Redbooks
ibm.com/redbooks

SAP Landscape Management 3.0 and IBM Power Systems Servers

Edmund Haefele

Arnold Beilmann

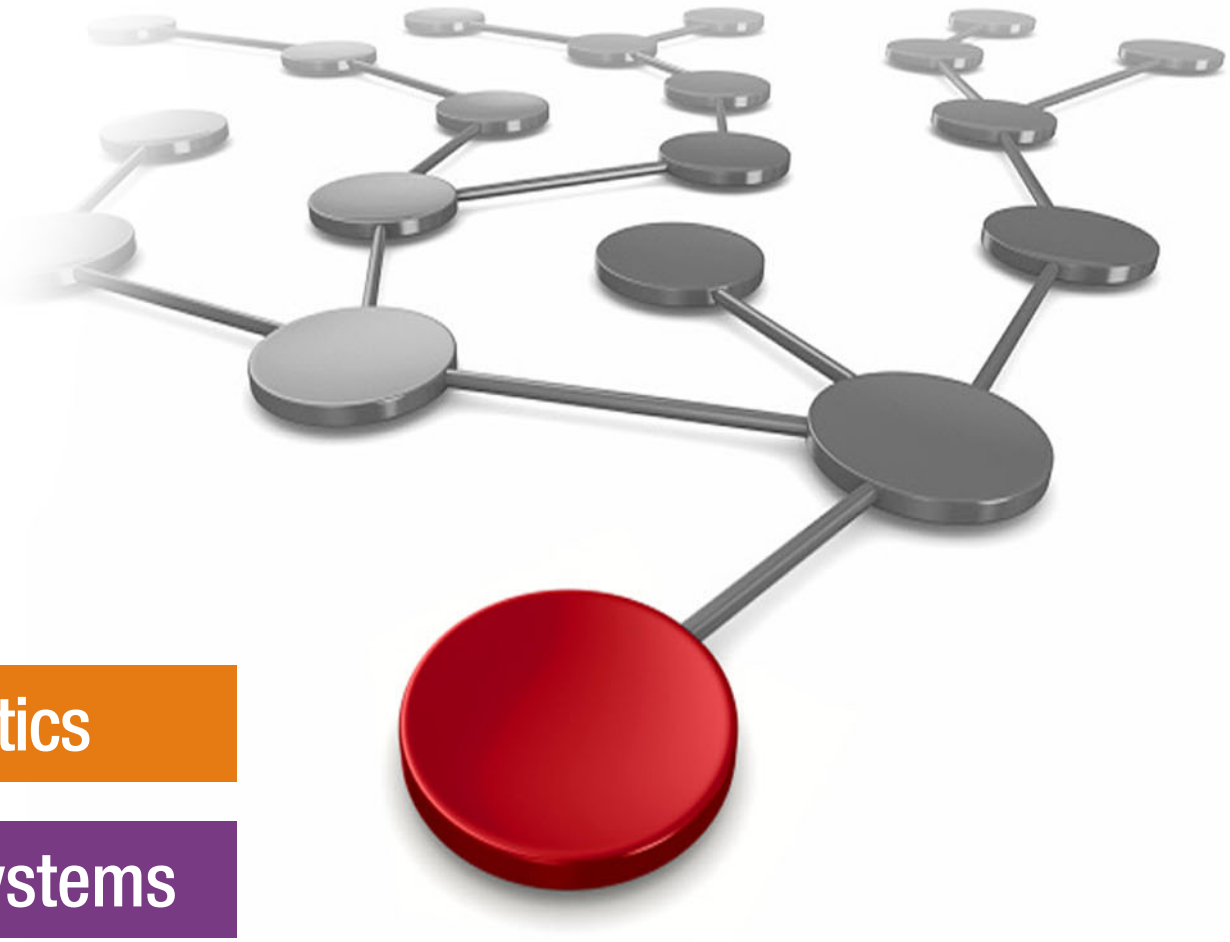

Analytics

Power Systems

IBM

Redpaper

IBM Redbooks

SAP Landscape Management 3.0 and IBM Power Systems Servers

November 2019

REDP-5568-00

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

First Edition (November 2019)

This edition applies to SAP Landscape Management (SAP LaMa) V3.0.

© Copyright International Business Machines Corporation 2019. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices v
Trademarks vi

Preface vii
Authors vii
Now you can become a published author, too! viii
Comments welcome viii
Stay connected to IBM Redbooks viii

Chapter 1. Introducing SAP Landscape Management 3.0 and IBM Power Systems servers 1
1.1 Introduction 2

Chapter 2. SAP Landscape Management 3.0 3
2.1 Overview 4
2.2 Introducing typical use cases with SAP LaMa 7
2.2.1 Use case 1: Managing complex and distributed SAP landscapes 7
2.2.2 Use case 2: Managing SAP HANA landscapes 7
2.2.3 Use case 3: Relocating adaptively installed SAP systems 8
2.2.4 Use case 4: SAP System Clone, SAP System Copy, and SAP System Refresh (including a refresh of the DB server) 9
2.2.5 Use case 5: Virtualization monitoring of Power Systems servers and LPARs . . . 11
2.2.6 Use case 6: Virtualization management of Power Systems servers and LPARs . 12
2.2.7 Use case 7: Operating system provisioning and image handling 12
2.2.8 Use case 8: Provisioning SAP Application Servers 13
2.2.9 Use case 9: Extensibility with user-defined operations and workflows 13
2.3 Setting up an SAP LaMa environment 14
2.3.1 Enabling SAP systems 14
2.3.2 Installing the SAP Landscape Management software 17
2.3.3 Installing and configuring an IBM PowerVC environment 18
2.3.4 Preparing SAP NetWeaver AS Java for communication with IBM PowerVC 19
2.3.5 Configuration steps in SAP LaMa 21
2.3.6 Discovering managed systems and hosts in SAP LaMa 26
2.3.7 Configuring SAP hosts in SAP LaMa for copy use cases 30
2.3.8 Configuring SAP systems in SAP LaMa for copy use cases 31
2.4 Running SAP System Copy in SAP LaMa 33
2.5 Outlook 43
2.5.1 SAP containerized applications 43
2.5.2 SAP HANA management 44

Related publications 45
IBM Redbooks 45
Online resources 45
Help from IBM 45

© Copyright IBM Corp. 2019. All rights reserved.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© Copyright IBM Corp. 2019. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM Cloud™	Redbooks (logo) ®
IBM®	Redbooks®	System Storage™

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

OpenShift, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

Preface

This IBM® Redpaper publication is part of a series of technical documentation to help the enablement of SAP on Linux for IBM Power Systems servers and IBM System Storage™ servers.

This book describes how by using SAP Landscape Management (SAP LaMa) 3.0 software that clients gain full visibility and control over their SAP and non-SAP systems, including the underlying physical, virtual, and cloud infrastructures. With SAP LaMa, you can automate repetitive tasks to manage critical applications across complex, hybrid IT landscapes.

This publication helps you to better control IT costs and increase business agility, for example, by freeing staff to focus on more strategic work rather than manual, error-prone tasks.

The target audiences of this book are architects, IT specialists, and systems administrators deploying SAP LaMa 3.0 whom often spend much time and effort managing and provisioning SAP software systems and landscapes.

Authors

This paper was produced in close collaboration with the IBM SAP International Competence Center (ISICC) in Walldorf, SAP Headquarters in Germany and IBM Redbooks®.

Edmund Haefele is a specialist in IT infrastructure for SAP environments with more than 20 years of experience in that field. Edmund is Level 2 Expert Certified as an IBM architect and worked both as an IT architect and specialist in many projects for both German and international clients. Edmund has a deep knowledge of IT systems, such as SAP HANA, servers, storage, network technologies and high availability (HA), backup and restore, disaster recovery (DR), service management, and automation. He also is a Certified Open Group Master Certified IT Architect. Edmund holds a diploma in Physics and a PhD in natural sciences from the University of Heidelberg.

Arnold Beilmann is a senior software developer for IBM and has more than 10 years of experience porting different SAP solutions to Power Systems servers. He is responsible for the architecture and implementation of monitoring and management solutions for Power Systems servers with SAP, Red Hat OpenShift, and IBM Cloud™ infrastructure-as-a-Service (IaaS) offerings. Arnold is developing IBM adapters for SAP LaMa and its predecessors (SAP NetWeaver Adaptive Computing Controller (SAP NetWeaver ACC) and SAP Landscape Virtualization Management (SAP LVM)) for 12 years. He is leading the integration of SAP LaMa with IBM Power Systems servers and is responsible for the architecture and development of this SAP and IBM solution.

© Copyright IBM Corp. 2019. All rights reserved.

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

Dino Quintero and Wade Wallace
IBM Redbooks, Poughkeepsie Center

Katharina Probst, Walter Orb, Tanja Scheller
IBM Germany

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 IBM Redbooks 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 papers to be as helpful as possible. Send us your comments about this paper 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, IBM Redbooks
  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

1

Introducing SAP Landscape Management 3.0 and IBM Power Systems servers

This document is part of a series of technical documentation for SAP on Power Systems servers and IBM System Storage servers.

The target audience is architects or administrators deciding on purchasing and deploying SAP Landscape Management 3.0 (SAP LaMa) on IBM Power Systems servers.

This chapter contains the following topic:

- Introduction

© Copyright IBM Corp. 2019. All rights reserved.

1.1 Introduction

Clients often spend much time and effort managing and provisioning SAP software systems and landscapes.

By using SAP LaMa software, clients gain full visibility and control over their SAP and non-SAP systems, including the underlying physical, virtual, and cloud infrastructures. With SAP LaMa, you can automate repetitive tasks to manage critical applications across complex, hybrid IT landscapes so that you can better control IT costs and increase business agility, for example, by freeing staff to focus on more strategic work rather than manual, error-prone tasks.

IBM PowerVC is an advanced virtualization and cloud management offering. Built on OpenStack, it provides simplified virtualization management and cloud deployments for IBM AIX®, IBM i, and Linux virtual machines (VMs) running on IBM Power Systems servers. You can use the offer to build a private cloud on the Power Systems servers and improve administrator productivity. It can further integrate with multi-cloud environments through higher-level cloud orchestrators.

SAP LaMa provides a built-in integration with IBM PowerVC so that the SAP Basis administrator can holistically manage critical SAP applications. The SAP Basis administrator gets a complete overview of the full landscape of storage, servers, logical partitions (LPARs), and SAP Application Server and Database Instances, and they can efficiently manage and operate those components in a Power Systems and Storage System landscape.

Chapter 2, “SAP Landscape Management 3.0” on page 3 provides information about integration aspects with Power Systems technologies.

2

SAP Landscape Management 3.0

This section describes the SAP Landscape Management (SAP LaMa) 3.0 integration aspects with IBM Power Systems technologies.

This chapter contains the following topics:

- Overview
- Introducing typical use cases with SAP LaMa
- Setting up an SAP LaMa environment
- Running SAP System Copy in SAP LaMa
- Outlook

© Copyright IBM Corp. 2019. All rights reserved.

2.1 Overview

This section gives an overview of SAP LaMa and focuses on key aspects for integration with IBM Power Systems technologies.

SAP LaMa is an application that runs on top of an SAP NetWeaver Application Server (SAP NetWeaver AS) for Java. It provides a centralized and simplified framework for managing SAP and non-SAP solutions and systems, including SAP HANA and SAP S/4HANA systems. The goal of SAP LaMa is to simplify the management of hybrid SAP landscapes that are deployed either on-premises bare-metal or virtualized servers, or in the cloud.

SAP LaMa enables the automation of repetitive, day-to-day administration tasks for the SAP Basis administrators, which include complex provisioning operations like copying SAP systems or installing SAP application servers. By using SAP LaMa and its front end, SAP Basis administrators have simplified access to the infrastructure elements that are required for fulfilling their tasks.

SAP recently added many extensions to SAP LaMa as a standard solution for managing and operating SAP HANA powered landscapes. Advanced SAP HANA operations like performing system replication setup, takeovers, failbacks, system copies, system refreshes, and minimizing business downtime during SAP HANA maintenance activities can be managed with SAP LaMa (*LaMa4HANA*).

SAP LaMa relies on the SAP Host Agent and SAP Adaptive Extensions for monitoring and management of the operating system (OS). Those components must be installed on each host that is managed by SAP LaMa. The SAP Host Agent package includes all required elements for centrally monitoring and managing any kind of OSs that are covered in the SAP NetWeaver support matrix, which you can see in the SAP Product Availability Matrix (PAM). The SAP Host Agent is complemented by SAP Adaptive Extensions. By using those methods, SAP LaMa performs various operations on the host, such as:

- Perform operations on SAP HANA systems and tenant databases.
- Discover SAP HANA systems.
- Provision SAP HANA tenant databases.
- Monitor SAP HANA systems and tenant databases.
- Reexport Network File System (NFS).
- Manage virtual host names and IBM Storage by using SAP Adaptive Extensions (SAPACOSPrep) for the IBM PowerLinux platform.

Three main integration points for SAP Partner Technologies are available for SAP LaMa:

- Virtualization and Cloud Management
- Storage Management
- Orchestration

IBM provides adapters for Virtualization Management and Storage Management that are directly integrated with the SAP LaMa software. These adapters convert the generic SAP LaMa calls to platform-specific commands that are then run by the appropriate infrastructure component. The integration with IBM PowerVC allows virtualization monitoring and management of IBM Power Systems servers and logical partitions (LPARs), including the provisioning of new LPARs. By using IBM PowerVC to make the virtualization data available, all the dependencies on which the SAP systems are installed and the physical server on which they are running can be identified and visualized. When using IBM back-end storage that is storage area network (SAN)-attached to the LPAR by using N-Port ID Virtualization (NPIV) and managed by IBM PowerVC, then adaptively installed SAP systems can be relocated between different LPARs, and “Storage-based” System Copy/Clone can be run.

For Orchestration, SAP LaMa provides a REST API that you can use to get or update information from or to SAP LaMa or to trigger actions in SAP LaMa remotely. This API is a generic interface because there is no IBM specific extension that is available. Clients can define their own custom operations and processes and trigger them externally.

Table 2-1 lists the communication layers between SAP LaMa and the managed components (SAP instances and infrastructure).

Table 2-1 Communication layers for SAP LaMa and managed components

Method	Description
Remote Function Call (RFC)	SAP LaMa triggers RFCs in the SAP system for running SAP Basis system management tasks (for example, Post Copy Automation (PCA) task lists).
SAP Instance Agent (**`sapstartsrv`**)	Start, stop, and get the status of the SAP Instance.
SAP Host Agent (**`saphostexec`**) and Adaptive Extensions	Runs commands on the OS level to mount or unmount file systems, acquire or release IP addresses, and define custom commands. Perform database actions and monitoring (**`sapdbctrl`**).
Virtualization Adapter (IBM PowerVC)	Full visibility across server and virtualization, and mapping with SAP application topology. Automated provisioning of LPARs. Trigger live migrations from one system to another. Power off and power on LPARs. Gather information about hosts and OSs statuses and details.
Storage Adapter (IBM PowerVC)	Gather details about the storage volumes of the SAP system and instances. Perform storage mapping (attach and detach storage volumes to and from a host). Manage consistency groups and snapshots for the volumes that are required for SAP System Copy and Clone scenarios.

Figure 2-1 illustrates the different communication paths.

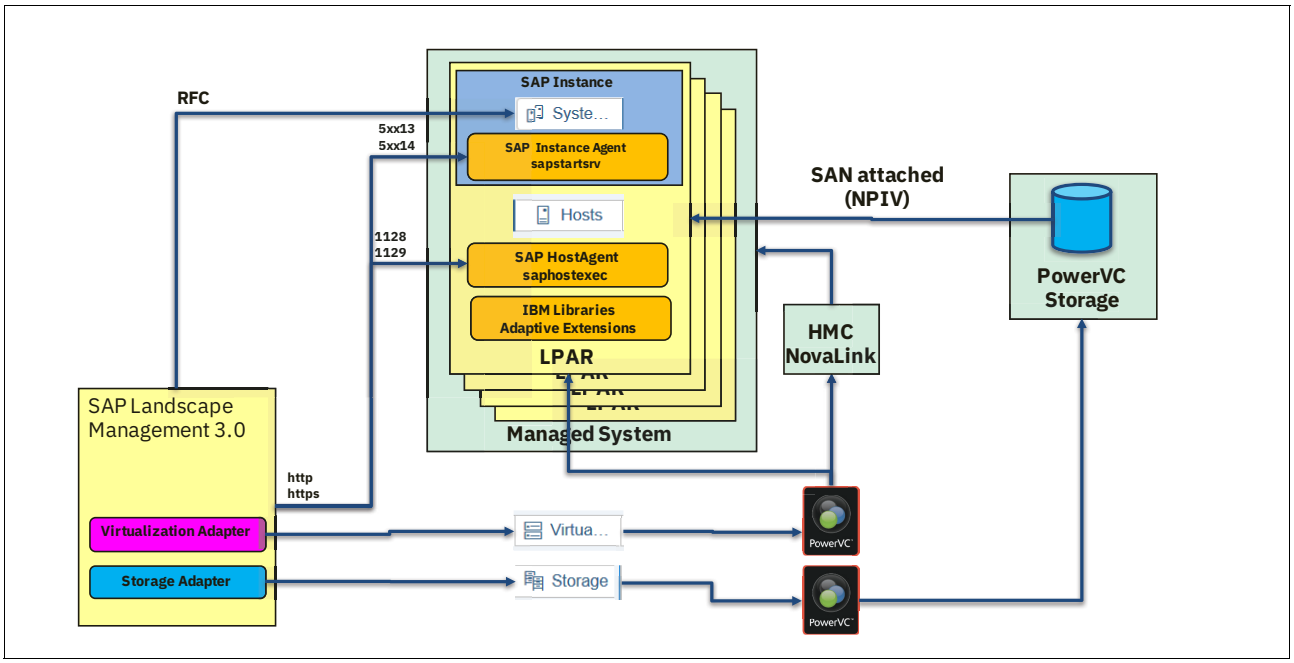

Figure 2-1 Communication between SAP LaMa and the managed components

With SAP LaMa capabilities like Custom Provisioning and Custom Hooks and Operations, you can extend standard workflows and create specific customer actions. SAP LaMa offers a *custom cloning* capability that you can use to create a system clone or system copy by integrating storage cloning procedures that are in at the client environment. Such an approach can be integrated into SAP LaMa: Pre-defined substeps in the *standard* workflow for the pre-clone, clone, and post-clone phases are replaced by custom-defined scripts or web service calls. The calls run on the source and target systems during the custom provisioning flow.

The web services and scripts must be registered as a *provider implementation definition* that references the script and web services execution that is triggered by the SAP Host Agent on the source or target system during the workflow step. The provider implementation definition references the location of the script or command to run on the LPAR, including all the parameters and arguments that are transferred from SAP LaMa to the script. For more information, see SAP Note 1465491 - Provider Implementation Definition and SAP Landscape Management 3.0, Enterprise Edition.

Then, you can create a custom provisioning process. In the custom provisioning process, you can replace a standard process step with the step that is referenced in the provider implementations. It is specified *when or under which circumstances* the custom step is run in an SAP LaMa operation.

Provider implementation definitions can be used in the context of custom hooks or custom operations. Custom hooks allow you to add more pre-, post-, or error processing steps to the standard workflows that are defined in SAP LaMa. Custom operations allow the triggering of user-defined activities directly from the SAP LaMa GUI or as part of mass operations. For example, the SAP Basis administrators can run certain custom-defined scripts with privileged permissions on the hosts without needing root access.

2.2 Introducing typical use cases with SAP LaMa

This section describes typical use cases that clients implement for the management of SAP systems. The example use cases focus on SAP systems that are deployed to an IBM Power Systems environment with Linux OS:

1. Management of complex SAP landscapes that are composed of several SAP systems that are running on-premises or in different public clouds.
2. SAP HANA landscape management.
3. Relocation of adaptively installed SAP systems.
4. SAP System Clone, SAP System Copy, and SAP System Refresh (including refreshes of DB Server-only).
5. Virtualization monitoring of IBM Power Systems servers and LPARs.
6. Virtualization management of IBM Power Systems servers and LPARs.
7. OS provisioning and OS image handling.
8. Provisioning of SAP Application Servers.
9. Extensibility with user-defined operations and workflows.

This section contains a short description of these use cases. Its subsections describe specific requirements for the implementation of use cases 3, 4, 5, 6, and 7. All other use cases are not specific to an IBM Power Systems environment and are described in the SAP LaMa documentation in detail.

