IBM® Storage

IBM Storage Solutions for SAP Applications
Version 1.5

IBM Storage Team

IBM

Contents

About this document

This IBM® Redpaper™ publication is intended as an architecture and configuration guide to set up the IBM System Storage™ for the SAP HANA tailored data center integration (SAP HANA TDI) within a storage area network (SAN) environment. SAP HANA TDI allows the SAP customer to attach external storage to the SAP HANA server.

The paper also describes the setup and configuration of SAP Landscape Management for SAP HANA systems on IBM infrastructure components: IBM Power Systems and IBM Storage based on IBM Spectrum® Virtualize.

This document is written for IT technical specialists and architects with advanced skill levels on SUSE Linux Enterprise Server or Red Hat Enterprise Linux (RHEL) and IBM System Storage.

This document provides the necessary information to select, verify, and connect IBM System Storage to the SAP HANA server through a Fibre Channel-based SAN. The recommendations in this Blueprint apply to single-node and scale-out configurations, and Intel and IBM Power based SAP HANA systems.

For more information about the storage systems that are certified for SAP HANA production, see the Certified and Supported SAP HANA Hardware Directory.

Scope

This Blueprint provides a solutions architecture and the related storage endpoint capabilities that interact with the following software and hardware components:

- IBM Spectrum Virtualize
- IBM SAN Volume Controller
- IBM FlashSystem® Family
- IBM Spectrum Protect
- IBM Spectrum Protect for Enterprise Resource Planning
- IBM Geographically Dispersed Resiliency for IBM Power Systems
- Linux operating system (SLES or RHEL)
- SAP Landscape Management
- IBM Power Systems
- IBM Power Virtualization Center (PowerVC)

This technical report does not:

- Provide performance analysis from a user perspective
- Replace any official manuals and documents that are issued by IBM
- Explain the installation and configuration of SAP HANA or SAP S/4HANA

1

Prerequisites

This technical paper refers to SAP documentation that is available for download from the SAP Help Portal.

It is assumed that you are familiar with and have basic knowledge of the following products:

• IBM FlashSystem
• IBM Spectrum Protect
• IBM Spectrum Protect for Enterprise Resource Planning
• Linux operating system (SLES or RHEL)
• SANs

Foundation: Data placement and data management

This section covers data placement and data management.

IBM storage architecture for SAP HANA TDI Environment

IBM System Storage must be connected through a SAN to the SAP HANA server. Figure 1 shows the relationship between Linux file system, LINUX LV, VG, and PV, and the Virtual Volumes (VD, VDISK) and Managed Disks (MD, MDISK, RAIDs/Arrays) of an IBM FlashSystem. The standard layout for an SAP HANA database includes the following components:

• One or more VDs are placed in the storage pool
• One or more LVs are in one volume group (VG)
• One VG is placed on one or more physical volumes (PVs, a Linux term)
• One PV is one virtual disk (VD, an IBM term)
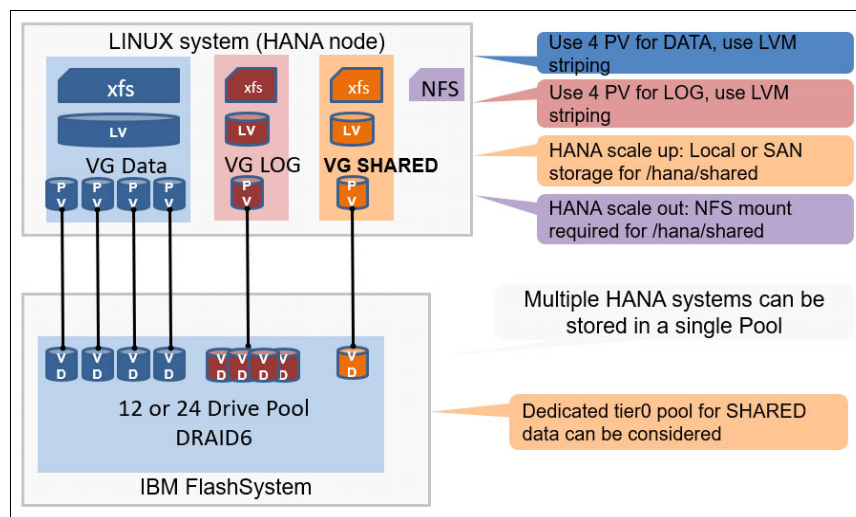• One or more VDs are placed in the storage pool

