entries that are owned by the BTS request are also automatically deleted.

## 8.3 BTS Manager window

You can use the BTS Manager window to configure the settings for the BTS and store all BTS logs. To open the BTS Manager window, click **Administrative** → **BTS Manager**.

### 8.3.1 General Information

Figure 8-5 on page 265 shows the connected clients and some of the internal task workload. You can toggle between User Worker Stats and Internal Worker Stats. If you have a long-running task, you can drill down to obtain information about the status by double-clicking the task.

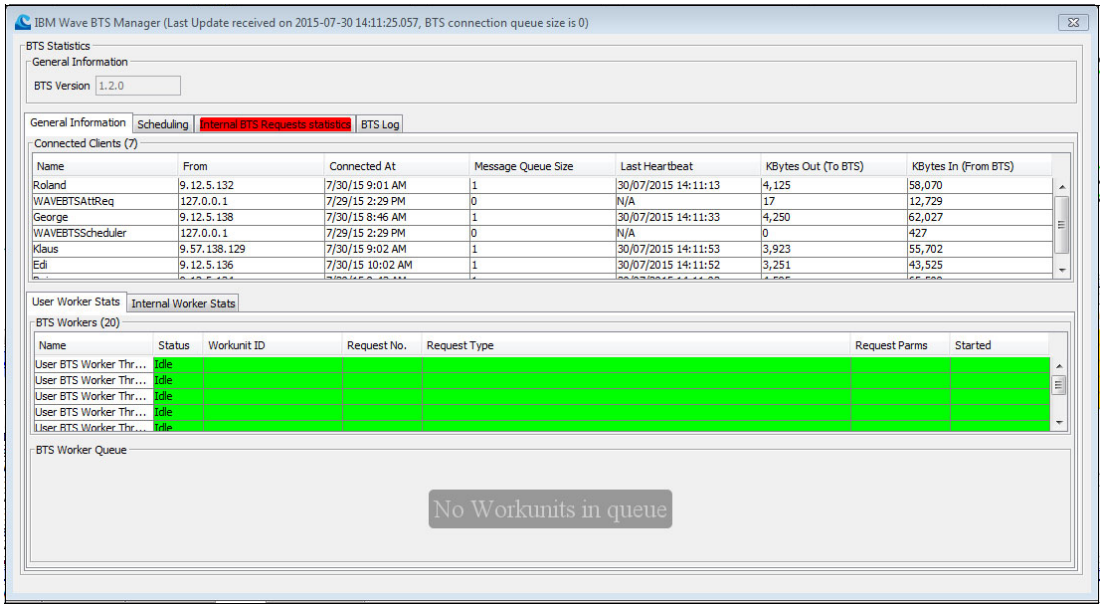


*Figure 8-5 BTS Manager - General Information*

## 8.3.2 Scheduling

The second tab in BTS Manager (Figure 8-6) displays scheduling. You also can use this tab to do parameter management.

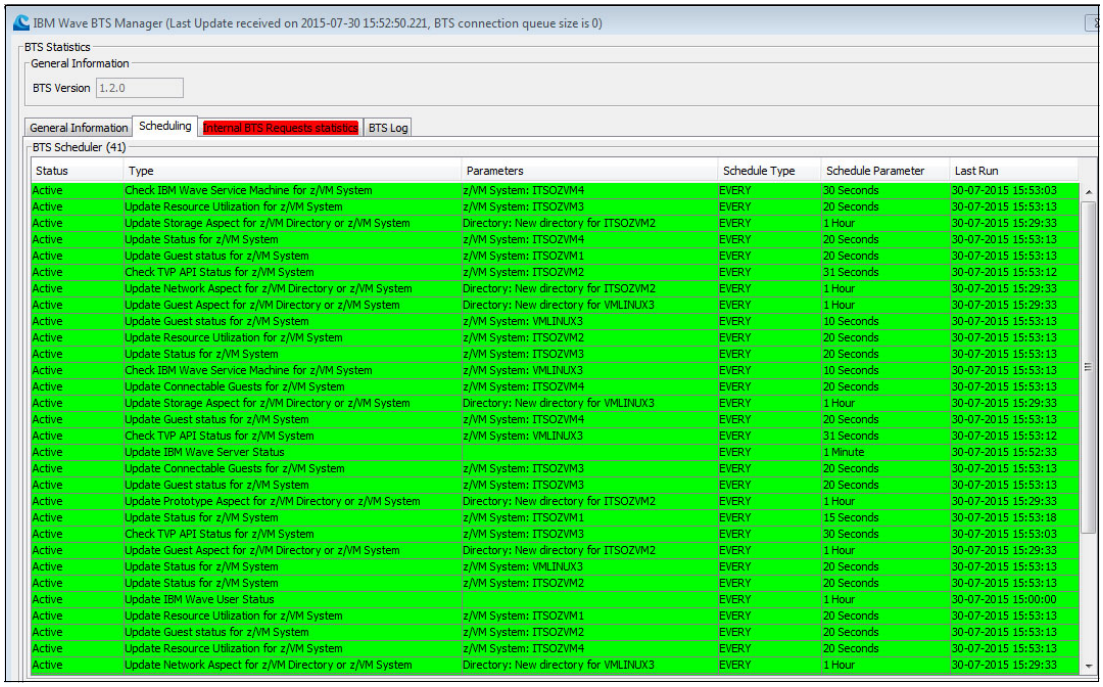


*Figure 8-6 BTS Manager - Scheduling*

This example shows that resource utilization for every managed system is updated every 20 seconds. If you feel that this is too often, double-click the Schedule Parameters in the appropriate row and type in a new value, as shown in Figure 8-7 and Figure 8-8, where we change the schedule update from 20 seconds to 44 seconds.

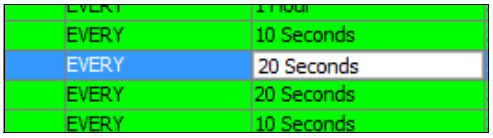


*Figure 8-7 Schedule Parameters*

Update

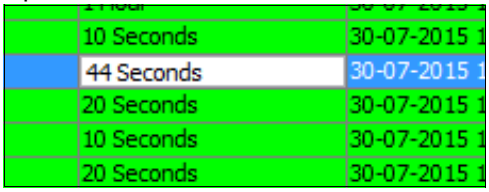


*Figure 8-8 Schedule parameters update*

You can also deactivate the tasks (Figure 8-9) or run a **`RUN NOW`**. *Deactivating the tasks can have unintended results*.

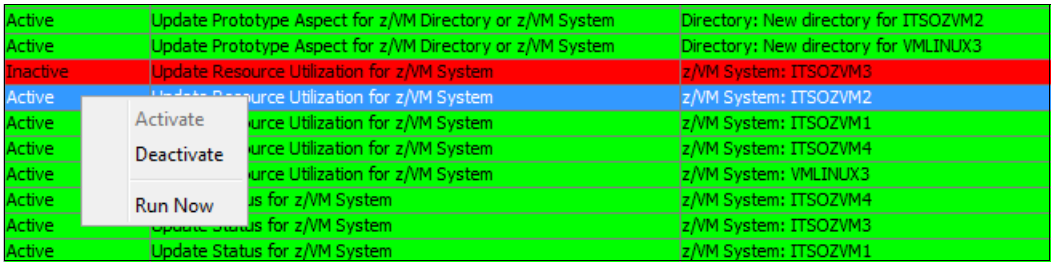


*Figure 8-9 Deactivate tasks*

For more information about the BTS Scheduler, see *IBM Wave for z/VM: Administration and Customization*, SC27-6118.