2.2.1 Use case 1: Managing complex and distributed SAP landscapes

An SAP landscape is composed of several SAP systems that are running on-premises or in different public clouds. SAP LaMa centralizes the management of different entities of the landscape within one user interface. The provided management view includes on-premises and cloud landscapes, physical hosts, virtual machines (VMs) (LPARs), OSs (hosts), SAP systems, and SAP instances. SAP LaMa provides functions to run basic operations (like start or stop) on SAP systems and SAP instances. These operations can be triggered also on groups of elements (mass operations).

2.2.2 Use case 2: Managing SAP HANA landscapes

SAP LaMa provides several options for management of SAP HANA landscapes:

- New SAP HANA tenant move with near-zero downtime maintenance
- Support for SAP HANA system and tenant rename
- End-to-end automation for SAP HANA system replication setup
- End-to-end automation for SAP HANA takeover and failback procedures
- Support for SAP HANA multitarget replication
- Near zero downtime maintenance for SAP HANA primary and secondary sites (such as SAP HANA upgrades, OS upgrades, or hardware maintenance)

Figure 2-2 illustrates new high availability (HA) and disaster recovery (DR) options with SAP HANA 2.0, which can be managed with SAP LaMa.

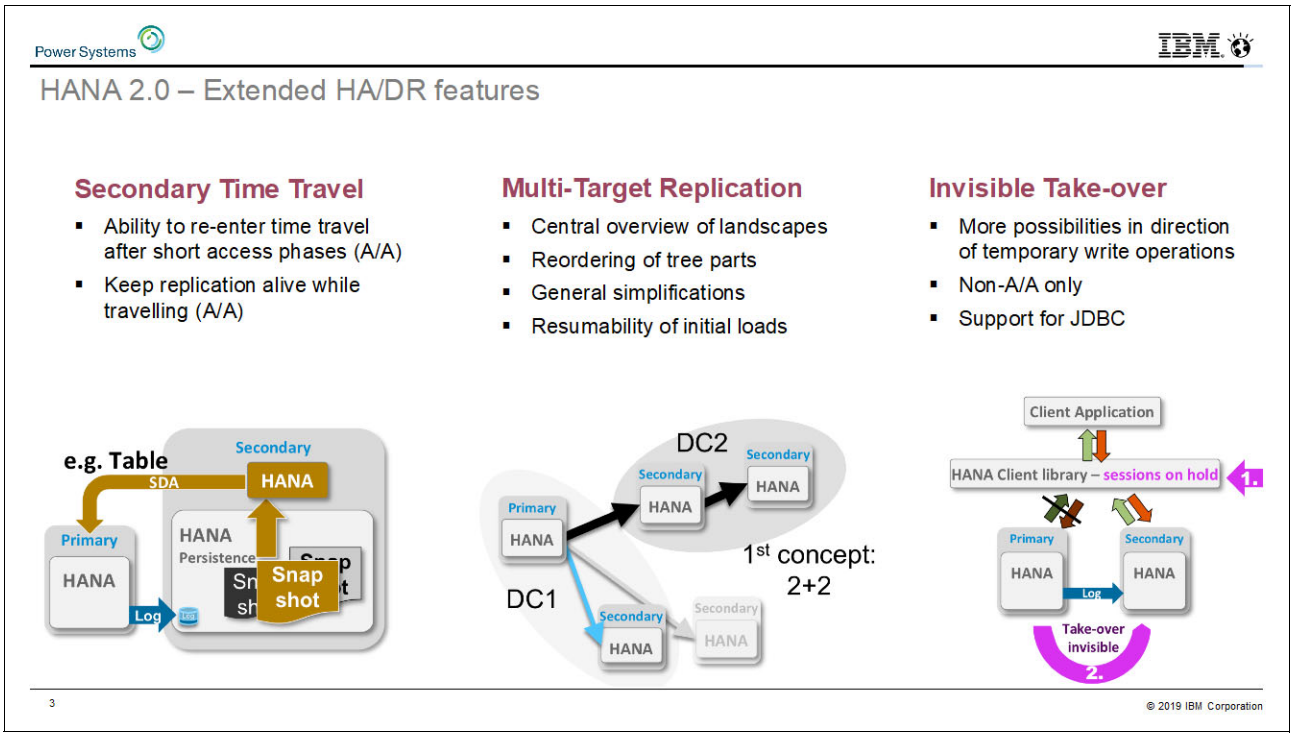

Figure 2-2 Extended HA and DR features in SAP HANA 2.0

2.2.3 Use case 3: Relocating adaptively installed SAP systems

SAP LaMa provides functions to move an entire SAP system or one of its instances from one host to another host. The SAP system or instance must be stopped on the source host and then restarted on the target host. The relocation is a combination of a detach task (*unprepare*) followed by an attach task (*prepare*) on the host. Detach and attach tasks invoke the IBM PowerVC storage adapter, so all the SAP system storage must be managed by IBM PowerVC. For this use case, IBM PowerVC is configured as storage adapter in SAP LaMa. Mass operations may relocate multiple SAP systems or instances immediately.

The end-to-end process invokes and automates the following steps:

1. Stop the SAP system or SAP instance.
2. Unprepare the SAP system or SAP instance on the source host. During this step, the following actions take place:
   - The SAP instance agent stops and unregisters from the host.
   - File systems (defined in the mount config) of the system or instance unmount.
   - The virtual IP addresses of the system or instance release from the network adapter on the source host.
   - Storage volumes reflecting the file systems unmap from the source host by using IBM PowerVC.

3. Prepare the SAP instance or system on the target LPAR. During this step, the following actions take place:
   - Storage volumes reflecting the file systems map to the target host by using IBM PowerVC.
   - All file systems (defined in the mount config) of the system or instance mount.
   - The virtual IP addresses of the system or instance are acquired in the network adapter of the target host.
   - The SAP instance agent registers on the host and starts.
4. Start the SAP system or instances on the target LPAR.

For this use case, the SAP systems or instances must be adaptively installed:

- Virtual IP addresses must be decoupled from the SAP instance from the host. Name resolution (forward and backward) for the virtual IP address and the associated virtual host name must be ensured on all involved hosts (source and target LPAR) and on the SAP LaMa system.
- The SAP instance requires an adaptive storage layout so that all required storage entities (volumes and file systems) of the SAP instance can be accessed on both the source and target LPARs. Access to the volumes in the storage system must be based on virtual Fibre Channel (NPIV) for both source and target LPARs.

2.2.4 Use case 4: SAP System Clone, SAP System Copy, and SAP System Refresh (including a refresh of the DB server)

SAP LaMa provides functions to copy an entire SAP system or the DB instance. The target system is either newly created (SAP System Clone or SAP System Copy) or an existing SAP system is refreshed with the actual content from another system. The system runtime environment is either identical to the source system (SAP System Clone) or the run time is changed to a new host name, new SAP system ID (SID), or new SAP instance number.

Figure 2-3 and Figure 2-4 illustrate the difference between SAP System Clone, SAP System Copy, and SAP System Refresh.

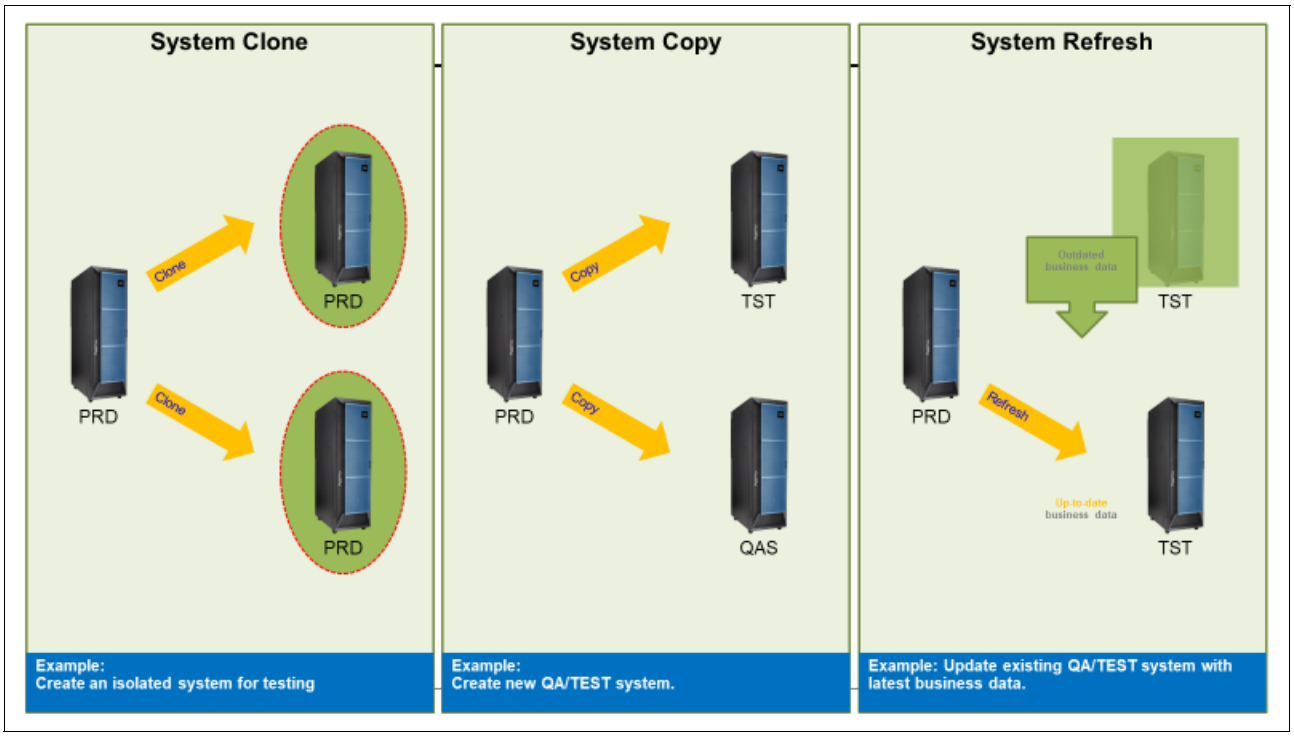

Figure 2-3 SAP System Clone, SAP System Copy, and SAP System Refresh

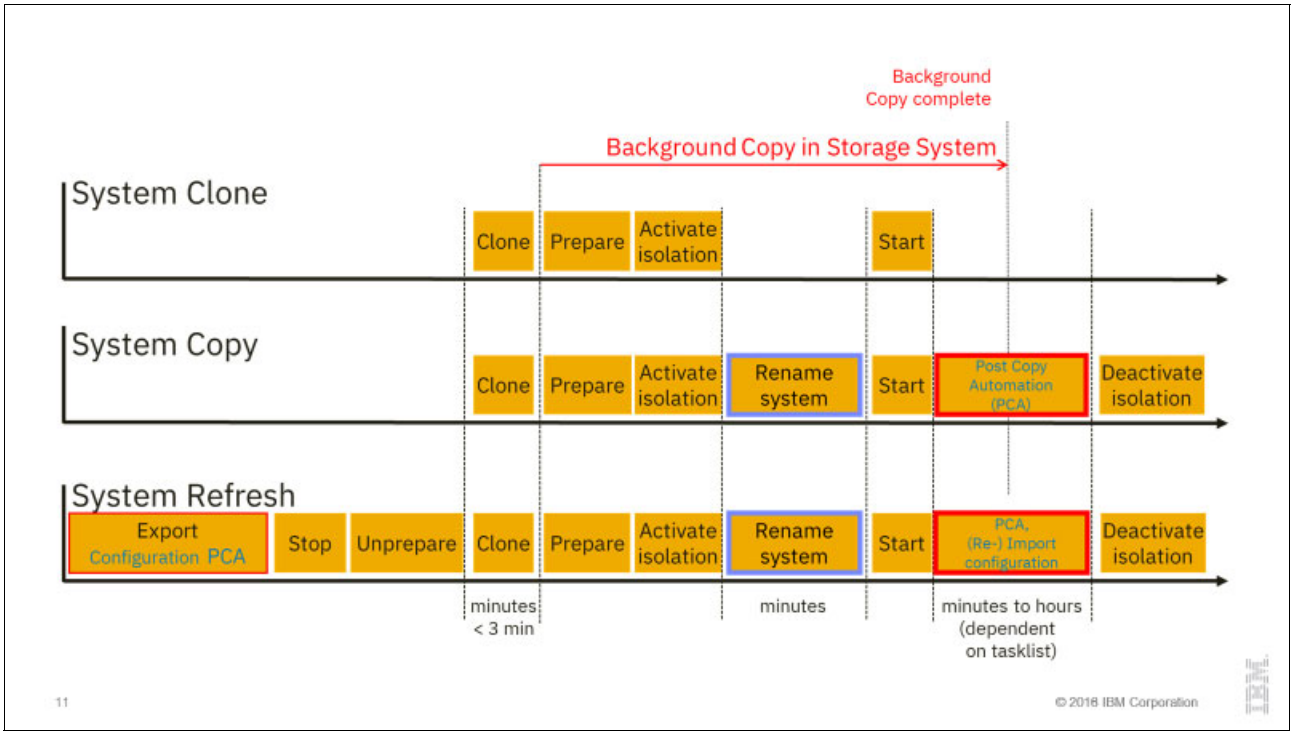

Figure 2-4 SAP LaMa execution steps for SAP System Clone, SAP System Copy, and SAP System Refresh

All storage volumes that are in the scope of the clone or copy process must be in one storage system and managed by IBM PowerVC. For this use case, the SAP systems or instances must be adaptively installed (see 2.2.3, “Use case 3: Relocating adaptively installed SAP systems” on page 8).

During the execution of this use case, SAP LaMa calls actions in the storage system by using IBM PowerVC to:

- Retrieve information about storage volumes of the SAP system or instances.
- Attach or detach storage volumes to and from the hosts in scope.
- Create consistency groups and snapshots for the volumes in scope.
- Verify the progress of the background copy.
- Target the host to import the cloned disk images and mount the file systems.

More preparation steps are required in SAP LaMa to enable those scenarios:

- Configure IBM PowerVC as a storage adapter.
- Define infrastructure elements in SAP LaMa configuration networks, user management, name server, and assignment (see 2.2.5, “Use case 5: Virtualization monitoring of Power Systems servers and LPARs” on page 11).
- Define SAP Software Provisioning Manager (SWPM) repositories for SAP System Copy and SAP System Refresh in SAP LaMa (see 2.2.5, “Use case 5: Virtualization monitoring of Power Systems servers and LPARs” on page 11).
- Configure discovered hosts and SAP instances (see 2.2.7, “Use case 7: Operating system provisioning and image handling” on page 12 and 2.2.8, “Use case 8: Provisioning SAP Application Servers” on page 13 (for example, enter RFC Users, configure network fencing exceptions, and so on)).
- Activate SAP System Clone, SAP System Copy, and SAP System Refresh in the SAP LaMa configuration for SAP system as an allowed use case (see 2.2.8, “Use case 8: Provisioning SAP Application Servers” on page 13).

2.2.5 Use case 5: Virtualization monitoring of Power Systems servers and LPARs

During the execution of this use case, SAP LaMa communicates with the virtualization manager and IBM PowerVC adapter to retrieve the monitoring data of IBM Power Systems physical servers and LPARs. SAP LaMa extends the landscape picture with information about the underlying IT infrastructure. SAP LaMa users see:

- Which physical server a particular SAP system is running.
- Which SAP systems are running on the same physical server.
- How many CPU and memory resources are available and how they are used (from the view of the hypervisor).

2.2.6 Use case 6: Virtualization management of Power Systems servers and LPARs

SAP LaMa triggers operations on IBM PowerVC to activate or deactivate an LPAR or to shut down the OS on it. SAP LaMa can also trigger an online relocation to another system by using Live Partition Mobility (LPM). For all these tasks, SAP LaMa monitors the execution process of the triggered operation and refreshes monitoring data. For this use case, IBM PowerVC is configured as the virtualization adapter in SAP LaMa. SAP LaMa users can perform the following actions:

- Activate an inactive LPAR.
- Deactivate an active LPAR.
- Shut down the OS on an active LPAR.
- Relocate an active LPAR to another IBM Power System server.

2.2.7 Use case 7: Operating system provisioning and image handling

SAP LaMa triggers operations on IBM PowerVC (which is configured as a virtualization adapter) to provision a new OS from an image that is available in the IBM PowerVC repository. The Power System servers where the new LPAR is created can be explicitly specified or are evaluated at run time due to the placement policy in IBM PowerVC. Further characteristics are defined by using a provisioning template that is chosen during the deployment process.

More preparation steps are required as follows:

1. Prepare image templates in IBM PowerVC:
   a. Install a golden image on an LPAR for capture (or capture an existing LPAR that meets the prerequisites).
   b. Capture golden images in IBM PowerVC.
   c. Define compute templates in IBM PowerVC.
2. Deploy the LPARs:
   d. Request IP addresses and host names for the new LPARs.
   e. Create DNS entries for the host names of the new LPARs.
   f. Deploy the new LPARs (provision a virtual host in SAP LaMa).

Defining content for the golden image for managed SAP systems

For this task, see “General Installation Information for Your Operating System” in SAP Note 2369910 - Installation of SAP Systems Based on the Application Server ABAP of SAP NetWeaver 7.0 to 7.03 on UNIX: SAP MaxDB and the references that are mentioned in it. From an SAP LaMa perspective, a few extra items must be considered for the Linux OS.

During an SAP System Copy, SAP LaMa uses the TCP/IP filtering capabilities that are bundled with the OS to isolate the target SAP system. To avoid impacting any other system or interface, outbound communication is blocked during a system copy, and is released only after the PostCopy Automation tasks complete. On Linux, those TCP/IP filters are included in the `iptables` packages (`iptables` must be part of the image).

SAP LaMa uses the SAP Host Agent on the OS to gather monitoring data and run actions on the host. Some specific actions are not included in the base SAP Host Agent package but are part of the SAP Adaptive Extensions package. Both SAP Host Agent and SAP Adaptive Extensions must be installed in the image. You can download SAP Host Agent and the SAP Adaptive Extensions for the OS from the SAP Software Download Center. Also, see the following SAP Notes for installation details:

- SAP Note 1031096 - Installing Package SAPHOSTAGENT
- SAP Note 1759181 - Installation of SPAACEXT.SAR

Note: These packages either need to be installed within the golden image before it is captured in IBM PowerVC or a pre-installation script must be created.

2.2.8 Use case 8: Provisioning SAP Application Servers

SAP LaMa can install SAP NetWeaver ASs on an existing system. In such a case, SAP LaMa triggers the SWPM on an LPAR to invoke the SAP application server installation.

As a prerequisite, the installation media (SWPM, InstMaster, SAP Kernel, and Database Client) must be extracted to a directory that can be attached to the LPAR (for example, by using NFS). Create a root directory for the repository, and copy the installation content to individual subdirectories. In the Infrastructure Repositories, add a configuration for the SWPM Configuration for Application Server Installation.

The profiles and global directories (`/sapmnt/<SID>`) of the SAP system must be available on the LPAR where the application server will be installed.

2.2.9 Use case 9: Extensibility with user-defined operations and workflows

SAP LaMa can be considered a framework for landscape management. In addition to predefined functions, SAP LaMa can extend the management software with user-defined operations and workflows and tailor the management software for specific needs. SAP LaMa provides the following extension points:

Custom Tabs	Define your own tabs in the SAP LaMa GUI to start more tools.
Custom Services	Detect, manage, and integrate more applications.
Custom Operations	Define your own action buttons in the SAP LaMa GUI.
Custom Hooks	Extend the functions of predefined activities (pre-, post, and error handling).
Custom Provisioning	Provide your own storage replication technology or integrate backup and restore.
Custom Links	Invoke a custom URL with context information.
Custom Validations	Define your own validations to pro-actively scan for problems.
Automation Studio	Create custom processes.

2.3 Setting up an SAP LaMa environment

At first, the SAP systems in scope must be enabled. To enable the environment for all the use cases, certain prerequisites must be fulfilled for the managed environment. The managed environment consists of the following components:

- SAP LaMa management system
- IBM PowerVC server

Figure 2-5 illustrates the steps that are required for the setup of the environment.