Figure 1   SAP HANA disk layout

IBM recommends the use of IBM SAN Volume Controller as a virtualization layer to enable centralized management of the storage systems and easy high-availability and disaster-recovery functionality, and IBM FlashCopy® backup and tiered-storage pooling.

Figure 2 shows the high-availability storage architecture for SAP, with IBM SAN Volume Controller configured as an enhanced stretched cluster.
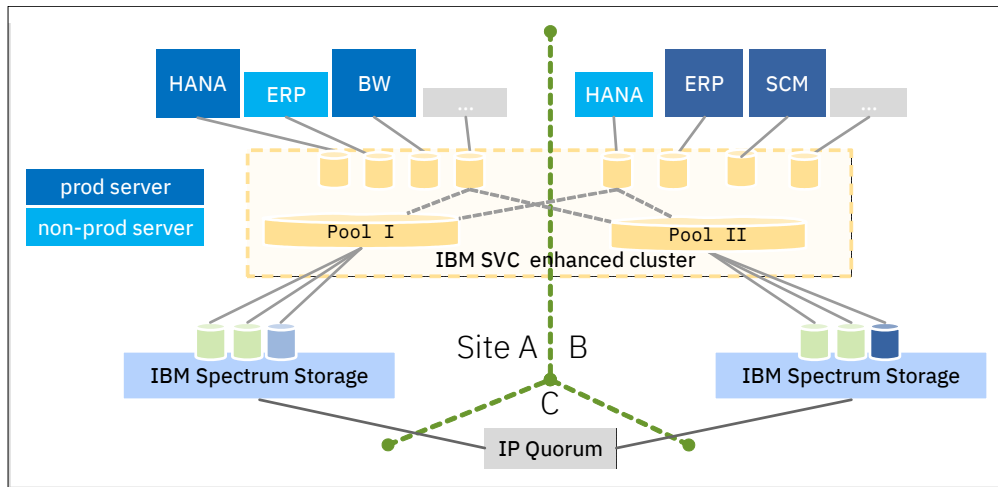

Figure 2   IBM Spectrum Virtualize architecture landscape diagram

All data, which is stored on IBM SAN Volume Controller as virtual volumes, is mirrored through IBM SAN Volume Controller virtual disk mirroring onto the physical storage system on sites A and B. These storage systems are a mix of IBM FlashSystem hardware. They are configured depending on the workload, based on the total number of SAP instances and the size of each SAP instance.

To set up a stretched storage cluster, a tiebreaker (quorum) mechanism is required. A requirement met in this case with an IP quorum.

Implementation requirements

Identify the appropriate size for your IBM Spectrum Virtualize environment, and then review requirements for hardware and software based on your SAP HANA landscape.

Use Table 1 to determine the capacity of the IBM Storage systems that you need, based on the number of SAP instances you manage. Both the total number of SAP instances and the database size in terabytes (TB) are shown in Table 1.

Table 1   Storage sizing for SAP HANA

Number of SAP HANA instances (max)	Size of each SAP HANA instance on disk (max)	IBM Storage system sizing
6 Prod + 30 Dev/Test	0.5 TB	Small
12 Prod + 60 Dev/Test	1 TB	Medium
12 Prod + 40 Dev/Test	2 TB	Large
20 Prod + 48 Dev/Test	4 TB	Extra Large
10 Prod + 45 Dev/Test	8 TB	Extra Large

Table 1 on page 3 guides you through your storage sizing. If you need more precise sizing, contact your IBM sales representative or IBM Business Partner for a sizing with the IBM myFlash4HANA tool, which is available to IBM sales team members and IBM Business Partners only.