## 8.3.3 Internal BTS Request Statistics

The tab that is shown in Figure 8-10 on page 267 provides information about how internal requests are running, how long they take, and what they are doing. You can drill down on each of these items by double-clicking the task.

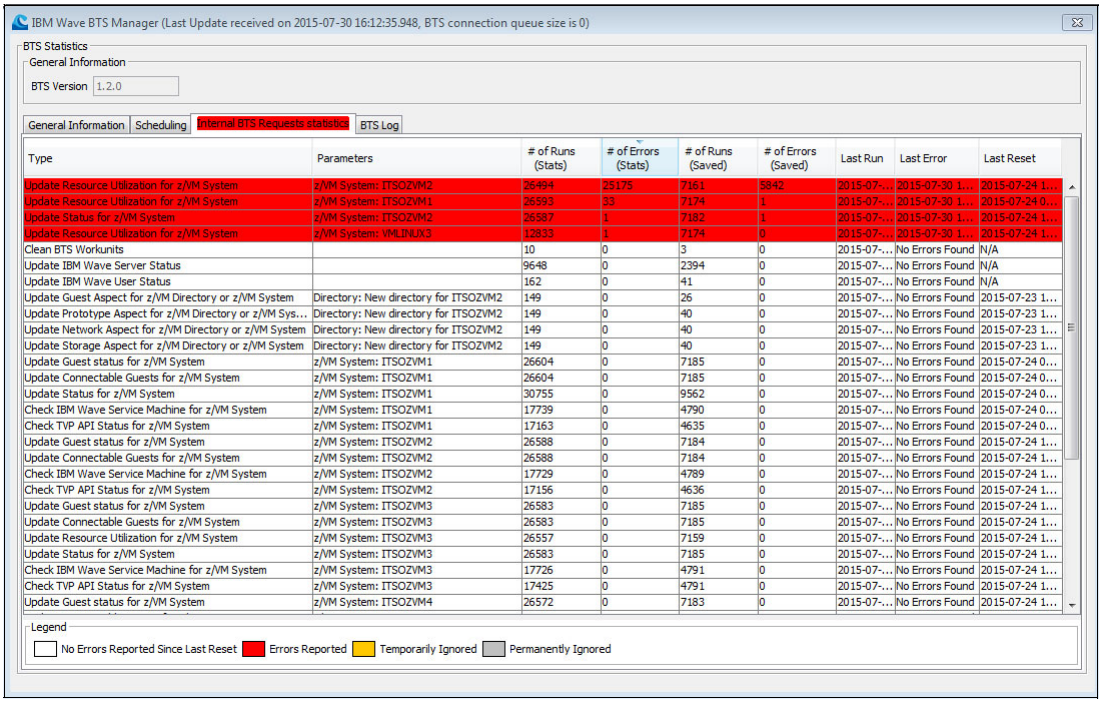


*Figure 8-10 Internal BTS Request Statistics*

You also can take several other actions on tasks in this table. You can reset statistics, toggle **Ignore Until Next Error Occurs**, toggle **Ignore Permanently**, or clean work units for stats, as shown in Figure 8-11.

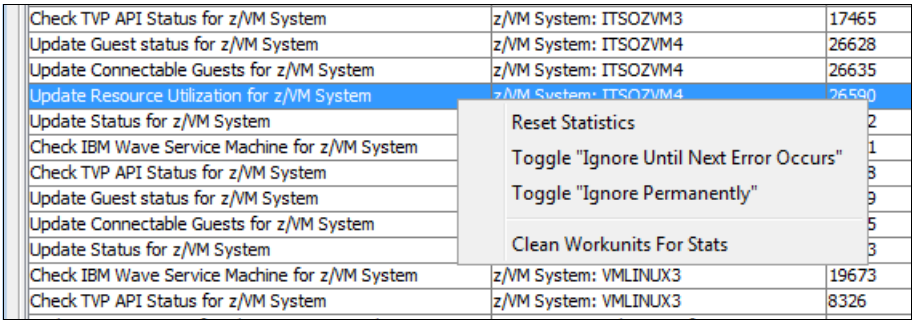


*Figure 8-11 Modify entries in Internal BTS Request Statistics*

You might decide that you want to keep only the last 10 entries of this event in the statistics. To do so, select **Clean Workunits for Stats**, as shown in Figure 8-12, select **Keep last __ Workunits**, and type 10 into its field.

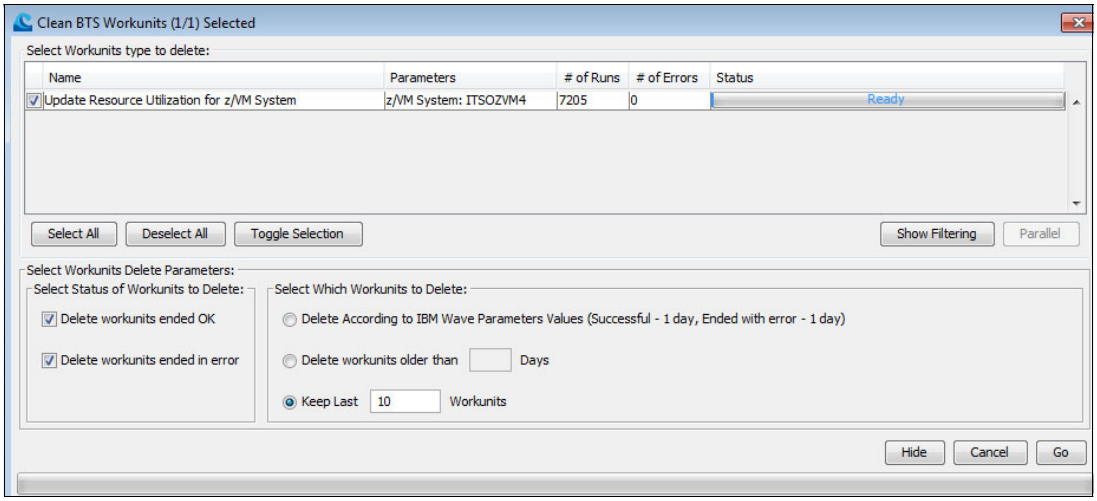


*Figure 8-12 Delete Entries - keep last 10*

You also might decide to delete only the work units in the Error field (after you correct the error).

## 8.4 Restarting the BTS

If you want to restart the BTS but do not want to restart the whole IBM Wave server, log in to the WAVESRV Linux guest through SSH and run the following commands:

- **`/etc/init.d/WAVEBackgroundServices stop`**
- **`/etc/init.d/WAVEBackgroundServices start`**

## 8.5 Dumping the BTS

If a dump of the BTS is required, log in to the WAVESRV Linux guest through SSH and run the following command:

```
/usr/wave/WAVEBackGround/WAVEDatacapture
```

This command generates a dump into `/tmp` called `sendToIBM*.tar.gz`, where * is a time stamp.

A

# Miscellaneous

This appendix contains examples, descriptions, or references to other publications that are relevant to the IBM Wave installation or usage.

This appendix includes the following topics:

- IBM Wave for z/VM fix pack installation
- Customization with other security managers
- Wave CLI helper file
- Wave flow charts

# IBM Wave for z/VM fix pack installation

This section provides short versions of the IBM Wave for z/VM (IBM Wave) installation readme files. The information in the readme files is not cumulative.

## IBM Wave V1R2 Fix Pack 0.5 readme file

There are several fixes. There are some new functions.