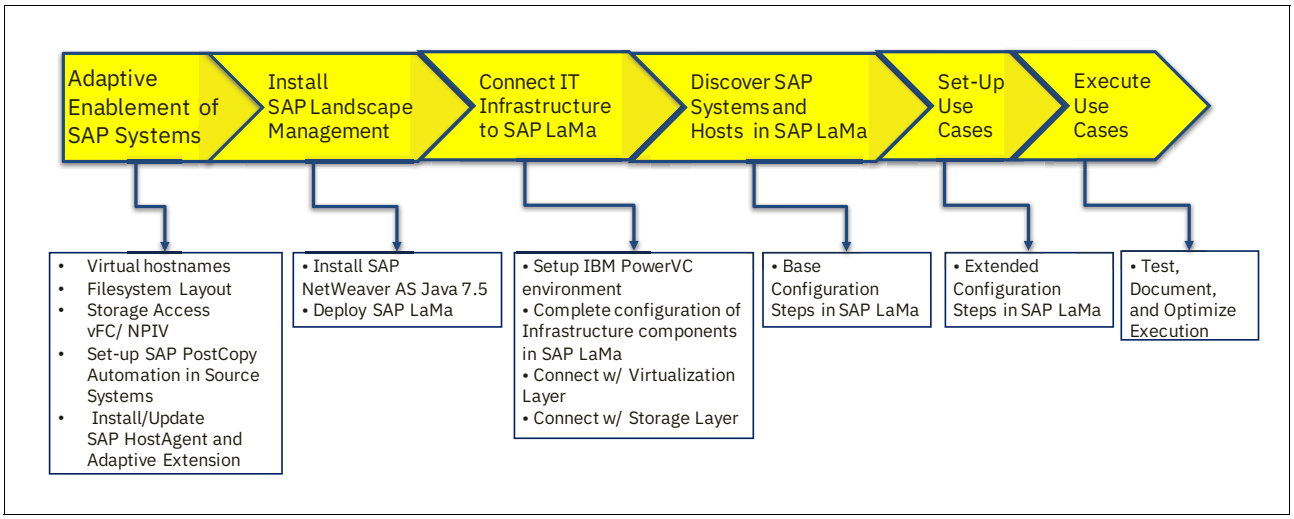

Figure 2-5 Roadmap for setup of SAP LaMa environment

2.3.1 Enabling SAP systems

This section describes the steps to enable SAP systems.

Preparing the managed hosts

On each LPAR that will be integrated into and managed with SAP LaMa, you must run the following preparation steps on the OS level:

- Installing the SAP Host Agent and SAP Adaptive Extensions.
- Prepare firewall fencing for target hosts.

Installing the SAP Host Agent and SAP Adaptive Extensions on all LPARs

Check or install the SAP Host Agent and the SAP Adaptive Extensions on all LPARs that are part of the environment.

Checking the version of the SAP Host Agent

As the root user, run the following command:

```
cd /usr/sap/hostctrl/exe
./saphostexec –version
```

The available version of the SAP Host Agent at the time of writing is SAP Host Agent 7.21 SP40.

If an update is required, complete the following steps:

1. Download the current patch of `SAPHOSTAGENT.SAR` for the OS platform (IBM AIX, Linux on Power Little Endian, or Linux on Power Big Endian) from the SAP Support Software Download Portal.

   On the website, select **Support Packages and Patches** → **By Category** → **SAP Technology Components** → **SAP Host Agent 7.21**.

2. Choose the appropriate OS platform:

   As the root user, copy the downloaded `SAPHOSTAGENT<SP-version>.SAR` archive to a temporary directory, and then run the SAP Host Agent **`upgrade`** command:

   ```
   /usr/sap/hostctrl/saphostexec -upgrade -archive <path to downloaded
   SAPHOSTAGENT<SP-version>.SAR>
   ```

Checking the version of the SAP Adaptive Extensions

To check the version of the SAP Adaptive Extensions, see SAP Note 1808793 - SAP Adaptive Extension Improvements and Fixes.

In addition to the SAP Host Agent, the IBM libraries and scripts for integrating the IBM components into the SAP LaMa are required. The libraries are delivered as part of the `SAPACEXT` package. The package contains the following parts:

- IBM AIX or Linux platform library (`libsaposprep.o` or `libsapacosprep.so`). The library includes OS-specific functions for activating and deactivating virtual IP addresses and handling of IPSec and `iptables` firewalls.
- IBM storage library (`libsapacosprep_ibm.o` or `libsapacosprep_ibm.so`). The library includes OS-specific functions for mounting and unmounting file systems.

To install, upgrade, or roll back SAP Adaptive Extensions, see SAP Help Portal.

Download the appropriate packages for the OS, which in this case is Linux on IBM Power LE, as shown in Figure 2-6.

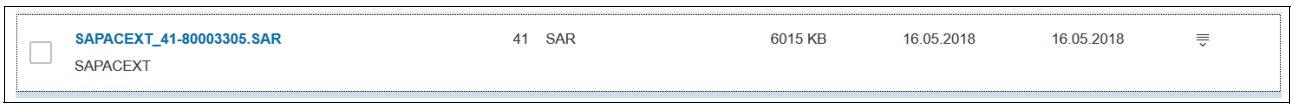

Figure 2-6 SAP software download for SAP Adaptive Extensions

The libraries are extracted to the directory `/usr/sap/hostctrl/exe`. More files are extracted to `/usr/sap/hostctrl/exe/operations.d` and `/usr/sap/hostctrl/exe/operations.d/IBM`. Create an extra directory for the trace and log files of the scripts in all the systems. The default location for the trace and log files is `/usr/sap/hostctrl/exe/operations.d/IBM/traces`. Ensure that the files under `operations.d` belong to the root user and are only writeable by the root user.

Checking and creating the trace directory

Create a directory for the trace and log files of the scripts in all the systems. The default location for the trace and log files is `/usr/sap/hostctrl/exe/operations.d/IBM/traces`.

Checking the iptables packages on the target LPARs

SAP LaMa uses the TCP/IP filtering capabilities that are bundled with the OS to isolate the target SAP system after the system copy (unless the PostCopy Automation is completed). On Linux, the TCP/IP filters are included in the `iptables` packages. Those packages must be installed in the OS for the target LPARs.

Preparing the SAP system

The SAP NetWeaver systems must be adaptively installed. An adaptively installed SAP NetWeaver system has the following attributes:

- Installed with a suitable file system layout that decouples the SAP instance from the server.
- Can independently relocate the different instances (database, central services, and primary application server) to different target LPARs.
- Installed by using virtual IP addresses.
- SAP-related OS users and groups have UIDs and GIDs that are identical and unique across the entire landscape.

Suitable file system layout

Figure 2-7 illustrates an example file system layout for a central SAP S/4HANA system.

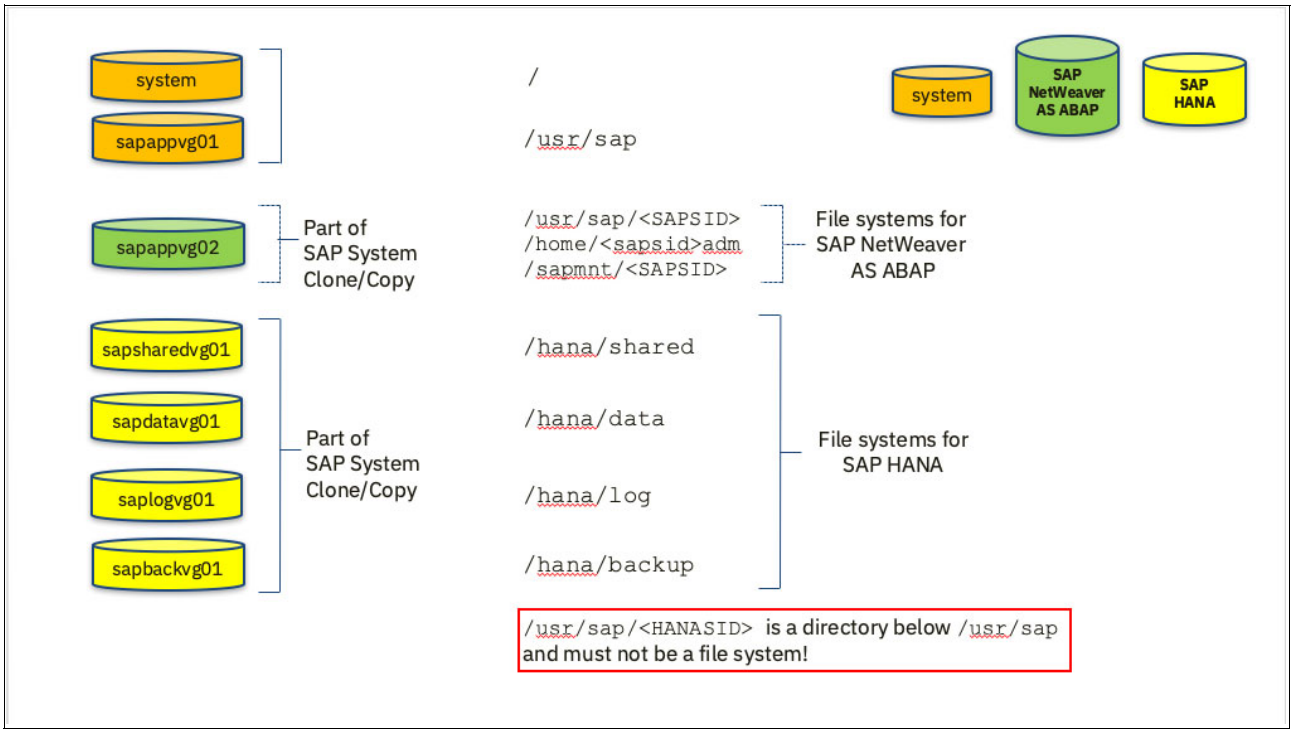

Figure 2-7 File system layout for SAP HANA that is suitable for SAP System Copy

The root file system / and the file system `/usr/sap` belong to the host and are not part of any relocation or SAP System Clone or SAP System Copy scenario. The SAP Host Agent and SAP Adaptive Extensions are installed in directories under the `/usr/sap` file system tree. The `/usr/sap` is either a file system or directory within the system disk that eventually is deployed to an extra volume. However, the file system `/usr/sap` always belongs to the host.

The SAP NetWeaver AS Advanced Business Application Programming (ABAP) has the components SAP Central Service and SAP priMary Application Server. They are installed in the directory trees under `/usr/sap/<SAPSID>`, and require the `SAPMOUNT` share `/sapmnt/<SAPSID>`.

The SAP HANA database is installed in `/hana/shared`, with the data volumes in `/hana/data`, log volumes in `/hana/log`, and a backup directory that is named `/hana/backup`. For recovery reasons, it is best practice to separate the different areas into different disk volumes.

For an SAP System Copy and SAP System Refresh scenario, it is important that `/usr/sap/<HANASID>` is a directory and not a file system.

Virtual IP addresses

Ensure that SAP instances are installed with their own virtual host names. The SAP instances must be decoupled from the host name of the OS on which they are running. All virtual IP addresses must be resolvable by using DNS for SAP LaMa and all participating hosts. There must be a unique relationship between a (virtual) host name and its IP address. Do not use a DNS alias, but assign a unique virtual host name with its own IP address.

Note: All host names and virtual host names must be resolvable uniquely for both SAP LaMa and the LPARs by using a *DNS A record* and not a *DNS CNAME*.

Managing a central user

Ensure that required SAP OS users are (or will become) available and that the OS user IDs and OS group IDs are unique and consistent across the landscape. Possible options can be Lightweight Directory Access Protocol (LDAP), Network Information Service (NIS), or other customer-specific solutions.

Installing SAP LaMa Enterprise Edition in the source system

For information about this topic, see SAP Note 19121110 - Install SAP Landscape Management Enterprise Edition License.

PCA task lists are used to manage all the steps that must be run before an SAP System Refresh or after an SAP System Copy or SAP System Refresh procedure.

PCA is part of the SAP LaMa Enterprise Edition software. The PCA content is integrated in the SAP NetWeaver (AS ABAP) software packages and is deactivated per default. For more information, see *ABAP Post-Copy Automation Installation Guide*. You can also use task list `SAP_INSTALL_PCAI_ENT` to install the add-on PCAI_ENT 1.0, which enables the execution of ABAP PCA task lists.

2.3.2 Installing the SAP Landscape Management software

A detailed description of an SAP LaMa installation is out of scope for this publication. For more information, see the appropriate documentation from SAP.

Before you start the installation, you must install an SAP system that is based on SAP NetWeaver AS Java. Download the appropriate installation guide and software packages for the target OS and database platform combination to install the SAP NetWeaver Java System. For more information, see SAP Note 2350235 - SAP Landscape Management 3.0 - Enterprise Edition.

After SAP NetWeaver AS Java is installed, the SAP LaMa software must be deployed. SAP LaMa software is delivered as a set of SCA files that contain the Java archives that are shown in Table 2-2. Deploy those SCA files by using the SAP Software Update Manager, which is as described in the *SAP Landscape Management 3.0, Enterprise Edition Installation Guide*.

Table 2-2 SCA installation files

SCA file	Description
VCM	This SCA file contains all functional coding and is the main component of SAP LaMa.
VCMCR4E	This SCA file contains open source and third-party libraries that are required to use the search functions or to generate and display reports in SAP LaMa.
VCMENT	This SCA file activates the enterprise functions within SAP LaMa.

After the software is deployed, perform the initial setup by using the configuration wizard, as described in the “Post-Installation” section of *SAP Landscape Management 3.0, Enterprise Edition Installation Guide*. For more information, see the following SAP sections in that document:

- Initial Setup Using the Configuration Wizard
- Landscape Configuration
- Landscape Preparation

2.3.3 Installing and configuring an IBM PowerVC environment

The installation and configuration of IBM PowerVC is out of scope for this publication. The assumption here is that IBM PowerVC is installed and configured, and all server and storage infrastructure components (HMC, Storage, SAN, and so on) are attached to it. For more information about how to set up an IBM PowerVC environment, see IBM PowerVC. You can also refer to *IBM PowerVC Version 1.3.2 Introduction and Configuration*, SG24-8199.

Creating an admin user for SAP LaMa on IBM PowerVC

An IBM PowerVC user with *administrator privileges* is required because SAP LaMa uses this access to run monitoring and management tasks in IBM PowerVC.

Creating a project for SAP LaMa in IBM PowerVC

A *project*, sometimes referred to as a *tenant*, is a unit of ownership. Most resources in IBM PowerVC, such as VMs, volumes, and images, belong to a specific project. Only users with a role assignment for a project can work with the resources belonging to that project. The `ibm-default` project is created during installation, but IBM PowerVC supports the creation of more projects for resource segregation. As a best practice, use an IBM PowerVC project and add only those resources to that project that are visible in SAP LaMa.

Preparing LPARs and storage volumes in IBM PowerVC after onboarding

If LPARs and their storage volumes were created by using tools outside of IBM PowerVC, they must be imported first. The importation is a default procedure that is described in IBM Knowledge Center. The following post-configuration steps are required:

- LPAR: Ensure that LPAR has at least one network adapter and an IP address. The IP address is used for association between OS data (retrieved by SAP Host Agent) and LPAR data that is gathered by IBM PowerVC. Add the IP address to the LPAR configuration in IBM PowerVC UI.
- Storage volume: Ensure that storage volume is assigned to a storage template. This initial assignment is done by using the IBM PowerVC CLI, as described in How to set the storage template when managing existing volumes.

2.3.4 Preparing SAP NetWeaver AS Java for communication with IBM PowerVC

This section describes the steps that are needed so that SAP NetWeaver may communicate with IBM PowerVC.

Checking and updating the SSL library within the SAP NetWeaver AS Java server

IBM PowerVC requires the usage of secure communication protocols TLS V1.1 or TLS V1.2 for any external communication. Older protocol versions have security exposures, so they are no longer supported and are disabled. Check that the SAP NetWeaver AS Java engine where the SAP LaMa code is running supports those required TLS versions.

For more information, see SAP Note 2284059: Update of SSL library within NW Java server SAP Note 2284059 - Update of SLL library within NW Java server and SAP Note 2569156 - How to get and validate SSLContext.properties file SAP Note 2569156 - How to get and validate SLLContext.properties file.

Changing the SSLCipherSuite in the SAP NetWeaver engine

Add the cipher suites to the property file `SSLConfig.properties`, as described in SAP Note 2284059 - Update of SLL library within NW Java server. Add the following lines:

```
cipherSuite=TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
cipherSuite=TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
```

Importing the IBM PowerVC certificate into a truststore of SAP NetWeaver AS Java

A communication setup that uses *server certificates* provides the most secure level of communication between SAP LaMa and IBM PowerVC. In this case, the IBM PowerVC certificate is used and the host name that is contained in the certificate is verified against the host name of the public key (which is stored in the SAP NetWeaver AS Java keystore), which is defined within the SAP LaMa virtualization manager definition.

The IBM PowerVC server certificate must be imported first into the keystore of the underlying SAP NetWeaver AS Java engine. For more information about how to import the public key of the self-signed IBM PowerVC certificate into the keystore, see Using the AS Java Key Storage. Check that the host name in the certificate is valid and matches the host name of the IBM PowerVC server. The host name that is referenced in the SAP LaMa virtualization adapter configuration also must correspond to the host name in the certificate. Remember the name of the certificate store view of SAP NetWeaver AS Java because this name is required during the IBM PowerVC adapter definition in SAP LaMa.

By default, IBM PowerVC uses a self-signed X.509 certificate to secure its web interface and REST APIs. Self-signed certificates are not signed by an independent certificate authority (CA). If there are enhanced security requirements, then the default self-signed certificate can be replaced with a new CA-signed certificate.

You can learn to how to obtain and set up CAs by referring to IBM Knowledge Center or Securing IBM PowerVC web interface with self-signed or CA-signed certificates.

Export the certificate file `/etc/pki/tls/certs/powervc.crt` from the IBM PowerVC server and import it into the keystore of the SAP NetWeaver AS Java by using the NetWeaver Administrator (NWA) tool. For more details, see Figure 2-8 and Figure 2-9.

To import the public key, complete the following steps:

1. Log in to NWA and select **Configuration** → **Certificates and Keys**, as shown in Figure 2-8.

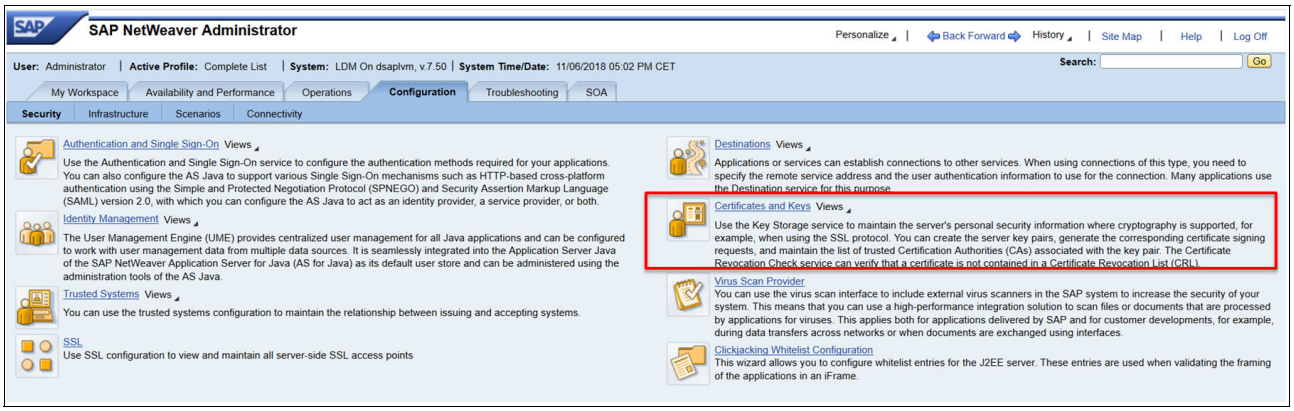

Figure 2-8 NetWeaver Administrator: Certificates and Keys

2. Select the appropriate keystore, for example, TrustedCAs, as shown in Figure 2-9.

Figure 2-9 TrustedCA view

This keystore view name is entered in to the Virtualization Manager or Storage Manager definition later.