Hardware requirements

This section covers the hardware requirements for small, medium, large, and extra large configurations. Table 2 on page 4, Table 3 on page 6, Table 4 on page 7, and Table 5 on page 8 contain the descriptions and quantities for IBM components that are used in the storage-configuration Blueprints.

The configurations are recommendations to meet the SAP HANA TDI performance and capacity specifications that are determined by SAP Quick Sizer.

The hardware specifications that are provided are current at the time of this writing. Part substitutions might be required, depending on hardware availability. Be cautious if you plan to substitute a smaller quantity of higher-capacity drives, particularly for the database. A smaller quantity of drives might not provide comparable performance.

The tables contain abbreviated part lists, which include only major components. Work with your hardware supplier to ensure that your configuration is complete.

The tables include recommended configurations for IBM FlashSystem and in addition the configuration of an enhanced cluster setup with IBM SAN Volume Controller and IBM FlashSystem, which is a high availability option for SAP HANA landscapes.

Also, IBM SAN Volume Controller is always certified for SAP HANA in combination with any certified backend storage system.

You must provide or acquire hardware that is based on suitable scale requirements. You can choose components that are equivalent to or better than those listed.

> **Note:** Part numbers for IBM warranties are not included here. There are standard offerings for 1-year and 3-year warranties. The warranties that are available for purchase depend on your configuration requirements.

A SAN switch must be included, with enough ports to match those of the physical SAP HANA server.

A SAN environment likely exists and might need to be extended.

Hardware requirements for small-sized landscapes

Entry enterprise IBM FlashSystem solutions are designed to provide simple, high-performance management of block-based storage. For this reason, IBM FlashSystem is suited for use by SAP HANA. The configuration is based on Table 2.

Table 2   Hardware requirements for small-sized landscapes

Blueprint component	Description	Quantity	Part number
IBM SAN Volume Controller	IBM SAN Volume Controller Storage Engine **Required for optional HA solution**		2147-SV2

	IBM SAN Volume Controller Initial Storage Engine		3001
	Additional IBM SAN Volume Controller Storage Engine		3002
	4-port, 32 Gb Fibre Channel card, per engine	1	AH1D
	IBM Spectrum Virtualize software for IBM SAN Volume Controller V8		5641-VC8
	IBM Spectrum Virtualize software for IBM SAN Volume Controller 3Y		5641-VCC
	IBM Enterprise Class Support for IBM SAN Volume Controller		5773-ECS
IBM Storage	IBM FlashSystem 5030 Controller		2072-3H4
	16 Gb Fibre Channel adapter pair		ALBB
	Controller cache upgrade		ALGA
	7.68 TB, 2.5-inch flash drive (solid-state drive [SSD])	12	AL82
IP quorum	Linux system, native or as a virtual machine (VM)		RHEL > V7.x SLES > V12.x

Hardware requirements for medium-sized landscapes

Table 3 lists the hardware requirements for medium-sized landscapes.

Table 3   Hardware requirements for medium-sized landscapes

Blueprint component	Description	Quantity	Part number
IBM SAN Volume Controller	IBM SAN Volume Controller Storage Engine **Required for optional HA solution**		2147-SV2
	Initial IBM SAN Volume Controller Storage Engine		3001
	Additional IBM SAN Volume Controller Storage Engine		3002
	4-port, 32 Gb Fibre Channel card, per engine	2	AH1D
	IBM Spectrum Virtualize software for IBM SAN Volume Controller V8		5641-VC8
	IBM Spectrum Virtualize software for IBM SAN Volume Controller 3Y		5641-VCC
	IBM Enterprise Class Support for IBM SAN Volume Controller		5773-ECS
	Cache upgrade 256 GB		ACH1
IBM Spectrum Storage™	IBM FlashSystem 5100 Controller		2077-4H4
	32 Gb Fibre Channel adapter pair		ACB9
	Controller cache upgrade		ACGF
	9.2 TB, 2.5-inch flash drive (FCM)	12	AES6
IP quorum	Linux system native, or as a VM		RHEL > V7.x SLES > V12.x

Hardware requirements for large-sized landscapes

Table 4 lists the hardware requirements for large-sized landscapes.