### New functions

New capabilities in IBM Wave V1R2 provide enhanced reporting and increased support for Linux distributions:

- Enhancements to the reporting functions include the following items:
  - Two new reports: Custom Attribute Report and FCP Connections Report.
  - Users can generate customized reports of resources on demand.
  - Report templates are more comprehensive and you can manually select which object attributes on which to report.
  - Reports provide more attributes for the various reporting resources so that reports can be generated with more meaningful data.
  - Reports are more customizable and allow custom attributes to be specified and reported.
  - Users can export reports to an external format by using the Export to CSV option so that they can view, manipulate, print, and work with data outside of IBM Wave.
  - Report management flows are more usable.

  Reporting features that include new attributes and resources continue to follow scopes and permissions as they are implemented for reporting currently. For example, if a user does not have permission to access DASD, they cannot view DASD data in a report.
- Initial support for the latest Linux distributions from Red Hat and SUSE

The following Linux operating system (OS) distributions are fully supported:

- Red Hat Enterprise Linux Server 5 (RHEL 5)
- Red Hat Enterprise Linux Server 6 (RHEL 6)
- SUSE Linux Enterprise Server 10
- SUSE Linux Enterprise Server 11

The following Linux OS distributions are partially supported:

- Red Hat Enterprise Linux Server 7 (RHEL 7)
- SUSE Linux Enterprise Server 12

The Linux WAVESRV must have one of the following OSes, as it is not supported on RHEL7 or SUSE Linux Enterprise Server 12 currently:

- Red Hat Enterprise Linux Server 5 + dependant packages
- Red Hat Enterprise Linux Server 6 + dependant packages
- SUSE Linux Enterprise Server 11+ dependant packages

When you clone RHEL7 and SUSE Linux Enterprise Server 12 guests, consider the following items:

- In Red Hat Enterprise Linux Server 7 and SUSE Linux Enterprise Server 12, there are some changes that are made to the way that machine IDs are handled.
- After cloning a Red Hat Enterprise Linux Server 7 or a SUSE Linux Enterprise Server 12 guest with IBM Wave V1R2, run the following command manually on the new guest and then restart:

  ```
  echo "" > /etc/machine-id
  ```

  This command ensures that a unique machine ID is assigned to the new guest. this task will be handled programmatically in a future fix pack.

For Red Hat Enterprise Linux Server 7 and SUSE Linux Enterprise Server 12, the following actions are not supported and are disabled:

- Manage storage.
- Run scripts.
- Connect to a virtual network segment.
- Disconnect from a virtual network segment.
- Refresh Linux data.
- Relocate by using Live Guest Relocation (LGR)

Although the connect and disconnect actions are not disabled, the actions affect only the z/VM guest configuration (adding and removing **NICDEF** statements). IBM Wave cannot change the configuration of the OS running on the guest (such as bring the interface up or down or add configuration files).

## Publication information:

There are new IBM Wave publications that are available with IBM Wave for z/VM V1R2 Fix Pack 0.5. IBM Wave publications are available in IBM Knowledge Center, found at:

http://www.ibm.com/support/knowledgecenter/SS6JTX/welcome

PDF formats are available at the following websites:

- *Administration and Customization*:

  http://publibz.boulder.ibm.com/epubs/pdf/c2761186.pdf

- *User Guide and Reference*:

  http://publibz.boulder.ibm.com/epubs/pdf/c2761196.pdf

## Migration information

When you perform a migration, be sure to consider the following actions:

- Change the password for the WAVEWRKS Service Machine.
- This fix pack introduces new password rules for IBM Wave users. To enforce these rules, passwords that are not compliant are marked as expired, forcing the user to change the password the next time they log in.
- During the migration from IBM Wave V1.1.0 to V1.2.0, all Linux guests will be marked as partially inited for IBM Wave and the clone action will be disabled for them. To re-enable the clone action for a Linux guest, run the “Init User for IBM Wave” action on each Linux guest. Linux guests that are associated with a prototype must be disassociated from the prototype before the “Init User for IBM Wave” action is run on them, and then reassociated with the prototype afterward.

- Add the **LNKE** option to your WAVEWRKC USER DIRECT file so that the cloning of guests functions correctly.

To make this change, for each z/VM system that is managed by IBM Wave, complete the following steps:

1. Ensure that no IBM Wave work units are running.
2. Open a 3270 session to your system and log in as MAINT.
3. Run **FORCE WAVEWRKC**.
4. Run **DIRM FOR WAVEWRKC GET**.
5. Run **RLIST**.
6. Locate WAVEWRKC DIRECT and run **RECEIVE / (REPL**.
7. Exit RLIST.
8. Run **XEDIT WAVEWRKC DIRECT**.
9. Replace "OPTION LNKNOPAS DIAG88" with "OPTION LNKNOPAS LNKE DIAG88".
10. Save and quit XEDIT.
11. Run **DIRM FOR WAVEWRKC REPL**.
12. Run **XAUTOLOG WAVEWRKC**.
13. From the IBM Wave menu, click **Auto Detect** → **Refresh** → **Schedule z/VM User Update**.
14. Verify that the USER DIRECT file is updated to have LNKE in the WAVEWRKC guest in IBM Wave.

**Important information about stale image volume group data that can be cached by SMAPI:**

Starting with Fix Pack 10, customers that use IBM Wave and making z/VM configuration changes outside of IBM Wave were directed, as a preferred practice, to run with LOHCOST, the SMAPI database server, disabled. Since then, SMAPI development has produced an updated practice.

IBM Wave uses z/VM SMAPI in some of its communications with the z/VM managed system. If updates to image volume groups are being made from outside of IBM Wave (SMAPI), manually updating the `EXTENT CONTROL` file or by using DirMaint commands, the `LOHCOST_GROUP` cache in the LOHCOST server becomes stale. This situation results in SMAPI returning older IBM Wave data. IBM engineers opened z/VM RFE 69890 for DirMaint to notify SMAPI of such changes that are made outside of IBM Wave or SMAPI in general.

Based on the fact that the LOHCOST cache becomes stale only when image volume groups are updated from outside of IBM Wave (SMAPI), as preferred practice, customers should disable only caching of image volume groups, but continue to cache the user directory entries.

For example, on z/VM V6.3, to accomplish this task, find the following statement in the `DMSSICNF COPY` file (on the MAINT 193 disk):

```
LOHCOST_Enabled = LOHCOST_DIRECTORY + LOHCOST_GROUP
```

Change this statement to the following string:

```
LOHCOST_Enabled = LOHCOST_DIRECTORY
```

Then, restart SMAPI to make this change active. To restart SMAPI, run the following commands:

- **FORCE VSMGUARD**
- **XAUTOLOG VSMGUARD**

If you have any further questions about LOHCOST or running a different z/VM version and want to verify these instructions, see the z/*VM Systems Management Application Programming* manual for the version of z/VM that your installation is running.

## IBM Wave for z/VM V1R2 is available through Fix Pack 0.5.

This is the next fix pack that is delivered since IBM Wave V1R1 Fix Pack 10 and it migrates you to IBM Wave V1R2 Fix Pack 0.5. The IBM Wave fix packs are cumulative and contain all previous APAR fixes from prior fix packs.

This fix pack is being distributed as a delta to be applied to the installed IBM Wave for z/VM V1.1 RPM.

To apply this fix pack, complete the following steps:

1. This fix pack contains the `IBM-WAVE-1.2.0.00.5.tar` file. You must copy this file to the root directory of WAVESRV server. From there, you run the following command:

   ```
   tar -xf IBM-Wave-1.2.0.00.5.tar
   ```

   This command creates the `IBM-Wave-1.2.0.00.5` subdirectory.