3. Import the public key that was exported from the IBM PowerVC server into the keystore (Figure 2-10).

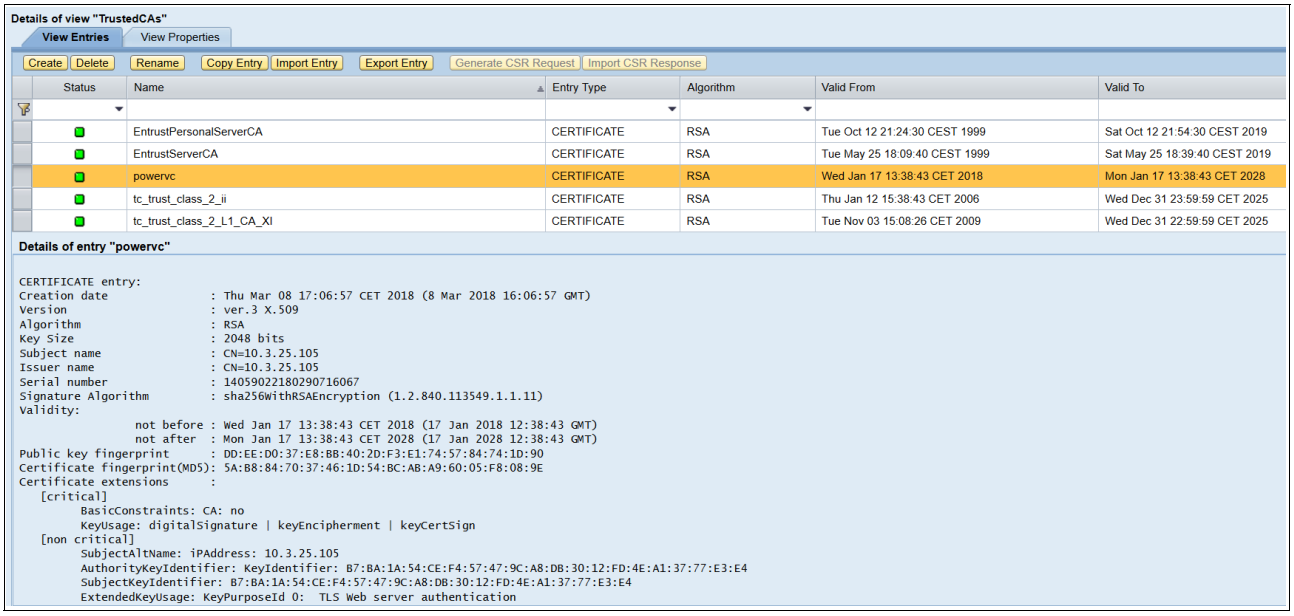

Figure 2-10 SAP NWA: Details of the TrustedCA view

2.3.5 Configuration steps in SAP LaMa

Before you can run the use cases, you must perform a set of configuration steps in SAP LaMa:

1. Describe all network configurations in the SAP LaMa configuration.
2. Set the OS user management settings.
3. Choose the assignment for the OS.
4. Define an SWPM repository for System Rename execution.
5. Define IBM PowerVC as a storage manager or virtualization manager.

In the following sections, these steps are described in more detail.

Networks

The properties for all networks, which are used by virtual or physical host names on managed SAP systems, must be defined in the SAP LaMa network definition. Network-based properties include a network name, the network subnet mask, and the broadcast address for the network.

Advanced configuration settings include the IP address of the gateway server, the DNS zone name, and an *IP range* if SAP LaMa assigns IP addresses and updates DNS entries for virtual host names during an SAP System Clone or SAP System Copy action. In that case, select the **Auto Assign IP Address** option in the System Copy roadmap (Figure 2-29 on page 35). During SAP System Copy execution, SAP LaMa selects an available IP address from the predefined IP range.

The range is defined by the “Update IP range” field of the network configuration:

- IP addresses are in dotted quad notation.
- The lower and upper parts of a range must be in each range and are separated by “-” or “:”.
- Multiple IP ranges can be separated by “;”.

The SAP LaMa network configuration is shown in Figure 2-11.

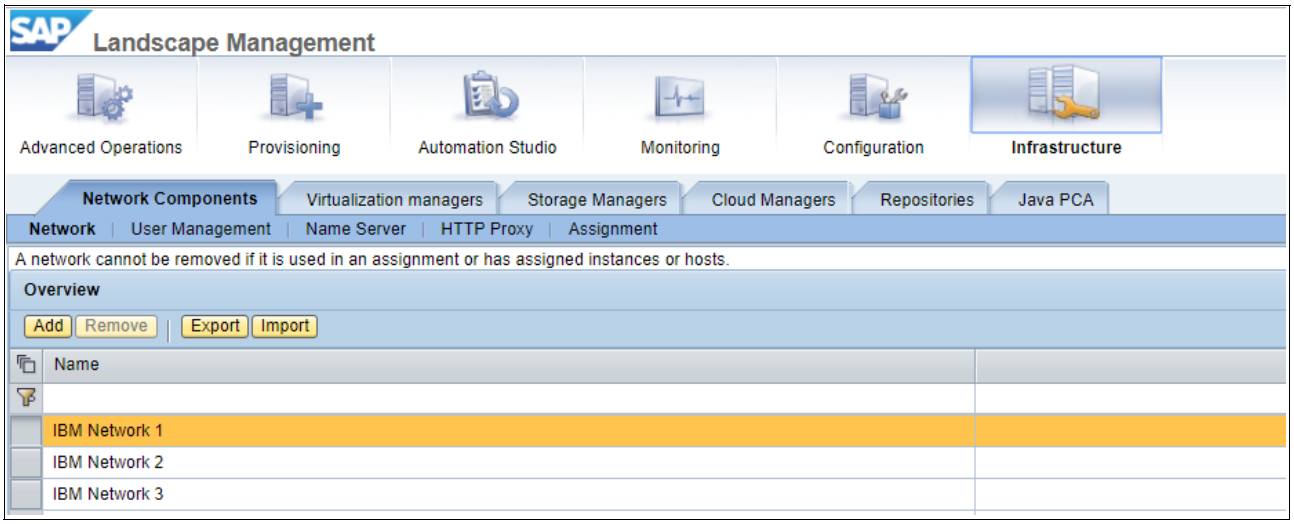

Figure 2-11 SAP LaMa network definition

For more information, see SAP Note 1572841 - Setup of name server update of SAP LVM, which describes the DNS name server integration with SAP LaMa.

User management

This step specifies how OS users and group handling for the SAP systems is managed. OS users that are required but do not exist on the target system are automatically created by SAP LaMa on the target system during a system copy or clone operation. SAP LaMa provides functions for managing users and groups by using LDAP or NIS. These settings must be configured.

Communication with LDAP or NIS software is done by the SAP Host Agent. Configuring LDAP and NIS in LaMa is out of scope of this publication, but it is described in *SAP Landscape Management 3.0 Enterprise Edition*.

Name server

SAP LaMa provides functions for automatic management of DNS entries by using a DNS update. You can set up the DNS management connection. The communication is set up by using the SAP Host Agent. This configuration is out of scope of this publication, but is described in *SAP Landscape Management 3.0 Enterprise Edition*.

Another option is to ensure that the DNS entries that are used for IP addresses and host names exist.

Assignment

After you configure the network, user management, and name server handling, associate these configurations. Then, SAP LaMa take the appropriate actions according to the different networks.

Defining SWPM repositories (AIX and LINUX_PPC64_LE) in SAP LaMa

In use cases that are based on SAP System Copy and SAP System Refresh, SWPM is required. SWPM is used to rename copied SAP System sources.

Ensure that the required SWPM software version is either locally available on the target system or can be mounted by using NFS from a central repository.

In this step, the required SWPM versions are specified according to the OS and NetWeaver releases (Figure 2-12).

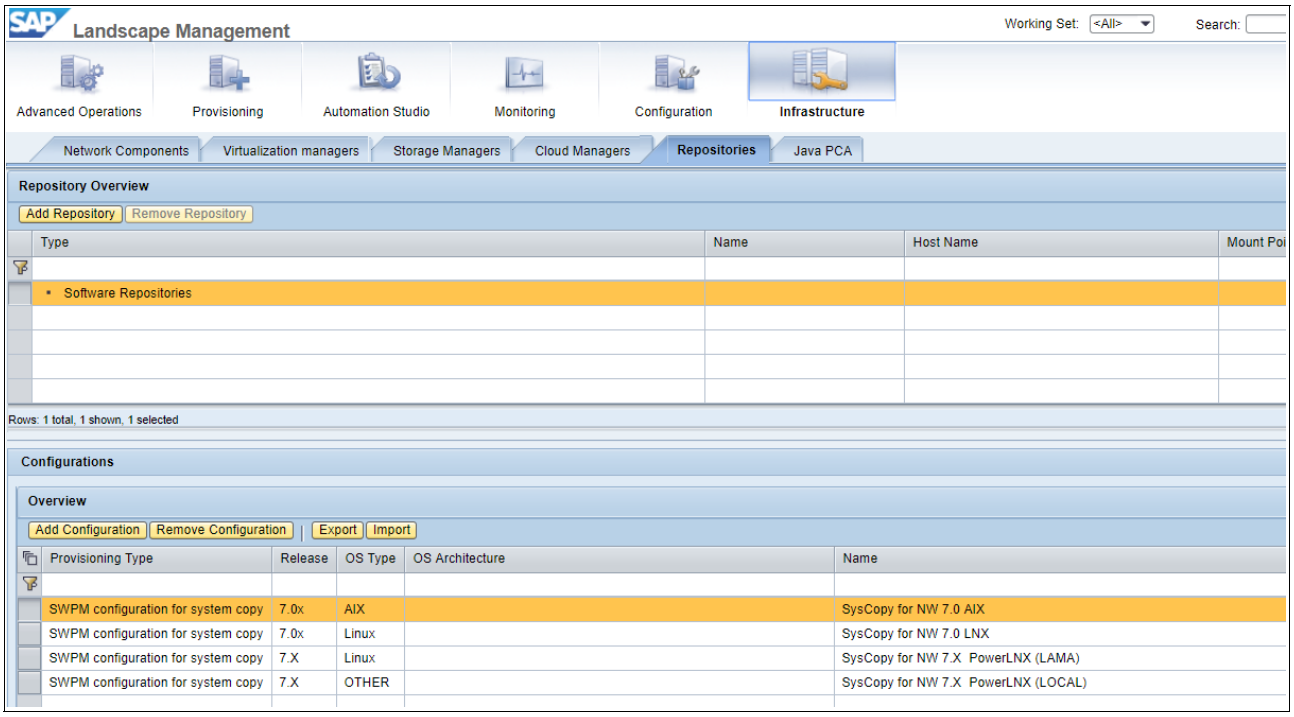

Figure 2-12 SAP LaMa Repository definition

Defining IBM PowerVC as a virtualization adapter in SAP LaMa

Complete the following steps:

1. Log in to the SAP LaMa system and select **Infrastructure** → **Virtualization managers** (Figure 2-13).

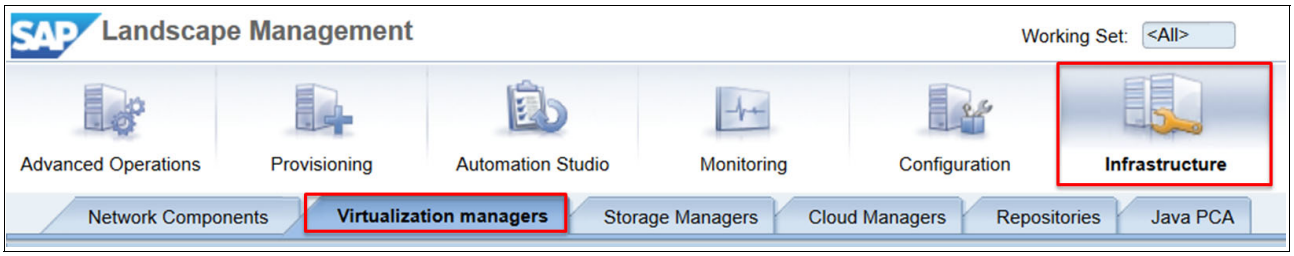

Figure 2-13 SAP LaMa Virtualization manager

2. Select **Add** to add a virtualization manager.

3. Select **IBM PowerVC Adapter** (Figure 2-14) and click **Next**.

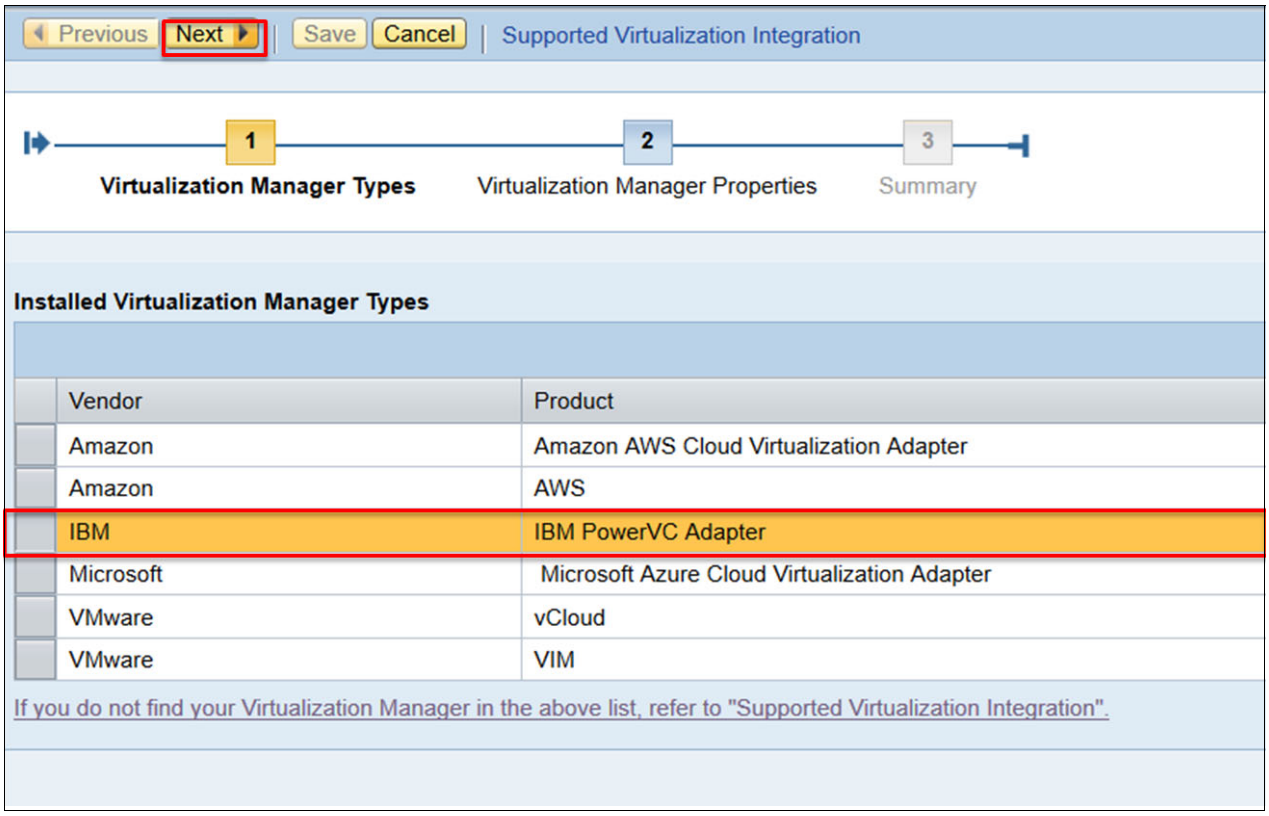

Figure 2-14 Selecting IBM PowerVC Adapter

4. Enter the configuration details (Figure 2-15).

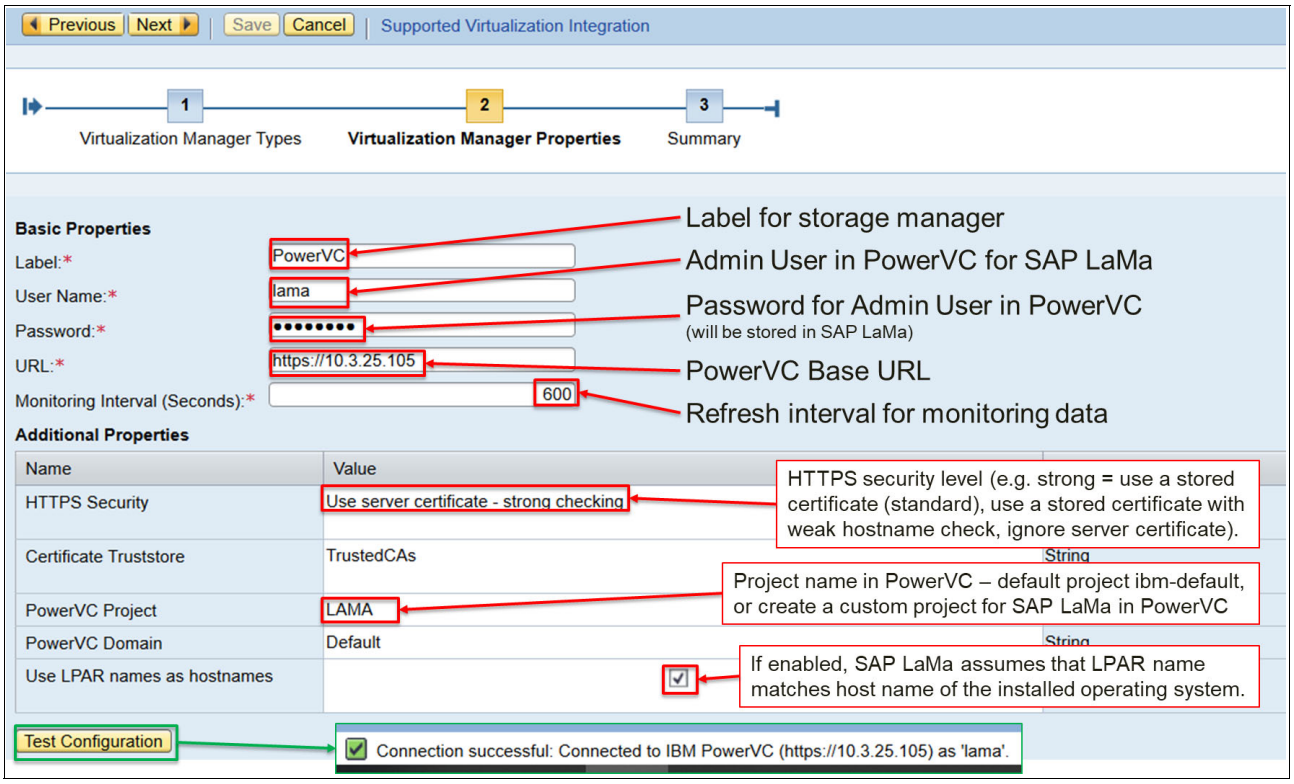

Figure 2-15 IBM PowerVC Virtualization Manager Properties

5. Click **Test Configuration**, and then click **Save**.

Defining IBM PowerVC Storage Manager

The configuration of the IBM PowerVC Storage Manager is done in a way that is similar to the virtualization manager definition that is described in “Defining IBM PowerVC as a virtualization adapter in SAP LaMa” on page 23 (Figure 2-16).