2. Run the following command:

   ```
   cd IBM-Wave-1.2.0.00.5
   ```

3. In the IBM-Wave-1.2.0.00.5 subdirectory, find the `doUpdate.sh` script file, and then run the following command:

   ```
   ./doUpdate.sh
   ```

4. You might be prompted with following string:

   ```
   Would you like to update Service Machines at z/VM?
   ```

   Type Y and press Enter.

5. The following string appears:

   ```
   Enter the Wave short Service machine (WAVEWRKS) password (press enter to use
   default)
   ```

   Type the current WAVEWRKS password and press Enter.

6. The following string appears:

   ```
   Enter the Wave short Service machine multi-write password for the 191 minidisk
   ```

   Press Enter to use the default.

7. The following string appears:

   ```
   Enter the Wave short Service machine Multi-write password for the 399 minidisk
   ```

   Press Enter to use the default.

You should now see several files being updated. When the update is complete, you should see the message that is shown in Figure A-1.

```
-----------------------------------
|
| IBM Wave Successfully updated
|
-----------------------------------
```

*Figure A-1 IBM Wave successfully updated*

The installation of the fix pack is complete.

## IBM Wave for z/VM V1R1 Fix Pack 10 readme file

The IBM Wave fix packs are cumulative and contain all previous APAR fixes from prior fix packs.

IBM Wave publications are now available in the IBM Knowledge Center, and are available at the following website:

`http://www.ibm.com/support/knowledgecenter/SS6JTX/welcome`

PDF formats are available at the following websites:

- *Administration and Customization*:

  `http://publibz.boulder.ibm.com/epubs/pdf/c2761185.pdf`

- *User Guide and Reference*:

  `http://publibz.boulder.ibm.com/epubs/pdf/c2761195.pdf`

As a preferred practice, customers that use IBM Wave and make z/VM configuration changes outside of IBM Wave should run with LOHCOST, the SMAPI database server, disabled. For instructions about how to disable SMAPI LOHCOST, see the z/*VM Systems Management Application Programming* manual for the version of z/VM that your installation is running.

For example, on z/VM V6.3, to disable LOHCOST caching of directory data and directory manager storage group data, set `LOHCOST_Enabled = 0` in the Server Configuration file.

IBM Wave uses z/VM SMAPI in some of its communications with the z/VM managed system. If updates to directory, minidisks, or extent control groups are being made from outside of IBM Wave, such as with DIRM commands, caches that are included in LOHCOST can become stale. This situation can result in SMAPI returning older IBM Wave older.

Recycling the SMAPI server (by clicking **IBM Wave** → **Administrative** → **Recycle API Servers**) after those updates or running with SMAPI LOHCOST turned off is a preferred practice.

This fix pack is being distributed as a delta to be applied to the installed IBM Wave for z/VM V1.1 RPM.

To apply this fix pack, complete the following steps:

1. This fix pack contains the `IBM-Wave-1.1.0.10.tar` file. You must copy this file to the root directory of the WAVESRV server. From there, run the following command:

   ```
   tar -xf IBM-Wave-1.1.0.10.tar
   ```

   This command creates the following subdirectory:

   ```
   IBM-Wave-1.1.0.10
   ```

2. Run the following command:

   ```
   cd IBM-Wave-1.1.0.10
   ```

3. In this subdirectory, you find a script file that is called `doUpdate.sh`. From this directory, run the following command:

   ```
   ./doUpdate.sh
   ```

4. You might be prompted with following string:

   ```
   Would you like to update Service Machines at z/VM?
   ```

   Type Y and press Enter.

5. The following string appears:

   ```
   Enter WAVE short Service machine (WAVEWRKS) password (press enter to use
   default)
   ```

   Type the current WAVEWRKS password and press Enter.

6. The following string appears:

   ```
   Enter WAVE short Service machine multi-write password for the 191 minidisk
   (press enter to use default)
   ```

   Press Enter to use the default.

7. The following string appears:

   ```
   Enter WAVE short Service machine multi-write password for the 399 minidisk
   (press enter to use default)
   ```

   Press Enter to use the default.

You should now see several files being updated. When the update is complete, you should see the message that is shown in Figure A-2


*Figure A-2 IBM Wave successfully updated*

The installation of the fix pack is complete.

## IBM Wave for z/VM V1R1 Fix Pack 9 readme file

There are several fixes. There are no new functions.

## IBM Wave for z/VM V1R1 Fix Pack 8.1 readme file

There are several fixes. There are some new functions.

### New functions

IBM Wave continues to support the industry standard TLS protocol and disables the usage of SSL.

With this fix pack installed, references to SSL in the GUI indicate that the TLS protocol is in use.

**Notes:** This fix pack should be installed only if you meet *all* of the following prerequisites:

- If you are using LDAP support for logging on to IBM Wave and the LDAP server is configured for SSL communication, your instance of LDAP must support the TLS protocols.
- If any of your managed z/VMs are configured for SSL communication, each of those z/VM instances must support the TLS protocols for the SMAPI, IBM Wave Service Machines, and 3270 ports.

These prerequisites can be tested by running the following command:

```
openssl s_client -connect XXX.XXX.XXX.XXX:Port# | egrep 'Protocol'
```

`XXX.XXX.XXX.XXX` is the IP address of your LDAP Server or your z/VM system, and `Port#` is the port number to test. This command will output TLS if your system is configured to support TLS.

If your system uses SSL and does not support TLS, contact your IBM Systems Service Representative for more information about this and future fix packs.

## IBM Wave for z/VM V1R1 Fix Pack 8 readme file

There are several fixes. There are some new functions.

### New functions

An IBM Wave user with the correct level of authorization can create, modify, and delete emulated devices (EDEVs) through the IBM Wave GUI client. Details about this support are in the “Storage Functions” section in the latest refresh of the *IBM Wave for z/VM User Guide and Reference*, found at:

`http://publibz.boulder.ibm.com/epubs/pdf/c2761193.pdf`

After installing this fix pack, any z/VM systems that are managed by IBM Wave must be auto-detected again in order for EDEV definitions on those systems to be made persistent after an IPL. To do the auto-detect again, in the Hardware Viewer, right-click the z/VM system, select **Update Details**, and then set the system status to Suspend (suspending it from IBM Wave activity). With the system still selected, click the **Auto Detect** menu and select **Run Auto Detect Wizard**.

To ensure that this persistency is being managed, check the systems' `AUTOLOG` profile `EXEC` for a reference to the following profile name:

```
EDEVPROF
```

If this reference is not present, the following command must be added to the `AUTOLOG` profile `EXEC` in order for the managed EDEV definitions to be persistent after an IPL:

```
Address Command 'EXEC WAVEAUTR EDEVPROF'
```

The same steps can be carried out to ensure the placement of the LAN Define (`LANPROF`), LAN Grant (`GRNTPROF`), and XAUTOLOG (`ACTPROF`) profiles, with the following commands:

- **`Address Command 'EXEC WAVEAUTR LANPROF'`**
- **`Address Command 'EXEC WAVEAUTR GRNTPROF'`**
- **`Address Command 'EXEC WAVEAUTR ACTPROF'`**

**Notes:** The WAVEWRKC Service Machine must be authorized to link read/write to target clone minidisks. In a system with the RACF security product, this action requires granting the WAVEWRKC OPERATIONS privilege.

VMSecure customers should make sure that they are at least at VMSecure Version 3.1. They should install fix pack R070842, along with individual fix R072116.

In addition to being current on z/VM service (that is, having the latest RSU), the z/VM V6.3 PTF UM34369 (APAR VM65578) should be installed. When this PTF is installed, a local modification should be made to the `DMSSICNF COPY` file (for SMAPI) to specify `LOHCOST_ENABLED=4`. The PTF and local modification apply only to z/VM V6.3 systems.

Init for Wave might leave guests without FCP devices showing as “This user is partially initialized for IBM Wave use”. This should not cause any functional issues.

**Note:** When migrating from Fix Pack 3 or earlier for FCP devices, you might need to delete and re-add the z/VM system in IBM Wave to free the FCP devices within IBM Wave.

## IBM Wave for z/VM V1R1 Fix Pack 7 readme file

There are several fixes. There are some new functions. The same information that is found in “IBM Wave for z/VM V1R1 Fix Pack 8 readme file” on page 277 is applied to this fix pack as well.

## IBM Wave for z/VM V1R1 Fix Pack 6 readme file

There are several fixes. There are some new functions.

### New functions

The WAVEWRKC Service Machine must be authorized to link read/write to target clone minidisks. In a system with the RACF security product, this requires granting the WAVEWRKC OPERATIONS privilege.

With Fix Pack 6, restrictions on Cross-System Cloning are lifted. Cross System-Cloning is supported on z/VM V 6.2 and later.

The following LDAP interface enhancements were made:

- The z/VM API Testing Application now includes an LDAP test utility.
- The group search attribute was changed to use a drop-down menu with potential attributes to use
- The LDAP directory search was changed to search initially for a returned user attribute as defined by the group search attribute.
- Group membership search was enabled by specific attribute, not just by member.
- SSL certificate was enabled to be a single file or directory.

To migrate to APAR LI78091 for each managed VM system, complete the following steps:

1. Change the Authorized API User password.
2. Change the read, write, and multi-passwords for the IBM Wave Service Machine 191 and 399 minidisks on both the VM side and the GUI (click **Administrative** → **Site Management** → **Update Minidisk Passwords**.)
3. Rerun the “Init for IBM Wave” action on all your managed guests.
4. Regenerate the host keys for cloned guests that were created before Fix Pack 6.

If a z/VM user/administrator is using SSH to access the IP addresses that are associated with the existing clones, those IP addresses and their previous SSH public keys are stored in the user's `known_hosts` file on their local system. When using SSH to access the clones with the regenerated SSH keys, SSH displays a message stating that the public keys of the guest are not the ones that are listed in the user's `known_hosts` file. To remove the old keys that are associated with the IP, the user runs **`ssh-keygen -R hostname`** in the Linux terminal, where **`hostname`** is the IP address.

**Note:** VMSecure customers should make sure that they are at least at VMSecure V3.1. They should install fix pack R067338 along with individual fixes R067266, R070893, and R072116. In addition to being current on z/VM service (that is, having the latest RSU), the z/VM V6.3 PTF UM34369 (APAR VM65578) should be installed. When this PTF is installed, a local modification should be made to the `DMSSICNF COPY` file (for SMAPI) to specify `LOHCOST_ENABLED=4`. The PTF and local modification apply to only z/VM V6.3 systems.

**Note:** As of Service Pack 6, under the "Optional Configuration" tab of the Cloning frame, the Regenerate SSH keys check box defaults to a checked state for Linux guests that are designated as such in IBM Wave. With this box checked, the RSA, RSA2, and DSA public and private SSH keys are regenerated and are not identical to the keys of the guest that is the origin of the clone. This check box can be cleared by the user, resulting in matching public and private SSH keys between the guest that is the origin of the clone and the resulting guest that is a clone. When creating multiple guest clones from one original guest at once, selecting the Regenerate SSH keys check box causes each clone to have unique public and private SSH keys.

**Note:** In Chapter 2, "Installatiom and Customization", of *IBM Wave for z/VM: Administration and Customization Version 1 Release 2*, SC27-6118-08, when you change the `AUTHFOR CONTROL` file, complete the following steps:

1. Add the following statements to the `AUTHFOR CONTROL` file:

```
ALL VSMGUARD * 140A ADGHMOPS
ALL VSMGUARD * 150A ADGHMOPS
ALL VSMWORK1 * 140A ADGHMOPS
ALL VSMWORK1 * 150A ADGHMOPS
ALL VSMWORK2 * 140A ADGHMOPS
ALL VSMWORK2 * 150A ADGHMOPS
ALL VSMWORK3 * 140A ADGHMOPS
ALL VSMWORK3 * 150A ADGHMOPS
ALL WAVEWRKS * 140A ADGHMOPSZ
ALL WAVEWRKS * 150A ADGHMOPSZ
ALL WAVEWRKC * 140A ADGHMOPSZ
ALL WAVEWRKC * 150A ADGHMOPSZ
ALL WAVEWRKL * 140A ADGHMOPSZ
ALL WAVEWRKL * 150A ADGHMOPSZ
```

2. Now, save the changed `AUTHFOR CONTROL` file and send it back to DirMaint by running the following commands:
   - **`DIRM FILE AUTHFOR CONTROL`**
   - **`DIRM RLDD`**
   - **`DIRM RLDC`**

**Note:** When migrating from Fix Pack 3 or earlier for FCP devices, you might need to delete and re-add the z/VM system in IBM Wave to free the FCP devices within IBM Wave.

## IBM Wave for z/VM V1R1 Fix Pack 5 readme file

There are several fixes. There are some new functions.

### New functions

Here are the new functions of this fix pack:

- Mixed case password support

  Instead of making passwords that are passed from the GUI client to the server uppercase, the case is not changed. You can use this function to use the mixed case passwords that RACF supports. If you change to mixed case, then you must update the TVI-API password by clicking **Administrative/Site Management** → **Update Authorized TVI-API Credentials** after you establish a mixed case password for that user.

- ext4 file system support

  IBM Wave supports creating a DASD partition (LVM or otherwise) on an ext4 file system for RHEL6. IBM Wave also supports extending an ext4 file system on RHEL6 Linux guests.

- Enhancements for SAN (FCP/SCSI) environments by using EDEVs

  IBM Wave supports the following functional features on configurations by using SCSI disks with emulated FBA devices (EDEVS):

  - Manage disks
  - Auto detect devices
  - Generate disk storage maps
  - Add a disk to a guest (add a partition to a guest)
  - Create a partition (file system)
  - Extend a partition (resize a file system)
  - Create or extend an LVM volume group
  - Create or extend an LVM logical volume
  - Manage spool devices
  - Manage page devices
  - Manage guests
  - Create a guest
  - Convert to prototype
  - Clone
  - Clone from prototype
  - IBM Wave installation
  - Support of the IBM Wave server itself by using directly attached FCP
  - Create service machines on EDEVS during auto detect
  - Bare Metal Installations

- Support of a new `PROFILE EXEC` exit to allow for site-specific configuration

  The `PROFILE EXEC` file on the WAVEWRKS service machine now starts a user-supplied exit that is named `XPRFEXIT EXEC`, if it exists. This exit may be used for a site-specific configuration, such as linking and accessing disks.

  If you have a site-specific configuration that is added to your `PROFILE EXEC`, you should relocate that logic to the exit. Do this task before updating your system with this service pack because the `PROFILE EXEC` is overwritten.

**Important information for changing the authorized TVP-API password on the z/VM TVP:**

It is possible for the password of an Authorized TVP-API user to be changed on z/VM by a system administrator. When doing so, special consideration must be taken to prevent the TVP-API credentials that are stored in the IBM Wave knowledge base from becoming invalid.