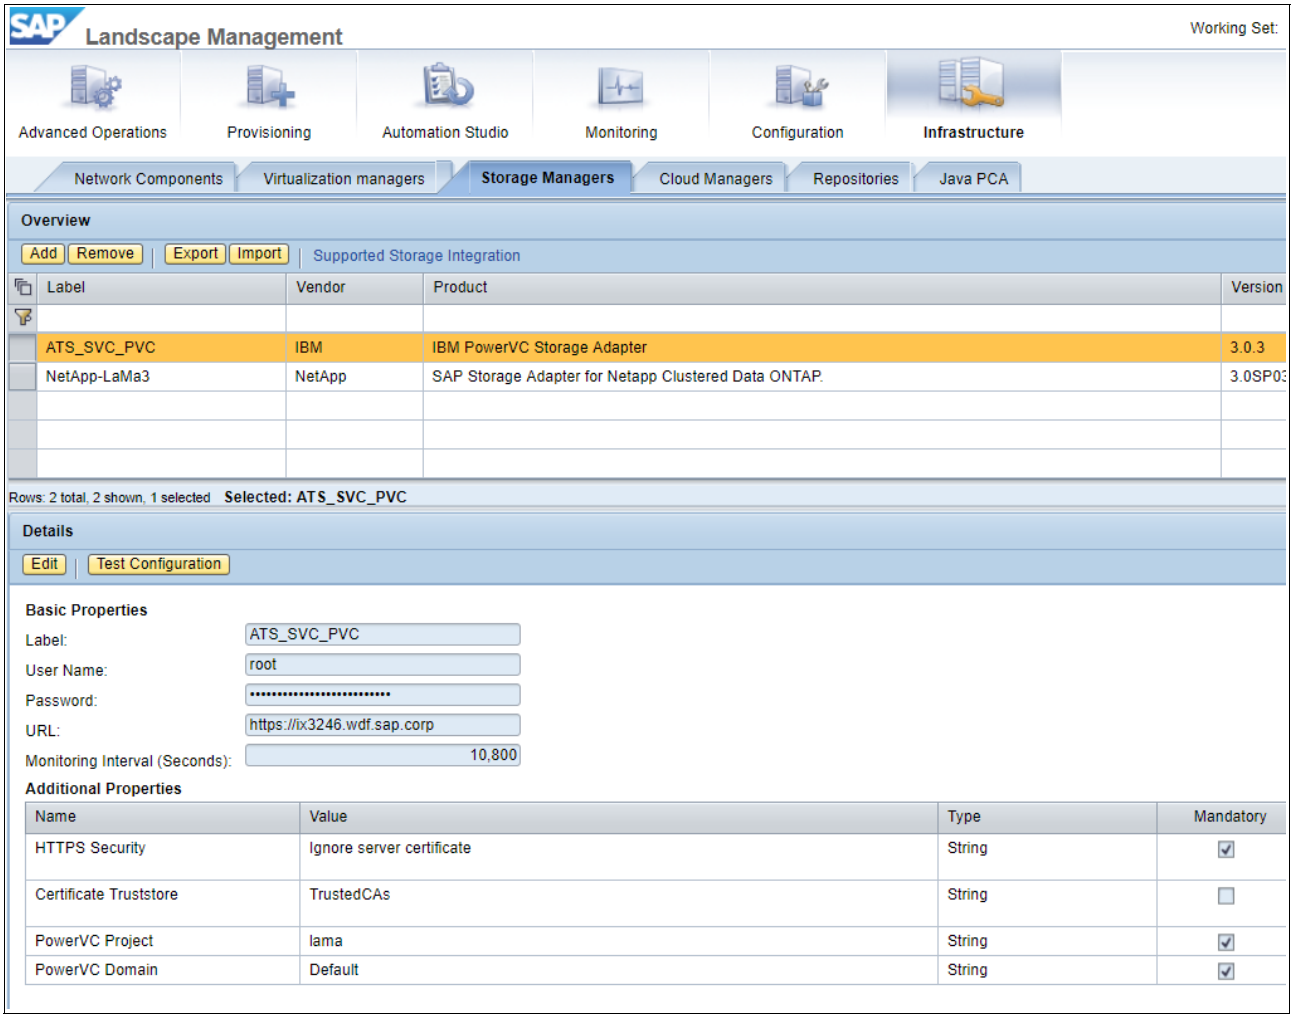

Figure 2-16 IBM PowerVC Storage Manager definition

2.3.6 Discovering managed systems and hosts in SAP LaMa

All SAP systems, their instances, and all the corresponding hosts must be discovered in SAP LaMa. One possibility is running a discovery where SAP LaMa communicates with SAP Host Agent instances running on these hosts. For this communication, you use sapadm user credentials. The discovered entities (SAP systems, SAP instances, and OS hosts) can be logically grouped.

To start discovery, complete the following steps:

1. Select **Configuration** → **Systems** and click **Retrieve new Instances and Hosts**. In the window that is shown in Figure 2-17, specify the detection method (for example, SAP Host Agent) and the entity types to discover (for example, instances and hosts). Depending on the landscape, enable a trigger to detect special entities (for example, shadow or stand-alone DBs, or SAP diagnostic agents).

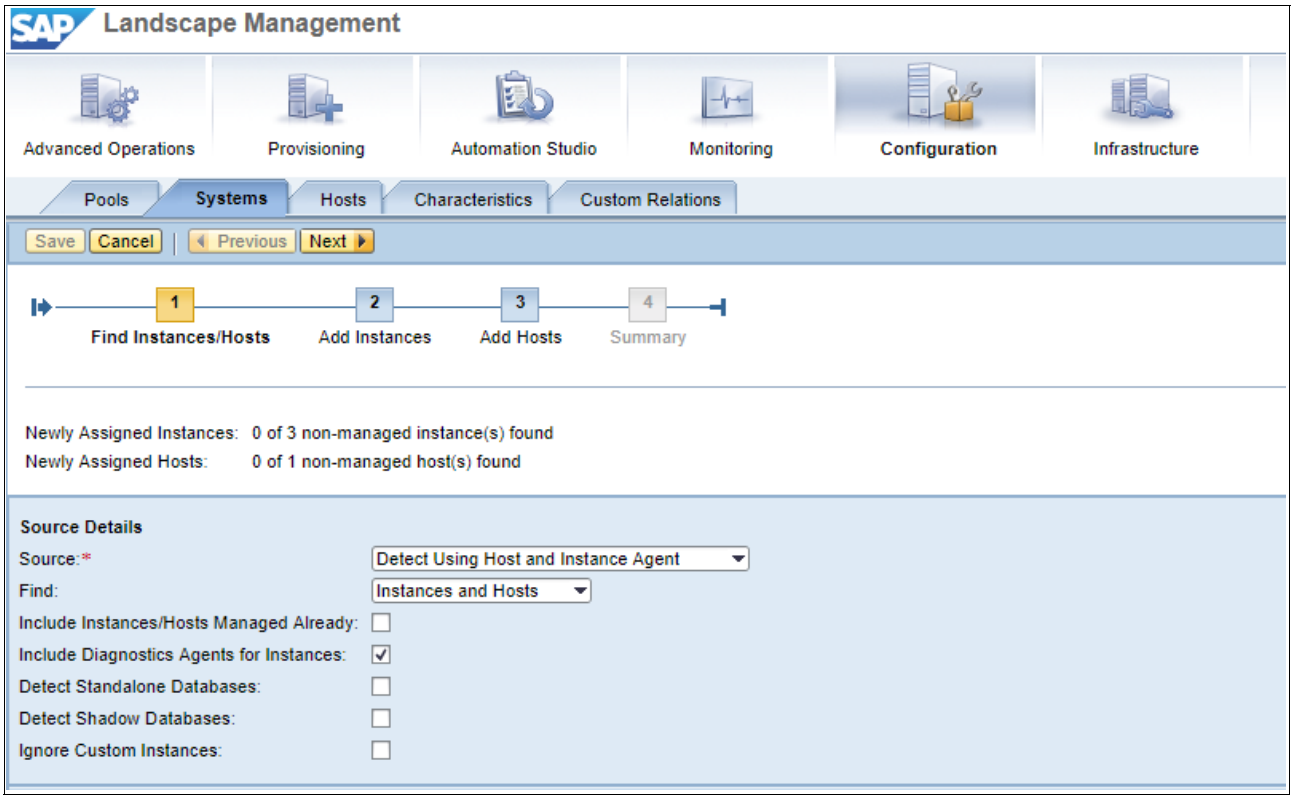

Figure 2-17 SAP LaMa discovery

2. Specify the list of OS hosts that will be scanned (by host name or IP address), the SAP Host Agent port that is used in the landscape, and credentials for communicating with SAP Host Agent, SAP instance agents, and with the SAP DB. Click **Next**, and the discovery process runs (Figure 2-18).

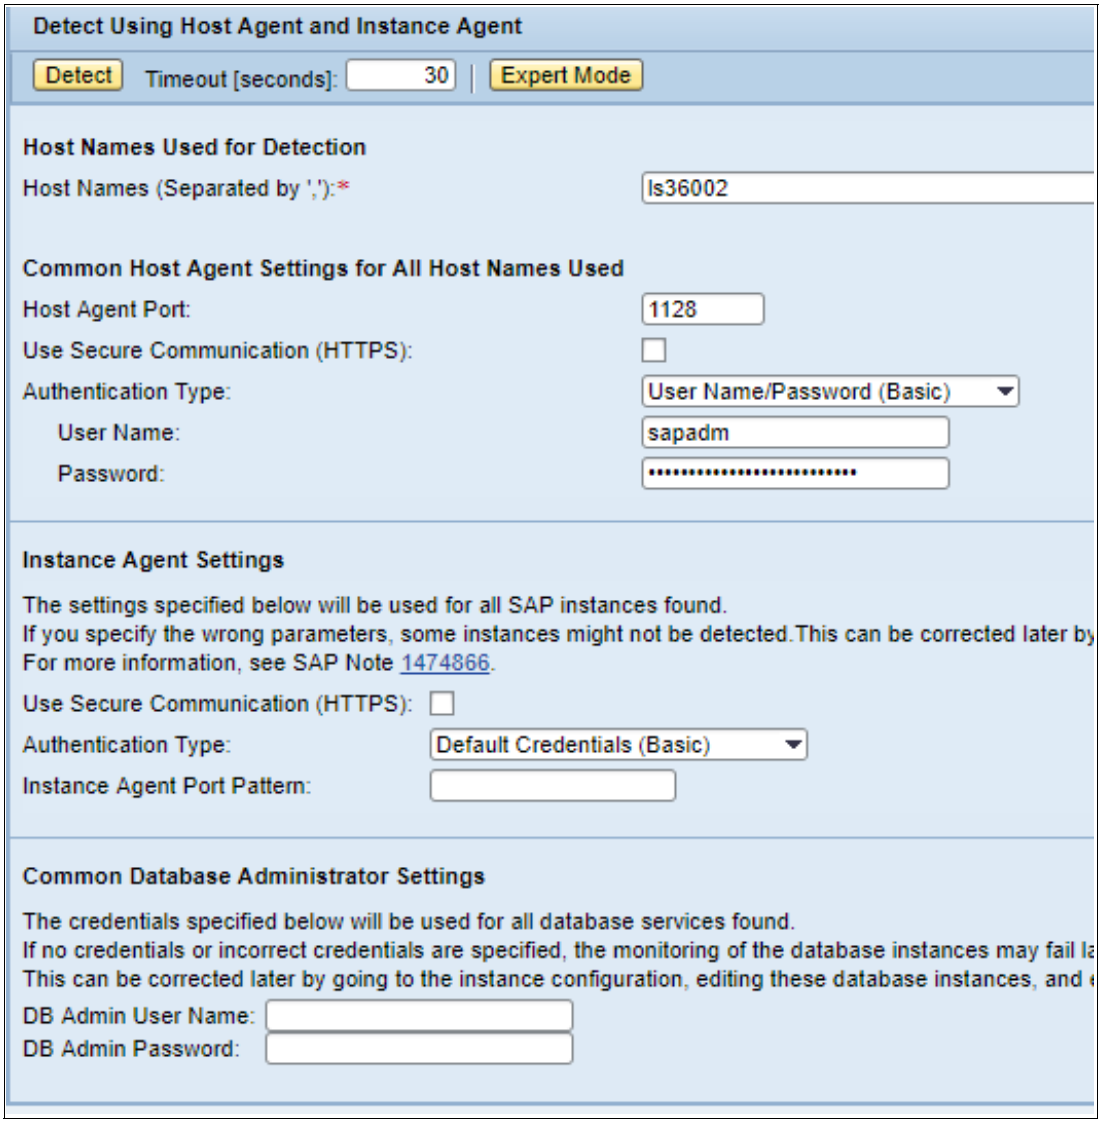

Figure 2-18 Discovery details

3. As result, a list of discovered SAP instances and hosts appear. Assign the instances to SAP systems (for example, by clicking **Auto Assign**) and group the hosts into SAP LaMa pools (see Figure 2-19 and to Figure 2-20). Click **save**, and the entities become managed by SAP LaMa.

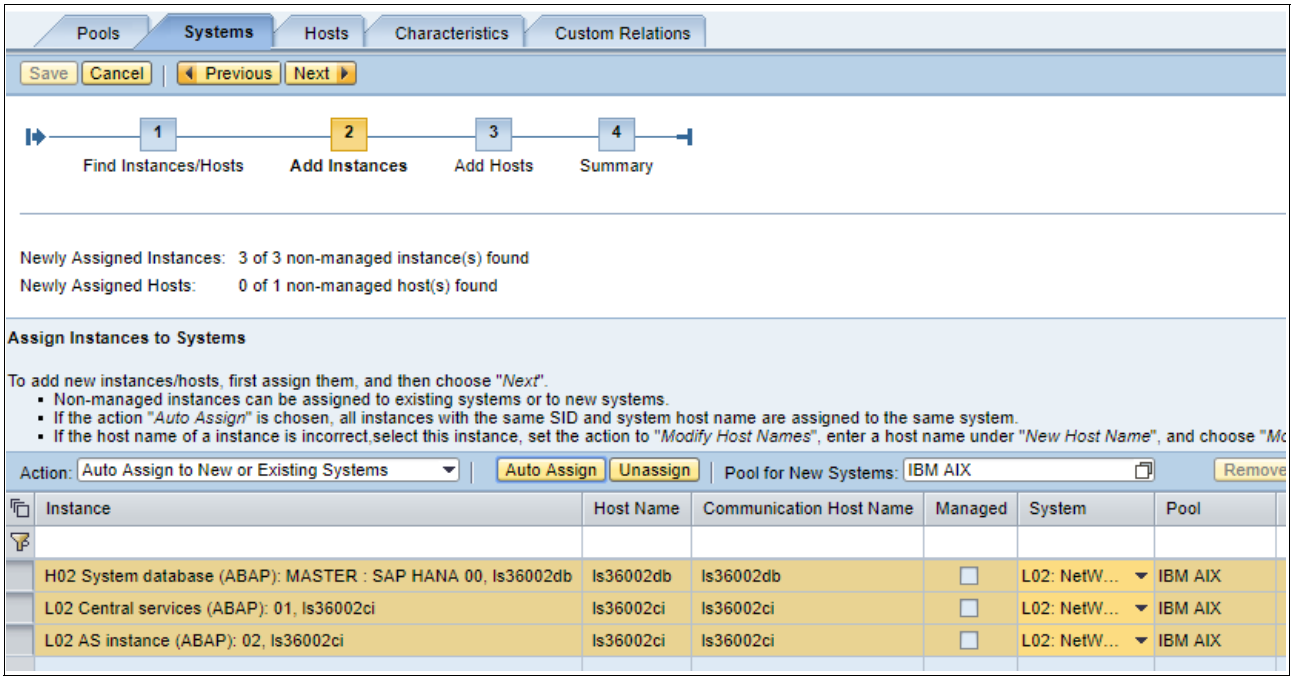

Figure 2-19 Assigning services to an SAP system

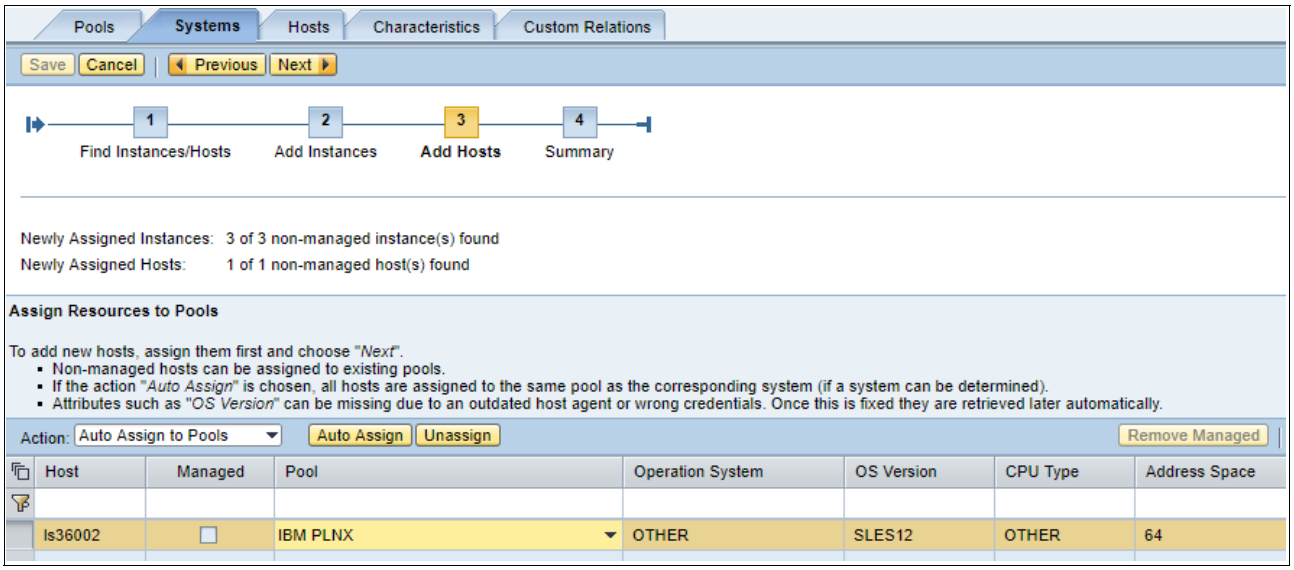

Figure 2-20 Assigning a pool

2.3.7 Configuring SAP hosts in SAP LaMa for copy use cases

For SAP System Clone, SAP System Copy, and SAP System Refresh use cases, you must do more configuration of host entities (see 2.3.1, “Enabling SAP systems” on page 14). Edit the configuration of the SAP hosts and enter the following settings (Figure 2-21):

1. Set SAP hosts to *AC enabled*.
2. Set an SAP host that acts as a target for SAP System Clone or SAP System Copy to *Isolation ready*.
3. Retrieve network interfaces that are used by SAP systems running on these hosts and assign them to the networks that are configured in SAP LaMa.
4. Retrieve *Host-Specific Storage Properties*.
5. Save the configuration.

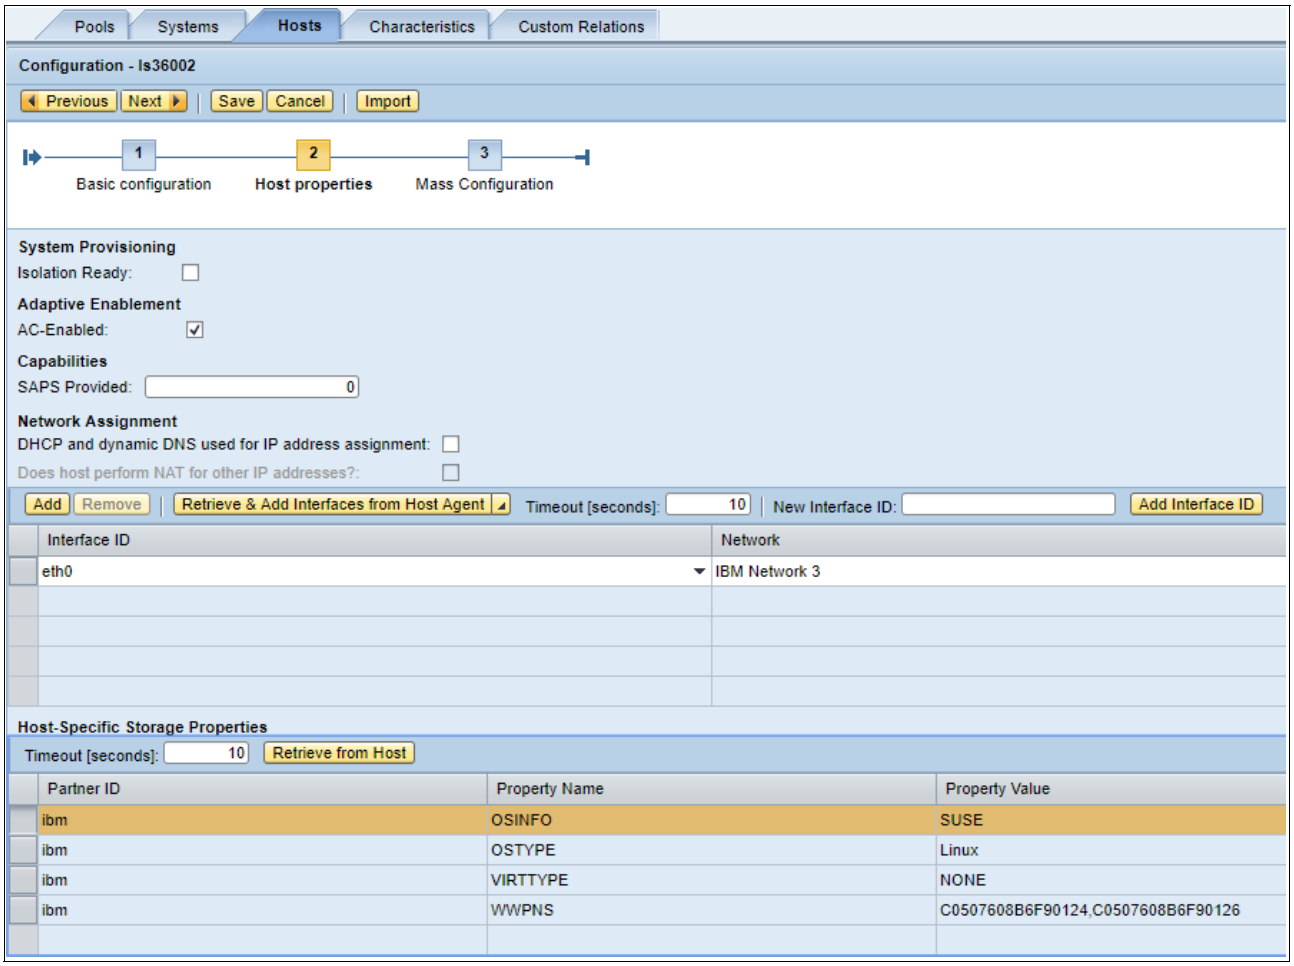

Figure 2-21 Host properties

2.3.8 Configuring SAP systems in SAP LaMa for copy use cases

For SAP System Clone, SAP System Copy, and SAP System Refresh use cases, you must do more configuration of SAP instances and SAP systems (see 2.3.1, “Enabling SAP systems” on page 14). Edit the configuration of each SAP instance and check the following settings (Figure 2-22, Figure 2-23 on page 32, Figure 2-24 on page 32):

- Check that the SAP DB communication credentials are complete and valid (relevant for SAP DB instances only).
- Each SAP instance is *AC Enabled*.
- Check that the *Required Host Type* is correct.
- Deactivate **OS Managed Mounts (Automounter)** and click **Retrieve Mount List**. Check that the list contains only SAP instance-related mount points. Remove all the other retrieved mount points.