Before changing the password for the Authorized TVP-API user on a managed z/VM system, the z/VM system in IBM Wave must first be suspended, and all outstanding scheduled tasks on the BTS should be allowed to either complete or terminate (you can view scheduled BTS tasks through the BTS Manager, as described in the “BTS Manager” section in the *IBM Wave for z/VM: User's Guide*).

If the z/VM system is a member of a single-system image (SSI) cluster with a shared directory or shared RACF database, then this task should be done for every z/VM system in the cluster. After the password is changed for the guest on z/VM, then the TVP-API credentials that are stored in IBM Wave for the managed z/VM system must be changed to reflect the new password.

If the z/VM system is a member of an SSI cluster with a shared directory or shared RACF database, then this task should be done for every z/VM system in the cluster. Only then should the z/VM systems be resumed in IBM Wave.

For information about how to suspend and resume a z/VM TVP in IBM Wave, see the “Update Details” section *in IBM Wave for z/VM: User's Guide and Reference*, SC27-6119-08.

## IBM Wave for z/VM V1R1 fix pack 4 readme file

There are several fixes. There are some new functions.

The same information that is found in “IBM Wave for z/VM V1R1 Fix Pack 5 readme file” on page 280 is applied to this fix pack as well.

## IBM Wave for z/VM V1R1 fix pack 3 readme file

There are several fixes. There is one new function.

### New function

Minidisk passwords for the service machines must now be specified by the user when a system is auto detected. If the service machines are created before the auto detect, then the passwords that are specified by the user must match the ones that are created on z/VM for the minidisks.

## IBM Wave for z/VM V1R1 fix pack 2 readme file

There are several fixes. There are no new functions.

# Customization with other security managers

This section describes z/VM customization for CA VM:Secure. For information about this product, see the following website:

http://www.ca.com/us/securecenter/ca-vm-secure-for-z-vm.aspx

## Customizing VM:Secure

When you use VM:Secure, you must define and authorize user IDs and resources. VM:Secure needs to be at Version 3.1 or later.

### CONFIG security

For the CONFIG security, consider the following items:

- GROUP GSMAPI

  The GSMAPI group is for SMAPI and contains `VSMWORK1`, `VSMWORK2`, `VSMWORK3`, `VSMREQIN`, and `VSMREQIU`.

- GROUP WAVEACIG

  In the `WAVEACIG GROUP RULES`, add the following line:

  ```
  ACCEPT * LINK 399 * (NOPASS
  ```

### System rules

For these system rules, adjust the name if you select different names for the Wave Service Machines.

```
ACCEPT * LOGON
ACCEPT ftp machine LINK * * (NOPASS
ACCEPT ftp machine DIAG88
ACCEPT ftp machine DIAGD4
ACCEPT ftp machine SPOOL
ACCEPT VSMWORK1 AUTOLOG (NOPASS
ACCEPT GSMAPI TAG * (GROUP
ACCEPT GSMAPI SPOOL (GROUP
ACCEPT GSMAPI DIAGD4 (GROUP
ACCEPT GSMAPI DIAG88 (GROUP
ACCEPT VMSERVR LINK * * (NOPASS
ACCEPT VMSERVU LINK * *(NOPASS
ACCEPT VMSERVS LINK * * (NOPASS
ACCEPT WAVEACIG LINK * *(GROUP NOPASS
ACCEPT WAVEACIG AUTOLOG (GROUP NOPASS
ACCEPT WAVEWRKS LINK (HISTORY NOPASS
ACCEPT WAVEWRKC LINK (HISTORY NOPASS
ACCEPT WAVEWRKL LINK (HISTORY NOPASS
ACCEPT WAVEWRKS AUTOLOG (HISTORY NOPASS
ACCEPT WAVEWRKC AUTOLOG (HISTORY NOPASS
ACCEPT WAVEWRKL AUTOLOG (HISTORY NOPASS
ACCEPT WAVEWRKS VALIDATE
ACCEPT WAVEWRKC VALIDATE
ACCEPT WAVEWRKL VALIDATE
ACCEPT WAVEWRKS DIAG88
ACCEPT WAVEWRKC DIAG88
ACCEPT WAVEWRKL DIAG88
```

- ACCEPT VSPROXY DIAG88
- ACCEPT VSMREQI6 DIAG88
- ACCEPT VSMEVSRV DIAG88
- ACCEPT PERFSVM DIAG88

### MAINT user rules

Here are the MAINT user rules:

- ACCEPT GSMAPI LINK 190 RR (GROUP NOPASS
- ACCEPT GSMAPI LINK 19E RR (GROUP NOPASS
- ACCEPT GSMAPI LINK 193 RR (GROUP NOPASS
- ACCEPT VSMWORK1 LINK CF1 MD
- ACCEPT VSMWORK1 LINK CF2 MD
- ACCEPT VSMWORK2 LINK CF1 MD
- ACCEPT VSMWORK2 LINK CF2 MD
- ACCEPT VSMWORK3 LINK CF1 MD
- ACCEPT VSMWORK3 LINK CF2 MD
- ACCEPT * LINK 190 RR
- ACCEPT * LINK 19D RR
- ACCEPT * LINK 19E RR

### TCPMAIN user rules

Here are the TCPMAIN user rules:

- ACCEPT GSMAPI LINK 0591 RR (GROUP NOPASS
- ACCEPT GSMAPI LINK 0592 RR (GROUP NOPASS

IBM Wave uses the z/VM FTP server to transfer files to minidisks that are owned by the IBM Wave Service Machines.

### VMRMAINT user rule

Here is the VMRMAINT user rule:

```
ACCEPT GSMAPI LINK 193 RR (GROUP NOPASS
```

### AUTHORIZ.CONFIG

Here are the rules that should be in the `AUTHORIZ.CONFIG` file:

- LIST *WAVEWRK WAVEWRKS WAVEWRKL WAVEWRKC
- GRANT * TO VSMWORK1
- GRANT * TO VSMWORK2
- GRANT * TO VSMWORK3
- GRANT * OVER *ALL TO WAVEWRKS
- GRANT * OVER *ALL TO WAVEWRKC
- GRANT * OVER *ALL TO WAVEWRKL

### Managers file changes

Add the VSMWORK1 machine as a manager, as shown in the following example:

```
MANAGER VSMWORK1 * POOL1 SKELETON VSMWORK1 GENERAL DEVTYPE VSMWORK1 3390
```

Add the user that will be passed to the SMAPI as a manager. Here are some examples with their corresponding commands.

For MAINT:

- `MANAGER MAINT * POOL1`
- `SKELETON MAINT GENERAL`
- `DEVTYPE MAINT 3390`

For Wave:

- `MANAGER WAVEWRKS * POOLX POOLY`
- `SKELETON WAVEWRKS GENERAL`
- `DEVTYPE WAVEWRKS 3390`
- `MANAGER WAVEWRKC * POOLX POOLY`
- `SKELETON WAVEWRKC GENERAL`
- `DEVTYPE WAVEWRKC 3390`
- `MANAGER WAVEWRKL * POOLX POOLY`
- `SKELETON WAVEWRKL GENERAL`
- `DEVTYPE WAVEWRKL 3390`

In this example, `POOLX` and `POOLY` are DASD pools in your environment. Ensure that IBM Wave is a manager for the DASD pools that are defined for Wave

For more information about customization with other security managers, see the following publications:

- *CA VM:Secure for z/VM*
- *CA VM:Director for z/VM Administration Guide*
- *CA VM:Secure for z/VM Rules Facility Guide*

# Wave CLI helper file

In the environment that we used for this book, we created/modified a file.