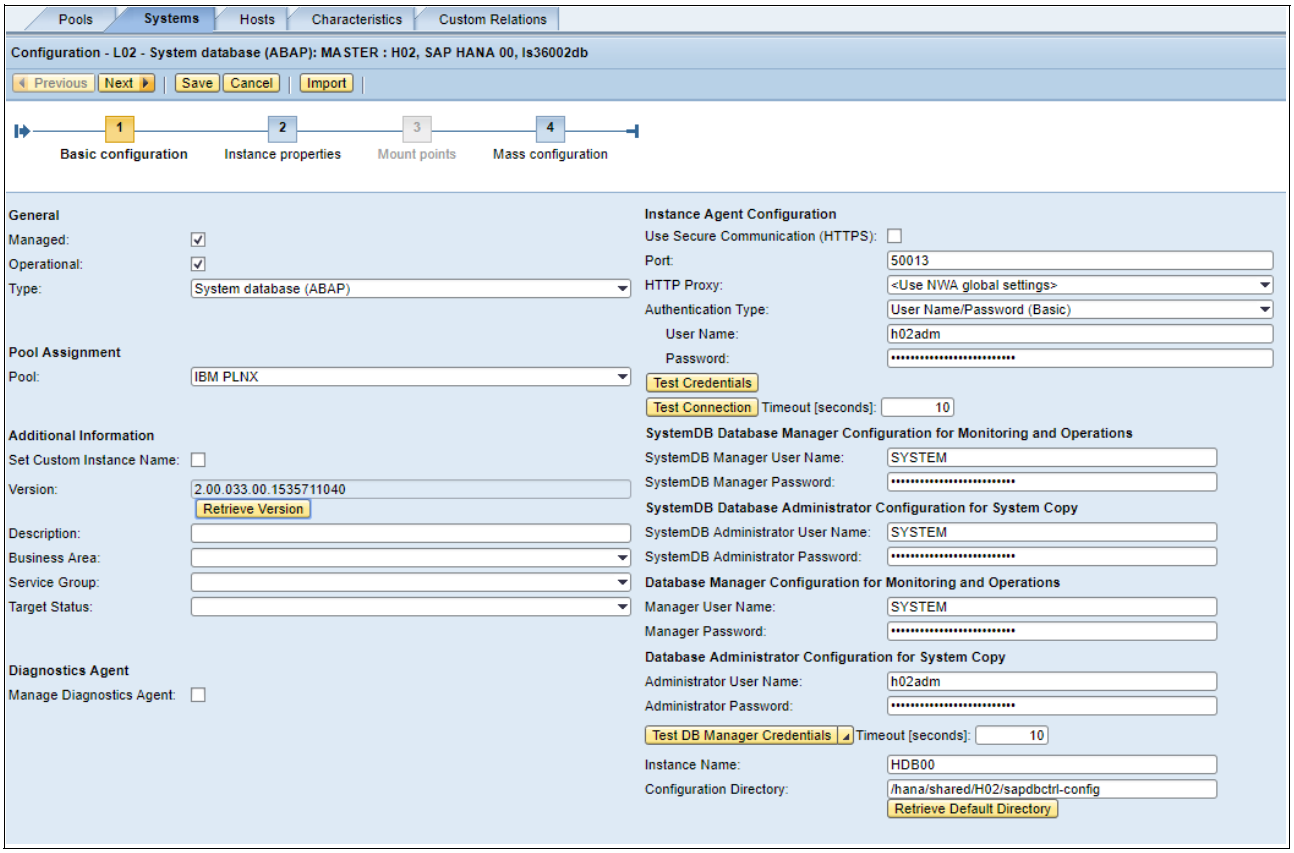

Figure 2-22 SAP instance settings: Basic configuration

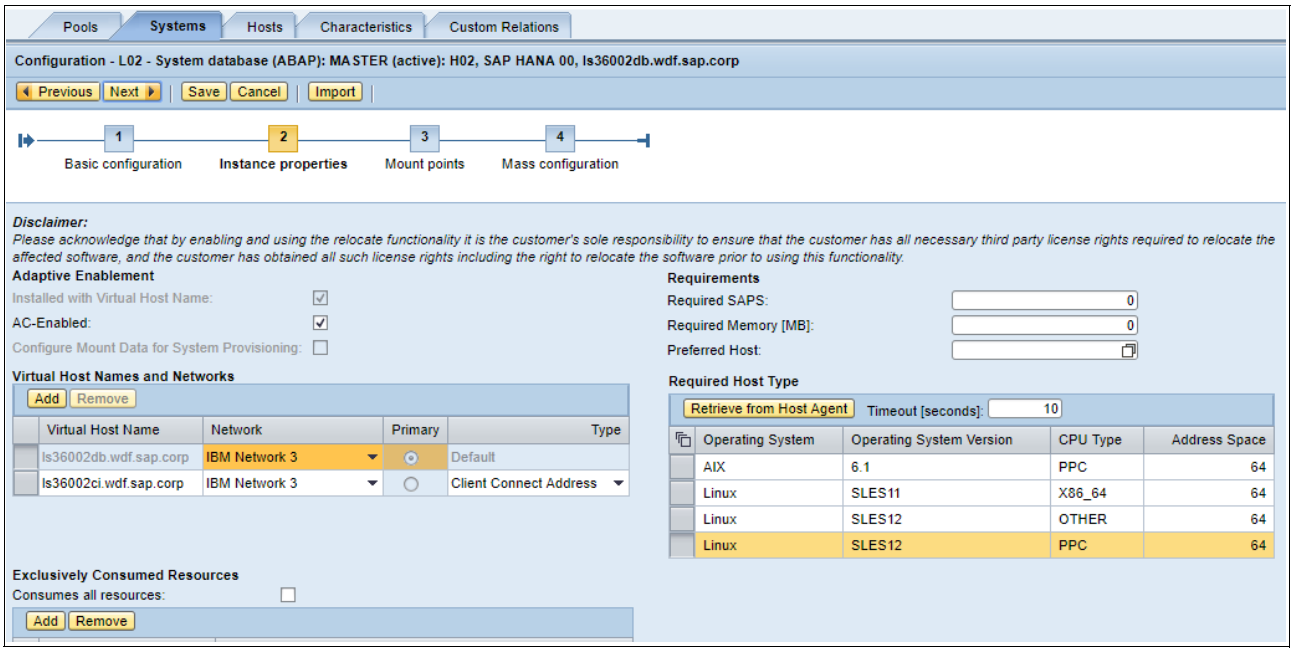

Figure 2-23 SAP instance settings: Instance properties

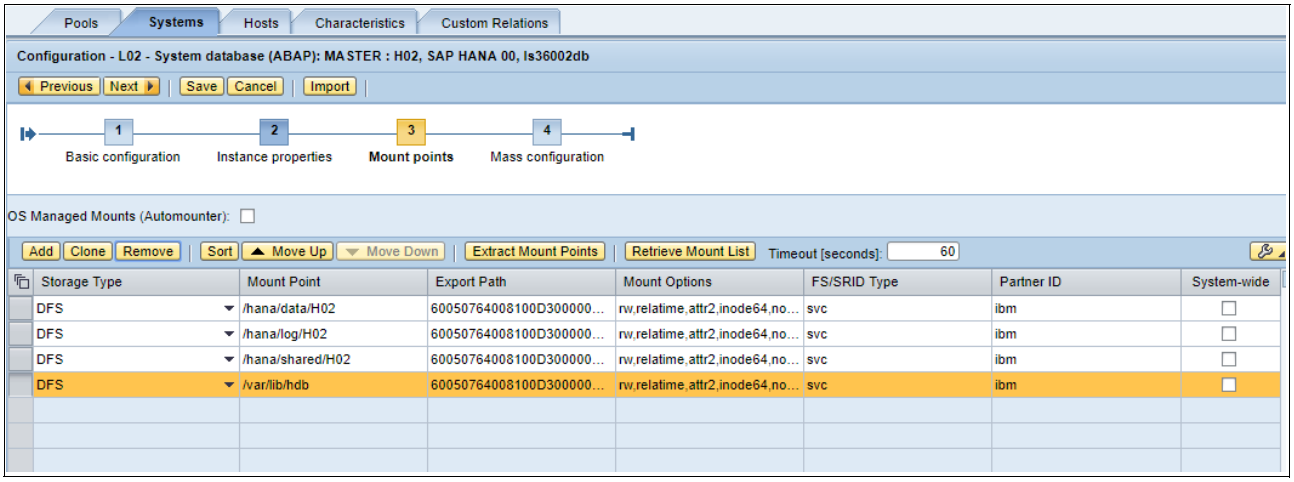

Figure 2-24 SAP instance settings: Mount points

Edit the configuration of the SAP system and check the following items:

- The SAP system can be used for cloning, copying, renaming, or other use cases depending on the relevant scenario.
- For copying, renaming, and refreshing scenarios, SAP PCA is triggered, so configure RFC configuration for these cases (Figure 2-25 on page 33).

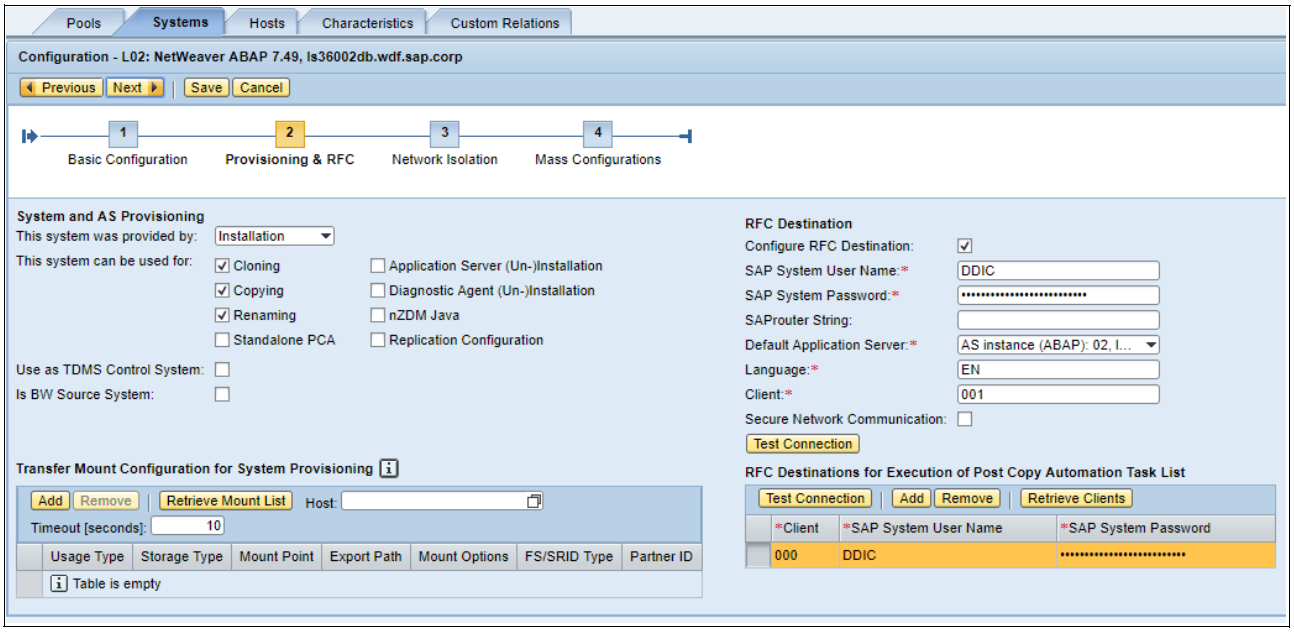

Figure 2-25 Instance settings (RFC user)

2.4 Running SAP System Copy in SAP LaMa

This section describes how to copy an SAP system in SAP LaMa. Other use cases (SAP System Clone and SAP System Refresh) are run in a similar way. The only difference is the requested configuration input in process windows. To initiate an SAP System Copy process, complete the following steps:

1. Click **Provisioning**, select the SAP source system, and click **Copy System** (Figure 2-26).

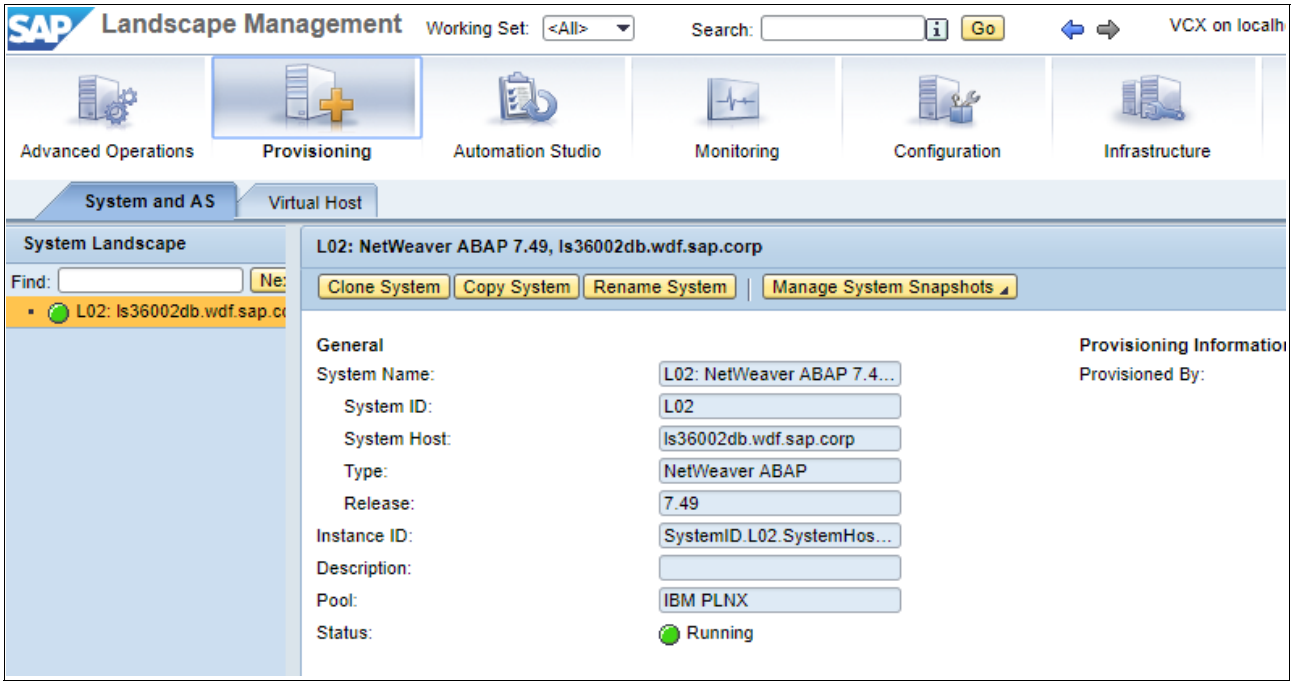

Figure 2-26 SAP System Copy roadmap

2. Specify data for the SAP target (SAP target SID, SAP target DB SID, and SAP LaMa pool for the target system). Provide the credentials for communicating with the system (Figure 2-27).

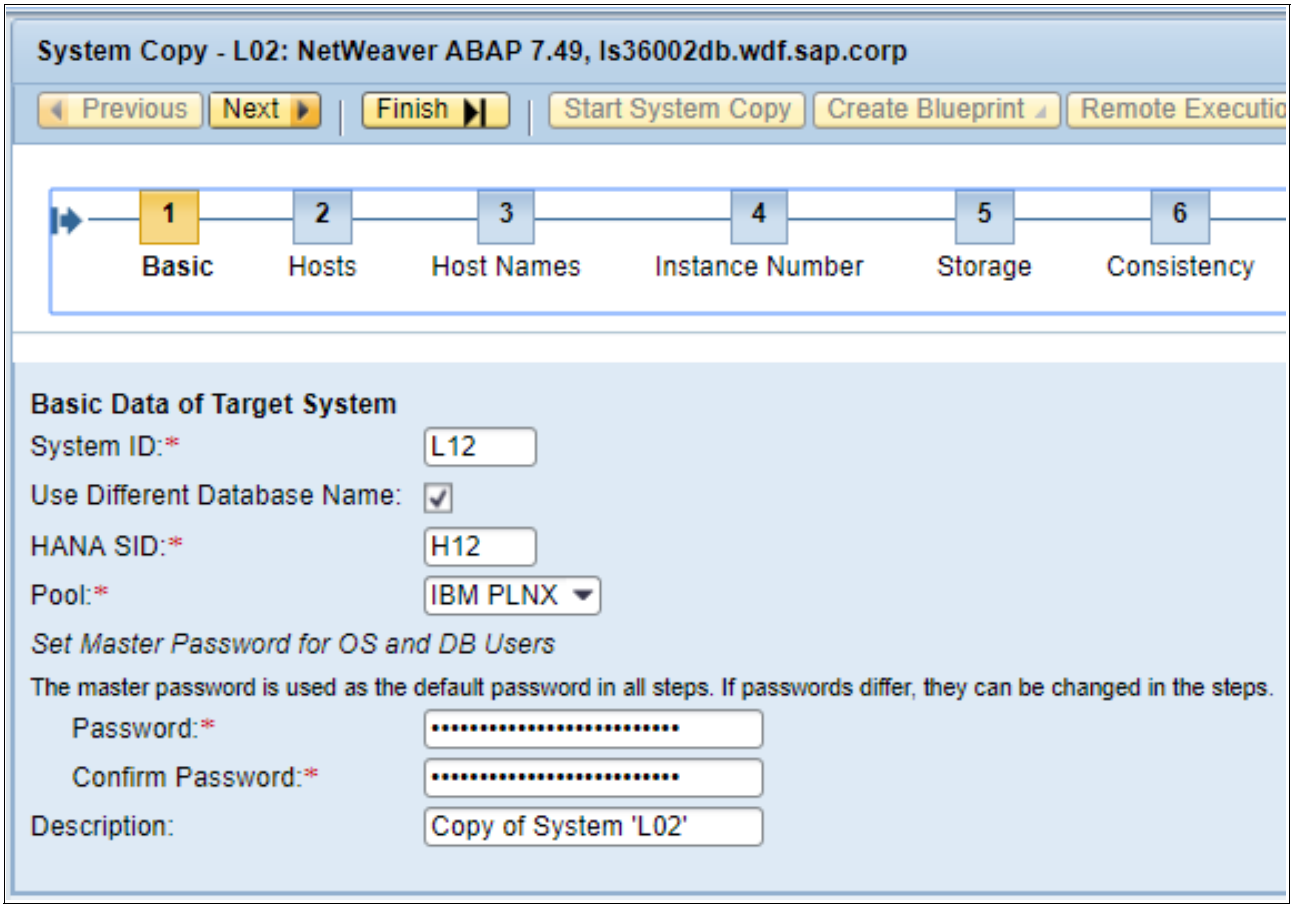

Figure 2-27 SAP System Copy roadmap: Basic data

3. Choose the host for the SAP target (Figure 2-28). This SAP host must be managed by SAP LaMa and configured as *AC Enabled* and `Isolation ready`. For more information and requirements, see 2.3.1, “Enabling SAP systems” on page 14.

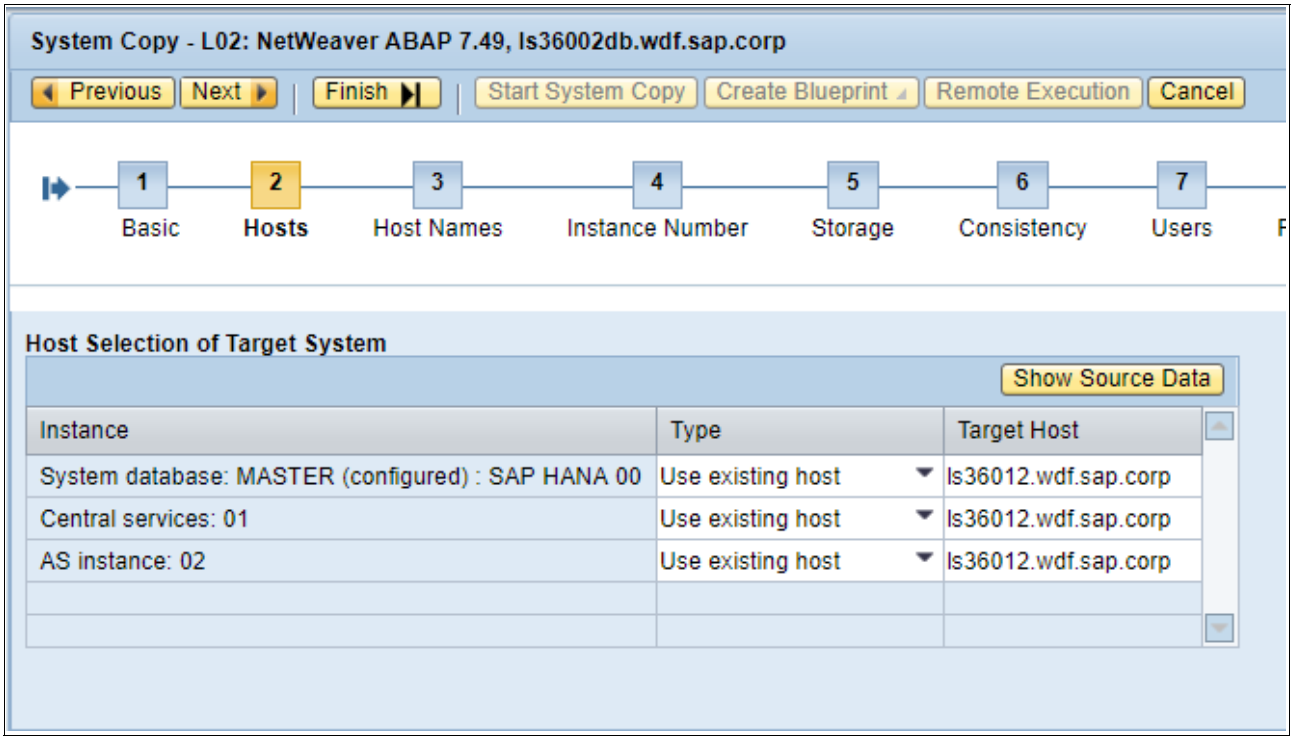

Figure 2-28 SAP System Copy roadmap: Selecting the target hosts

4. Specify the virtual host names for the SAP target instances (Figure 2-29). Virtual host names can be shared among the instances of the same system (if they are in the same host).

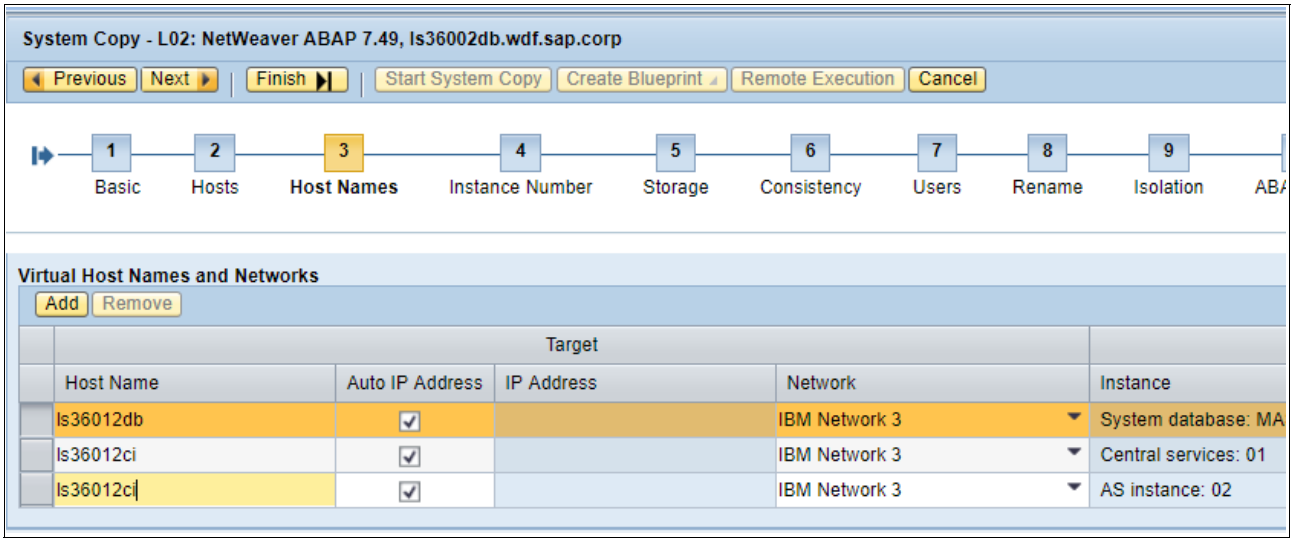

Figure 2-29 SAP System Copy roadmap: Defining the virtual host names

5. Define the instance numbers for the instances (Figure 2-30).

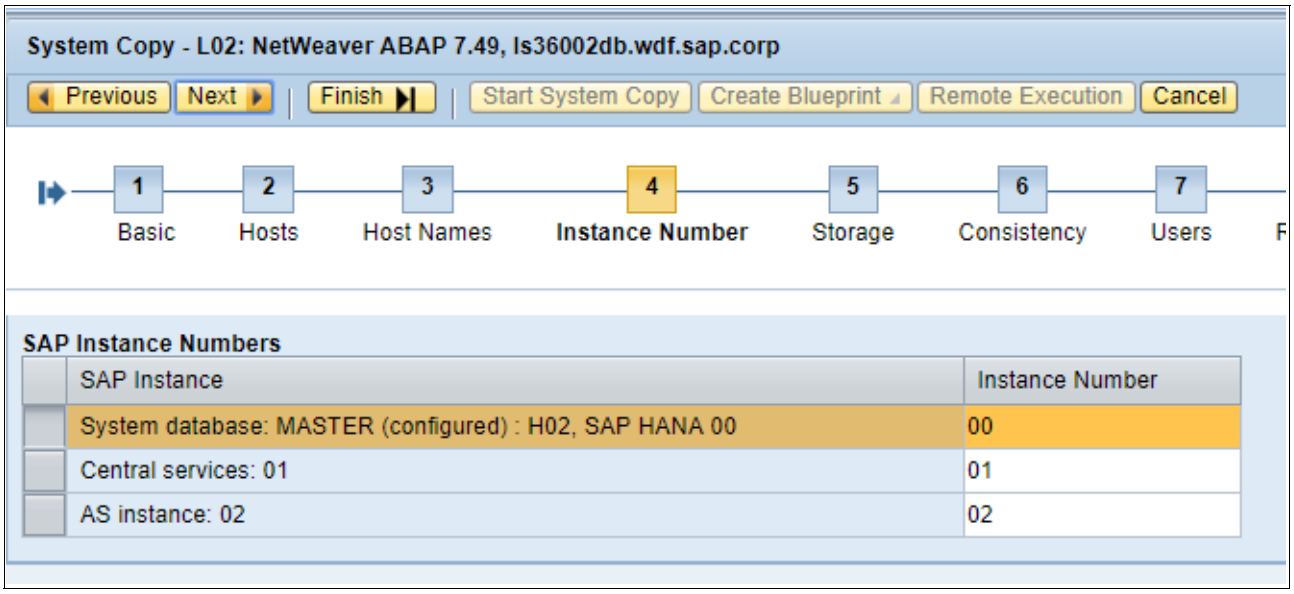

Figure 2-30 SAP System Copy roadmap: Defining the SAP instance numbers

6. Check the storage configuration (Figure 2-31). You can change the proposed names for OS volume groups and OS logical volumes.

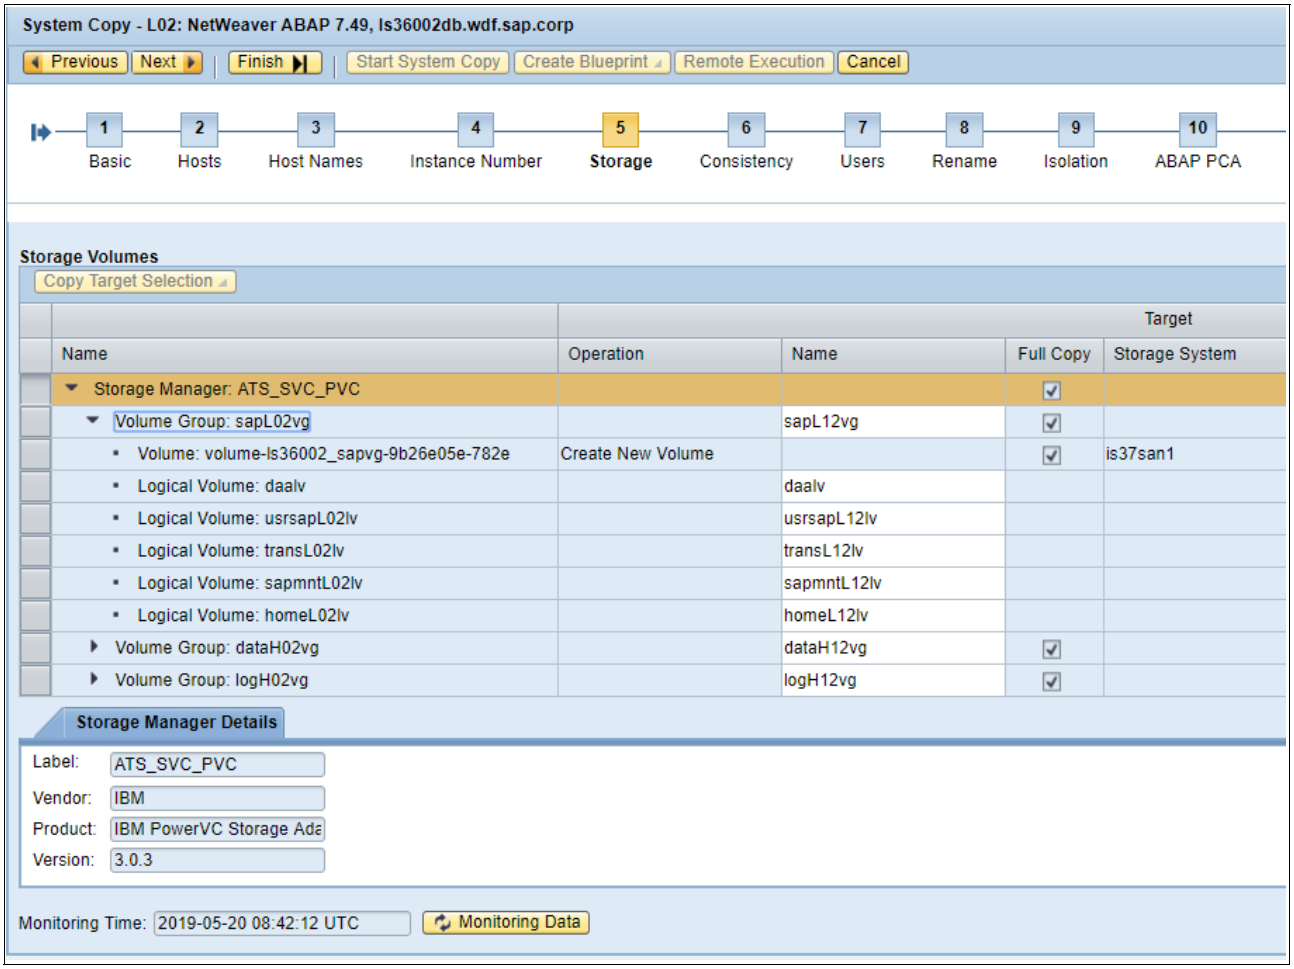

Figure 2-31 SAP System Copy roadmap: Checking the storage definitions

7. Specify the database consistency for the SAP source system (online or stop the SAP source system first) in case it is possible to define run time when the snapshot is triggered (Figure 2-32).

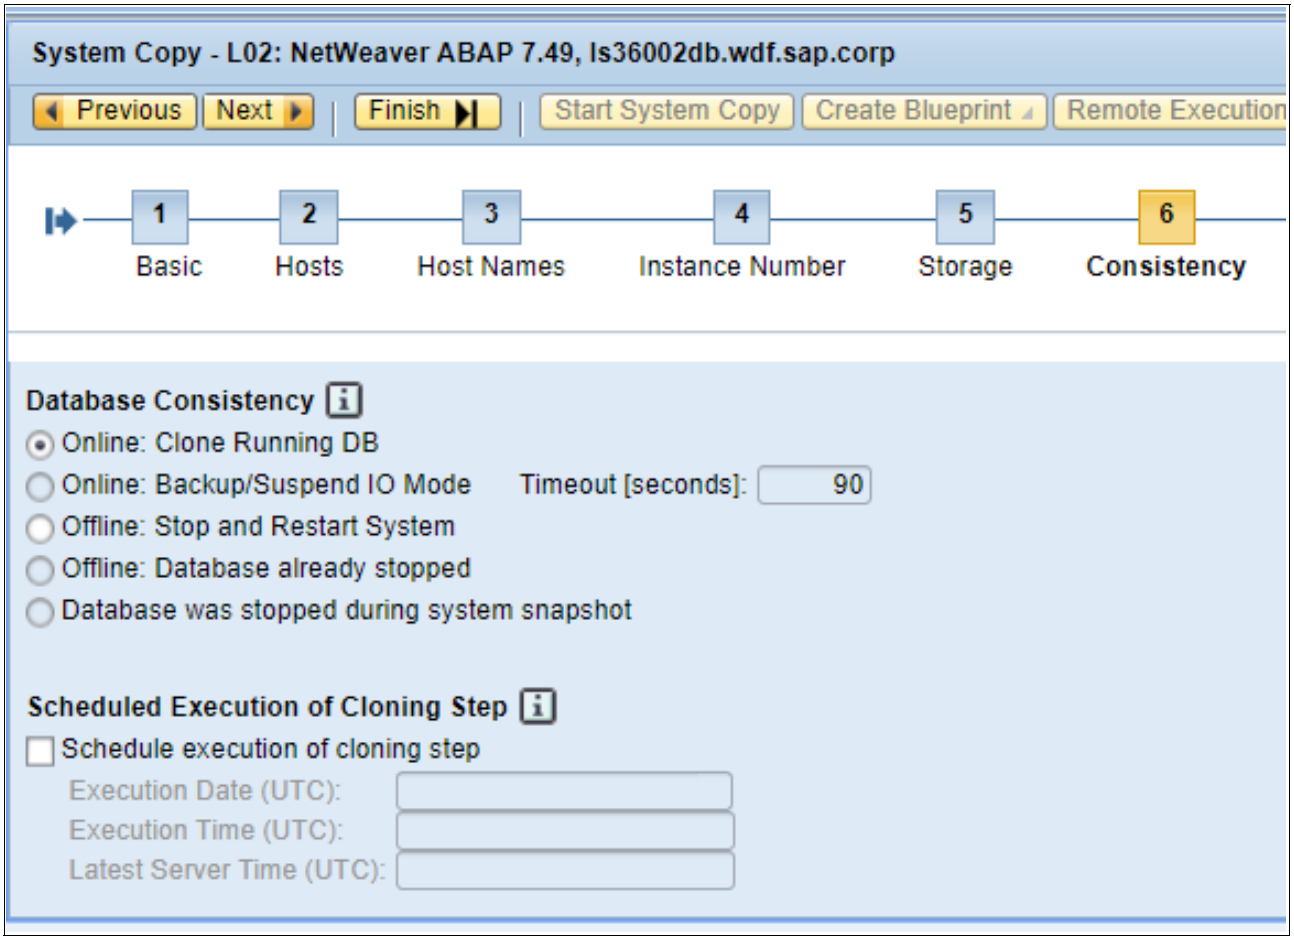

Figure 2-32 SAP System Copy roadmap: Consistency

8. Define the OS users that are required for the SAP target system. Here, the users are created by LDAP. Use the configuration of users and groups (Figure 2-33).

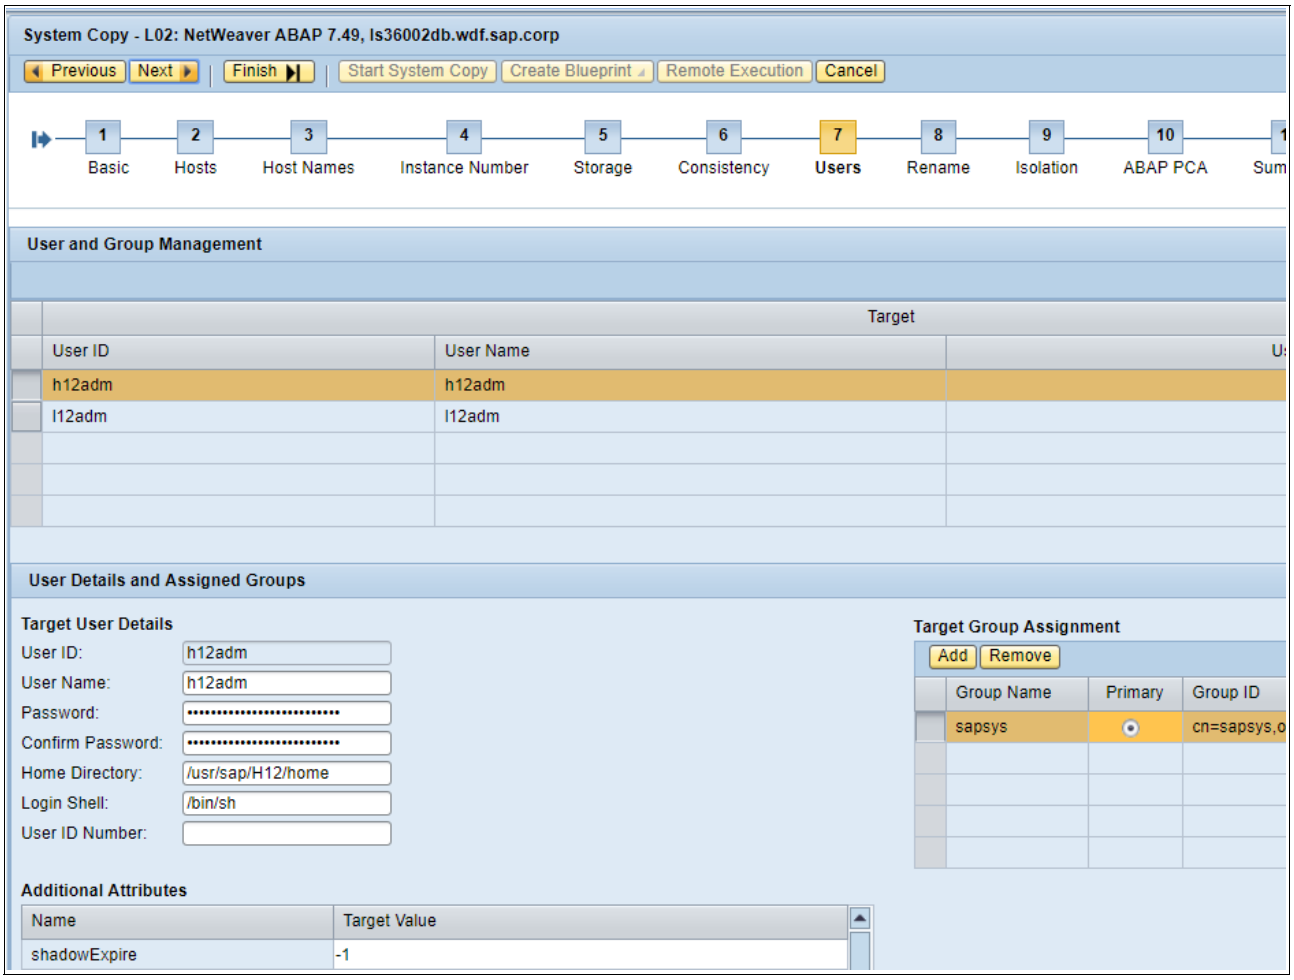

Figure 2-33 SAP System Copy roadmap: OS users and groups

9. View the SWPM settings. SWPM is started during the copy process and renames the SAP system to a new SID (Figure 2-34).