### WAVECLI.BAT

We modified the `wavecli.bat` file to better match our requirements. Our file is shown in Figure A-3 on page 285.

```
@echo off

rem Licensed Materials - Property of IBM
rem 5648-AE1
rem (c) Copyright IBM Corp. 2014.
rem Version:1.20-itso.01
rem *-------------------------------------------------------*
rem |                                                       |
rem |  1) Use the variables to improve readability          |
rem |     of the script                                     |
rem |  2) The actual command  and it parameters are entered |
rem |     as the positional parameters %1 %2 %3 %4 etc      |
rem |                                                       |
rem *-------------------------------------------------------*

set j1=WAVECommon120.jar
set j2=log4j-1.2.16.jar
set j3=mariadb-java-client-1.1.1.jar
set j4=jdom.jar
set j5=commons-collections-3.2.jar
set j6=commons-dbcp-1.2.2.jar
set j7=commons-pool-1.4.jar
set j8=jsch-0.1.42.jar
set j9=WAVECLI120.jar

set javaCP=%j1%;%j2%;%j3%;%j4%;%j5%;%j6%;%j7%;%j8%;%j9%

set c01=%1
set c02=%2
set c03=%3
set c04=%4
set c05=%5
set c06=%6
set c07=%7
set c08=%8
set c09=%9

shift
shift
shift
shift
shift

set c10=%5
set c11=%6
set c12=%7
set c13=%8
set c14=%9

java -cp %javaCP% com.CSL.WAVE.CLI.WAVECLI %c01% %c02% %c03% %c04% %c05% %c06%
%c07% %c08% %c09% %c10% %c11% %c12% %c13% %c14%
echo.
```

*Figure A-3 Modified wavecli.bat file*

# Wave flow charts

Figure A-4 through show Figure A-7 on page 289 show flow charts for IBM Wave planning, preparation, installation, and setup.

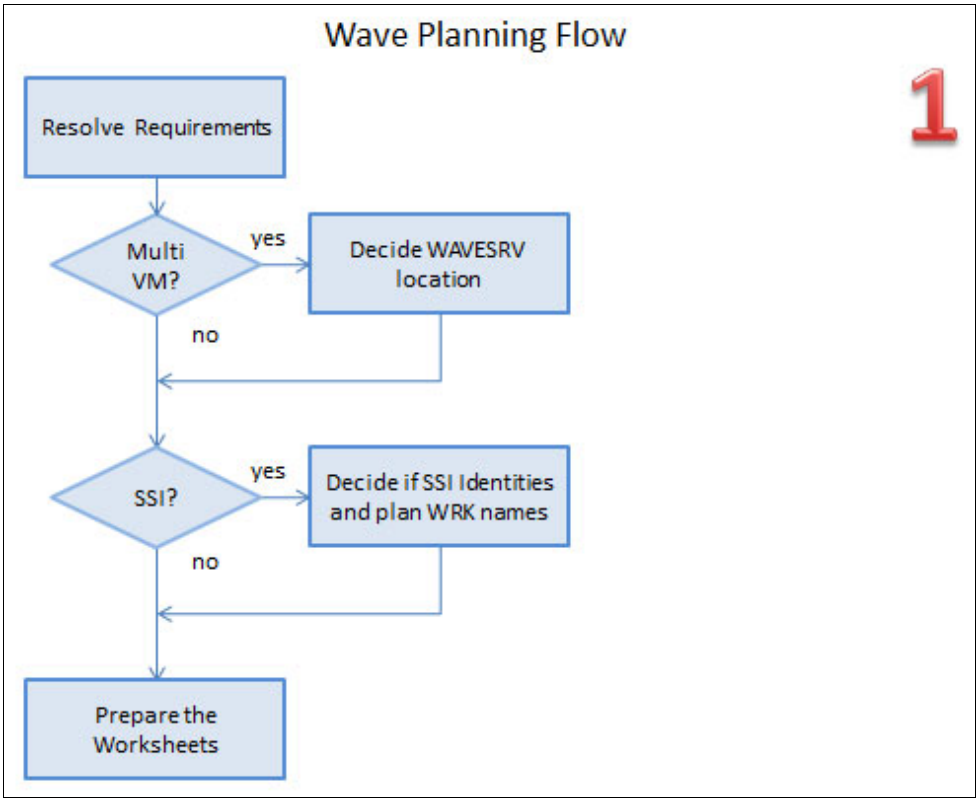


*Figure A-4 Wave Planning Flow*

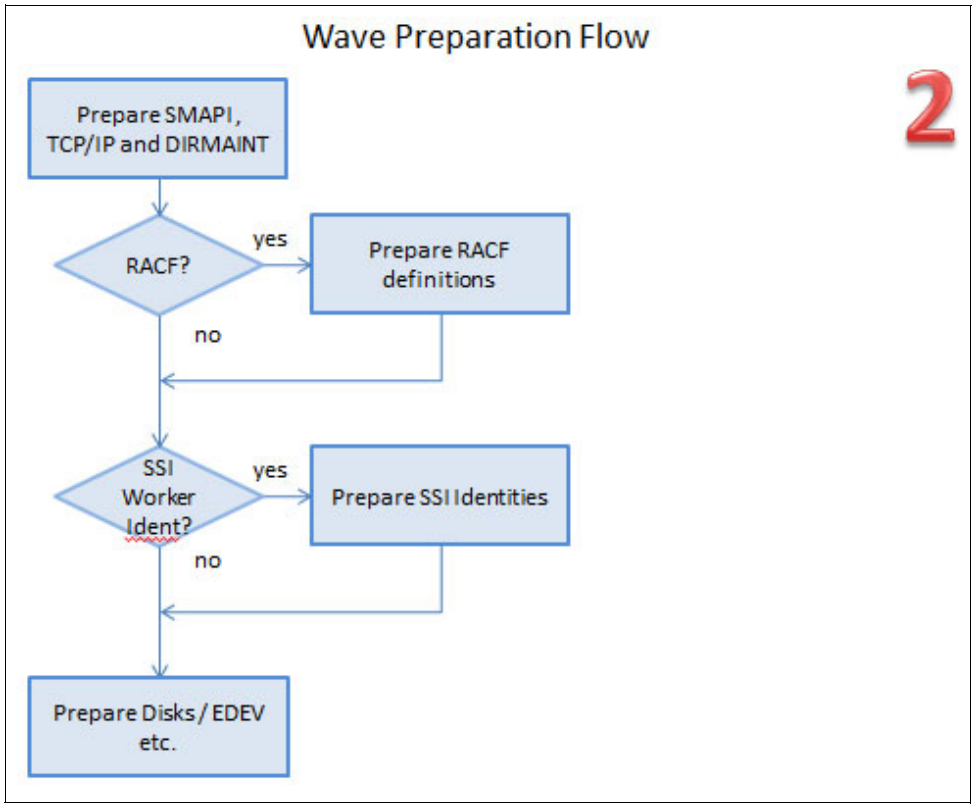


*Figure A-5 Wave Preparation Flow*

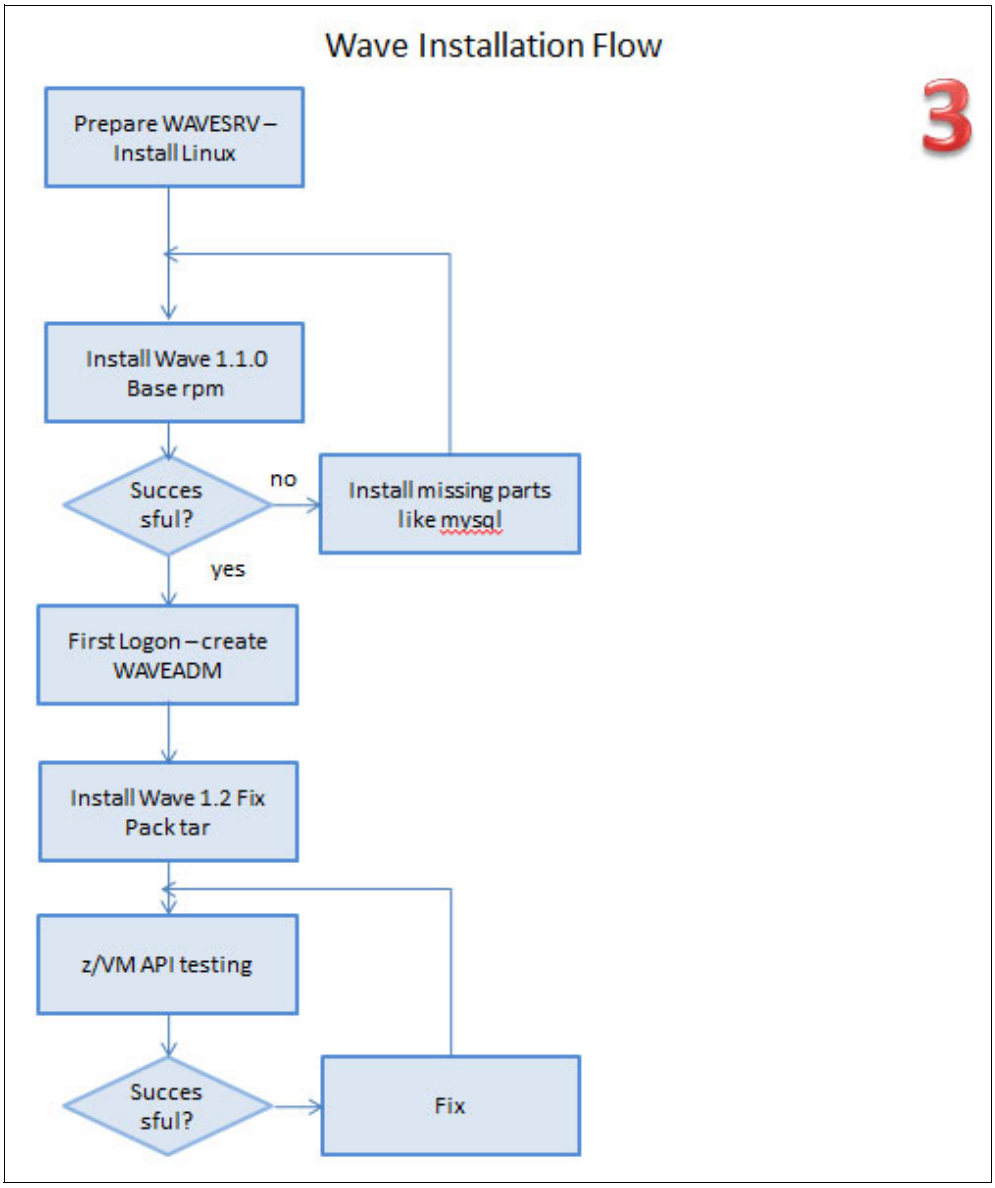


*Figure A-6 Wave Installation Flow*

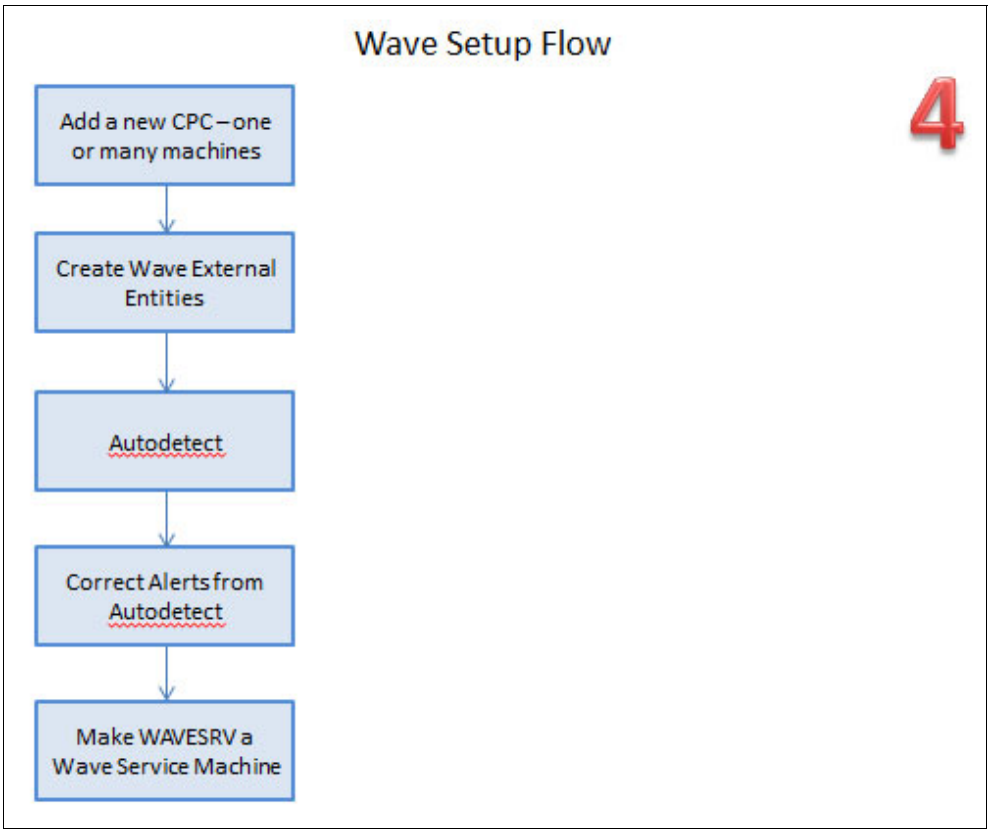


*Figure A-7 Wave Setup Flow*

# Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this book.

## IBM Redbooks

The following IBM Redbooks publications provide more information about the topic in this document. Some publications that are referenced in this list might be available in softcopy only:

- *Set up Linux on IBM System z for Production*, SG24-8137
- *Using z/VM v 6.2 Single System Image (SSI) and Live Guest Relocation (LGR)*, SG24-8039
- *The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3*, SG24-8147

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials at this website:

http://www.ibm.com/redbooks

## Other publications

The following publications are also relevant as further information sources:

- *IBM Wave for z/VM: Administration and Customization*, SC27-6118
- *IBM Wave for z/VM: User Guide and Reference*, SC27-6119

## Online resources

The following websites are also relevant as further information sources:

- IBM Offering Information page (announcement letters and sales manuals):

  http://www.ibm.com/common/ssi/index.wss?request_locale=en
- IBM Wave for z/VM product page:

  http://www-03.ibm.com/systems/z/solutions/cloud/wave/index.htmlDescription2

## Help from IBM

IBM Support and downloads:

http://www.ibm.com/support

IBM Global Services

http://www.ibm.com/services

Redbooks

IBM Wave for z/VM Installation, Implementation, and Exploitation

SG24-8192-01

ISBN 0738441139

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

SG24-8192-01

ISBN 0738441139

Printed in U.S.A.