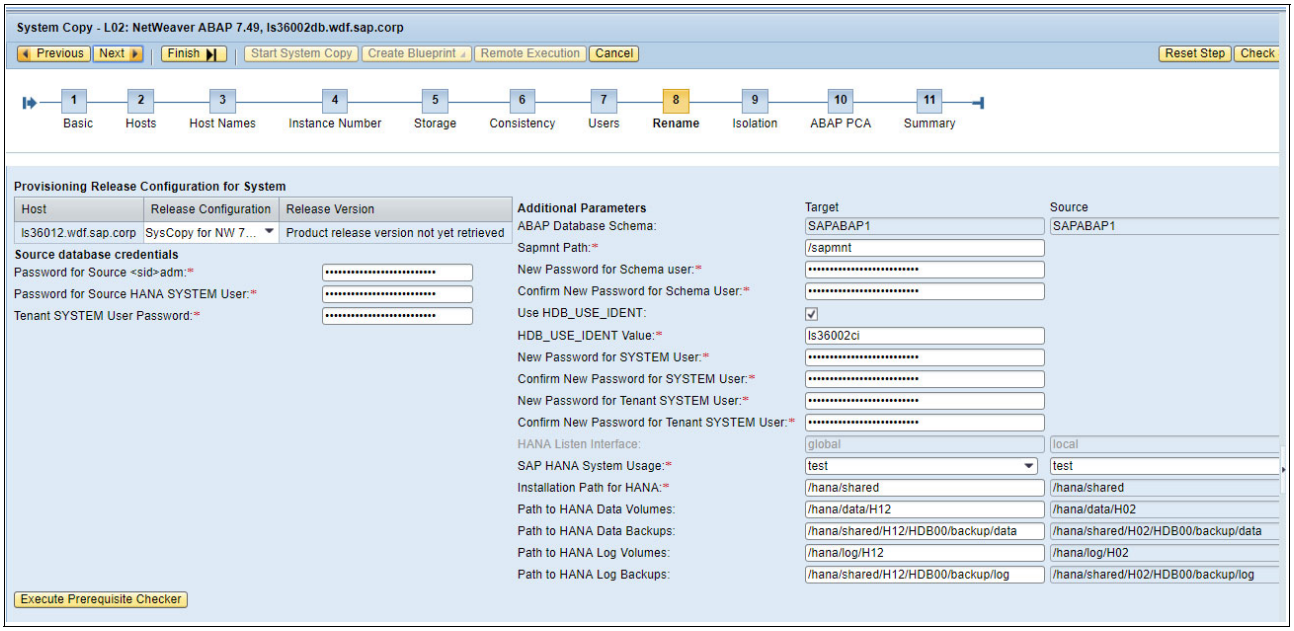

Figure 2-34 SAP System Copy roadmap: System rename settings

During the copy process, the target SAP system must be isolated to ensure SAP landscape consistency. For the isolation, SAP LaMa uses `iptables` in the SAP target host.

10. Define allowed outgoing connections, which are not blocked but remain open during the isolation process (for example, the RMC port that is used for communicating with the HMC). See Figure 2-35.

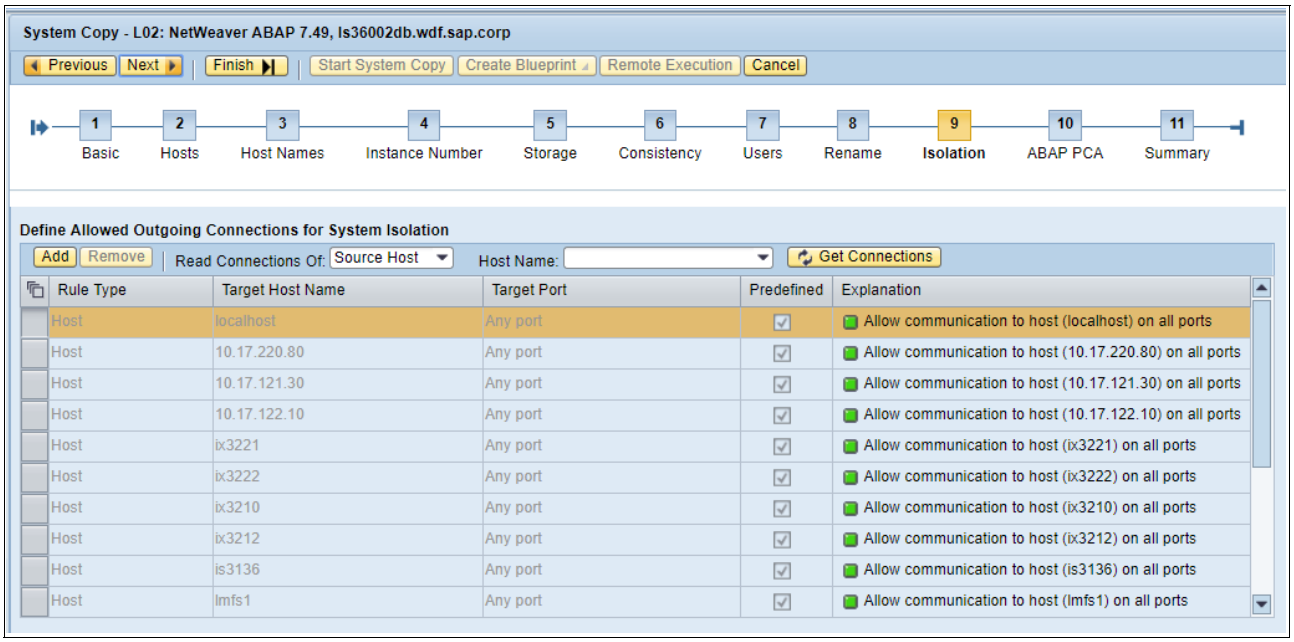

Figure 2-35 SAP System Copy roadmap: Fencing

To ensure SAP landscape consistency, the copied SAP system must be *clean*. To achieve this state, SAP SWPM must do certain tasks on the OS level (configured in step 8), and SAP PCA must do certain tasks in the SAP system itself.

11. Specify a PCA task list for the various SAP clients (Figure 2-36).

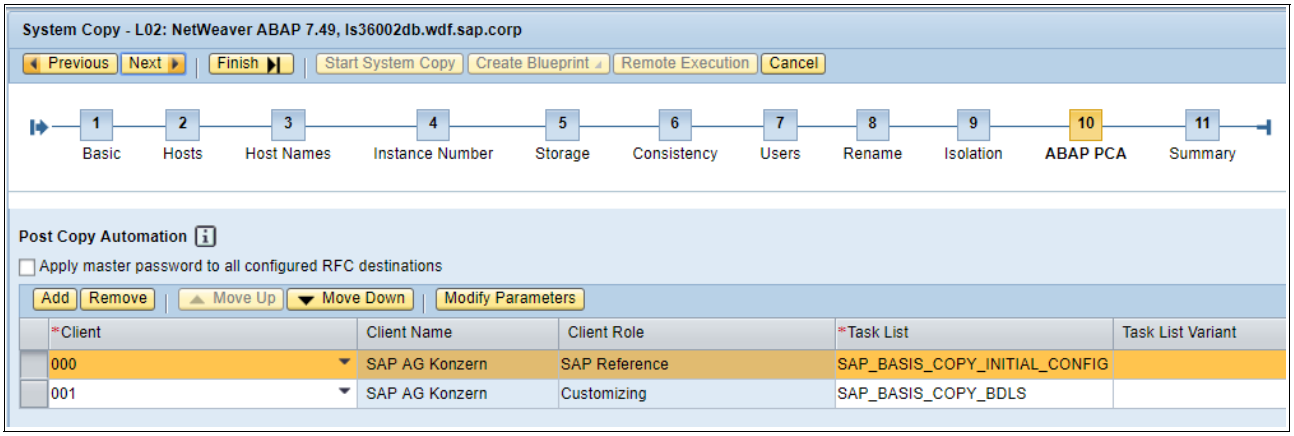

Figure 2-36 SAP System Copy roadmap: PCA settings

12. Either start the SAP system copy directly, create a blueprint and start the configured process later, or obtain a URL to trigger the process remotely (Figure 2-37).

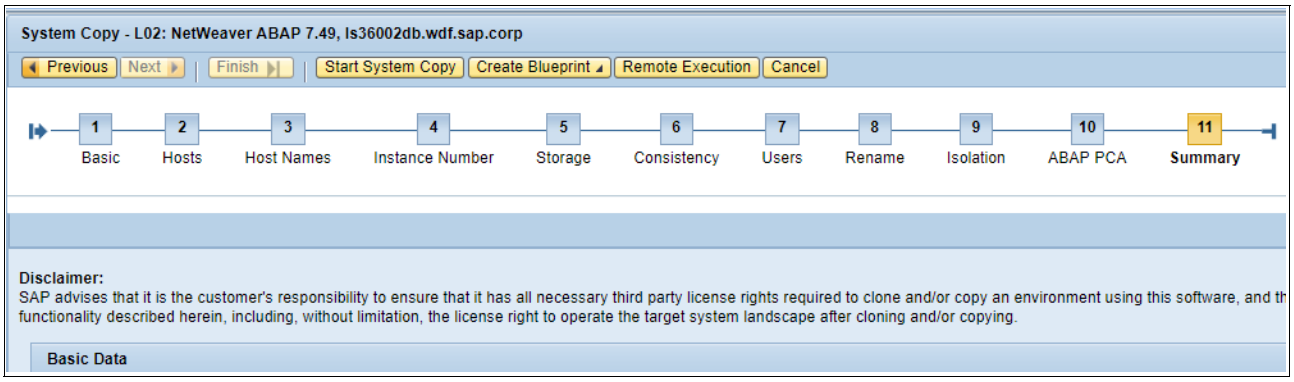

Figure 2-37 SAP System Copy roadmap: Summary

13. After starting the SAP system, select **Monitoring** → **Activities** to monitor the progress. Check the process status of each single step of this activity. You can click the step name to provide detailed logging information (see Figure 2-38 and Figure 2-39 on page 43).

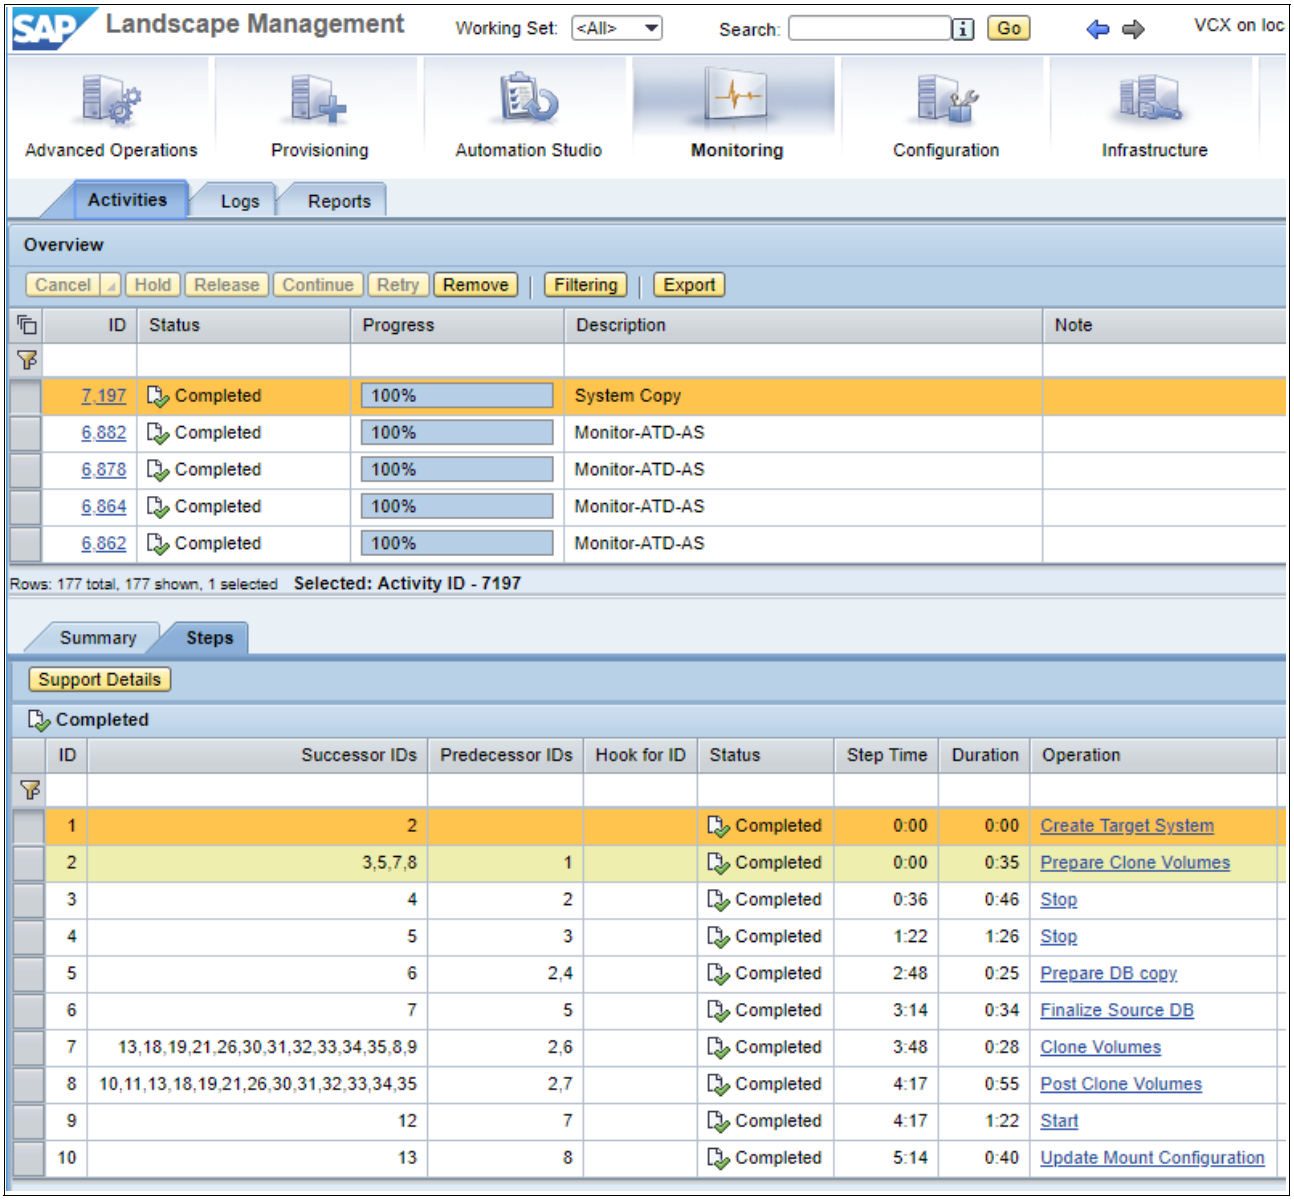

Figure 2-38 Monitoring view: Steps (1 of 2)

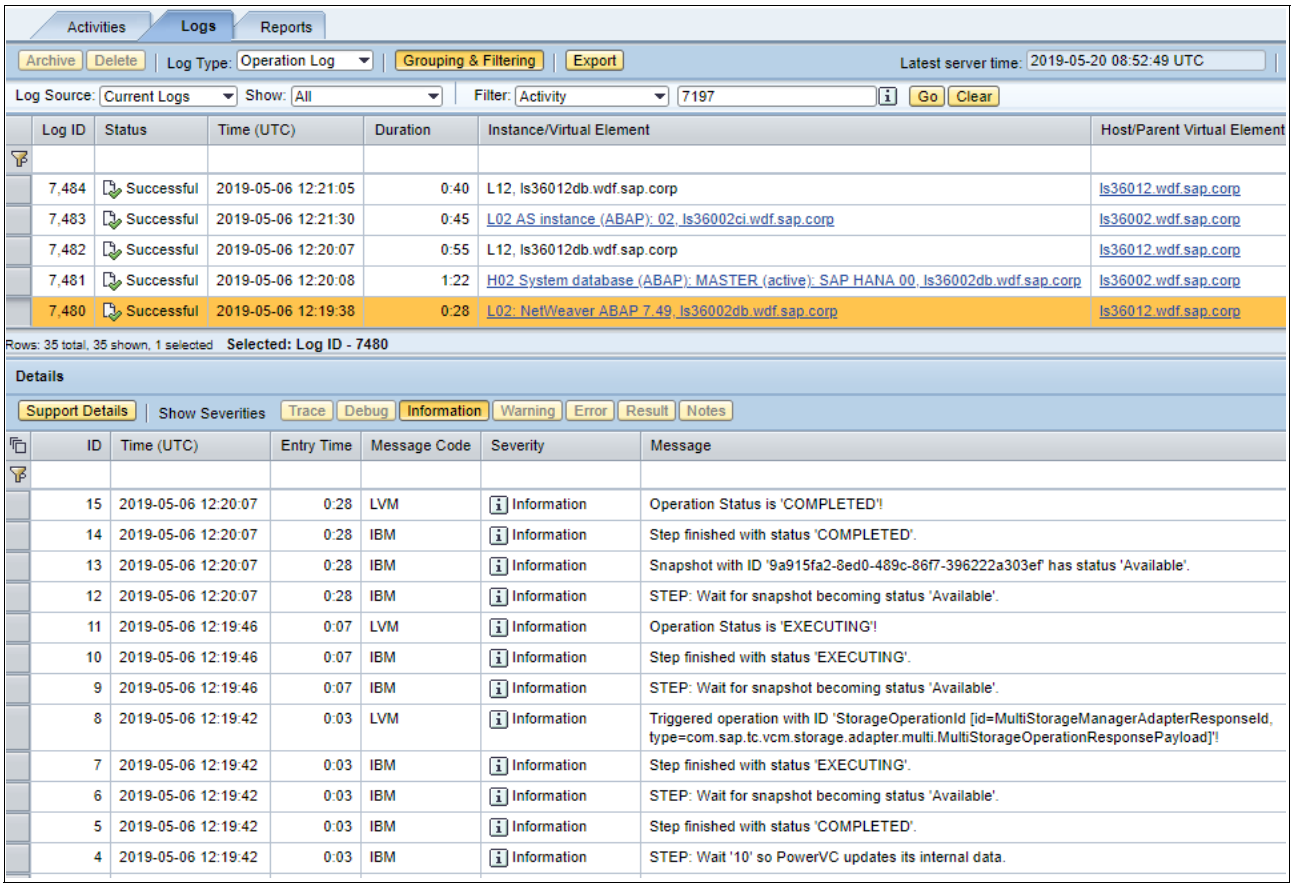

Figure 2-39 Monitoring view: Steps (2 of 2)

2.5 Outlook

SAP and IBM continuously develop more features for this solution, such as container support, HANA, and operational enhancements to reduce planned downtime. This section describes more use cases of SAP LaMa on IBM Power Systems servers with or without IBM System Storage.

2.5.1 SAP containerized applications

SAP is committed to the strategic trend of containerization. Containers are becoming a fundamental building block for modern application development. For example, SAP HANA Express Edition, which is a streamlined version of the SAP HANA platform, is available as a Docker container image for Linux and Intel processor-based platforms. You can use SAP HANA Express Edition to build and deploy modern applications that use up to 32 GB of memory so that developers can perform application development on their personal computers or in the cloud.

However, SAP NetWeaver Systems and SAP HANA databases usually are installed in bare-metal servers or VMs. These VMs remain in the data centers for many years, and must be managed and maintained by IT administrators.

One IBM Cloud key differentiator is that IBM Cloud uses Kubernetes to orchestrate container deployments, but IBM Cloud can also provision VMs for multicloud environments by using IBM Cloud Automation Manager. Cloud Automation Manager is part of IBM Cloud and can manage various cloud endpoints, including IBM PowerVC.

2.5.2 SAP HANA management

SAP and IBM introduced support for HANA scale-up copy, clone, and refresh functions and further basic management in 2018. Since then, the integration of SAP HANA has progressed and downtime-minimizing features have been general available in April 2019.

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 paper.

IBM Redbooks

The following IBM Redbooks publication provides more information about the topics in this document. This publication might be available in softcopy only.

- *IBM PowerVC Version 1.3.2 Introduction and Configuration*, SG24-8199

You can search for, view, download, or order this document and other Redbooks, Redpapers, web docs, drafts, and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM PowerVC documentation

  https://developer.ibm.com/powervc/
- SAP Landscape Management (SAP LaMa) 3.0 Enterprise Edition

  https://help.sap.com/viewer/lama_help
- SAP Support Software download portal

  https://support.sap.com/swdc

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

© Copyright IBM Corp. 2019. All rights reserved.

REDP-5568-00

ISBN 0738458228

Printed in U.S.A.

Get connected