Table 4   Hardware requirements for large-sized landscapes

Blueprint component	Description	Quantity	Part number
IBM SAN Volume Controller	IBM SAN Volume Controller Storage Engine **Required for optional HA solution**		2147-SV2
	Initial IBM SAN Volume Controller Storage Engine		3001
	Additional IBM SAN Volume Controller Storage Engine		3002
	4-port, 32 Gb Fibre Channel card, per engine	2	AH1D
	IBM Spectrum Virtualize software for IBM SAN Volume Controller V8		5641-VC8
	IBM Spectrum Virtualize software for IBM SAN Volume Controller 3Y		5641-VCC
	IBM Enterprise Class Support for IBM SAN Volume Controller		5773-ECS
	Cache upgrade 384 GB		ACH2
IBM Spectrum Storage	IBM FlashSystem 7200 Controller		2076-824
	32 Gb Fibre Channel adapter pair	2	ADB4
	Controller cache upgrade		ACGB
	9.6 TB, 2.5-inch flash drive (FCM)	24	ADS6
IP quorum	Linux system, native or as a VM		RHEL > V7.x SLES > V12.x

Hardware requirements for extra large-sized landscapes

For even larger SAP HANA deployments, IBM FlashSystem solutions are a good fit with respect to maximum performance and density (TB per rack unit). Also, the configuration for IBM SAN Volume Controller includes spare nodes for increased availability.

Table 5 lists the hardware requirements for extra large-sized landscapes.

Table 5   Hardware requirements for extra large-sized landscapes

Blueprint component	Description	Quantity	Part number
IBM SAN Volume Controller	IBM SAN Volume Controller Storage Engine **Required for optional HA solution**		2147-SV2
	Initial IBM SAN Volume Controller Storage Engine		3001
	Additional IBM SAN Volume Controller Storage Engine		3002
	Standby IBM SAN Volume Controller Storage Engine		3003
	4-port, 32 Gb Fibre Channel card, per engine	3	AH1D
	IBM Spectrum Virtualize software for IBM SAN Volume Controller V8		5641-VC8
	IBM Spectrum Virtualize software for IBM SAN Volume Controller 3Y		5641-VCC
	IBM Enterprise Class Support for IBM SAN Volume Controller		5773-ECS
	Cache upgrade 384 GB		ACH2
IBM Spectrum Storage	IBM FlashSystem 9200 SFF controller		9846-AG8
	32 Fibre Channel adapter pair	3	AHB4
	Cache upgrade	1	ACGB
	19.2 TB NVM Express (NVMe) IBM FlashCore modules	24	AHS7
IP quorum	Linux system, native or as a VM		RHEL > V7.1 SLES > V12.x

Physical installation and setup of IBM Storage Systems

This section describes how to set up and configure the storage systems that are outlined. The physical setup of the appropriate IBM storage systems is described in the following documents:

- IBM SAN Volume Controller library and related publications

- IBM System Storage SAN Volume Controller Models 2145 and 2147 Hardware Installation Guides:

   – IBM FlashSystem 5000 documentation
   – IBM FlashSystem 7200 documentation
   – IBM FlashSystem 9200 documentation

Configuring the storage landscape

This section describes configuring the IBM storage systems and the Linux operating system in preparation for SAP HANA.

Figure 3 shows the following logical storage relationship:

- One file system (here, XFS) is built on one logical volume (LV)
- One or more LVs are in one volume group (VG)
- One VG is placed on one or more physical volumes (PVs, a Linux term)
- One PV is one virtual disk (VD, an IBM term)
- One or more VDs are placed in the storage pool
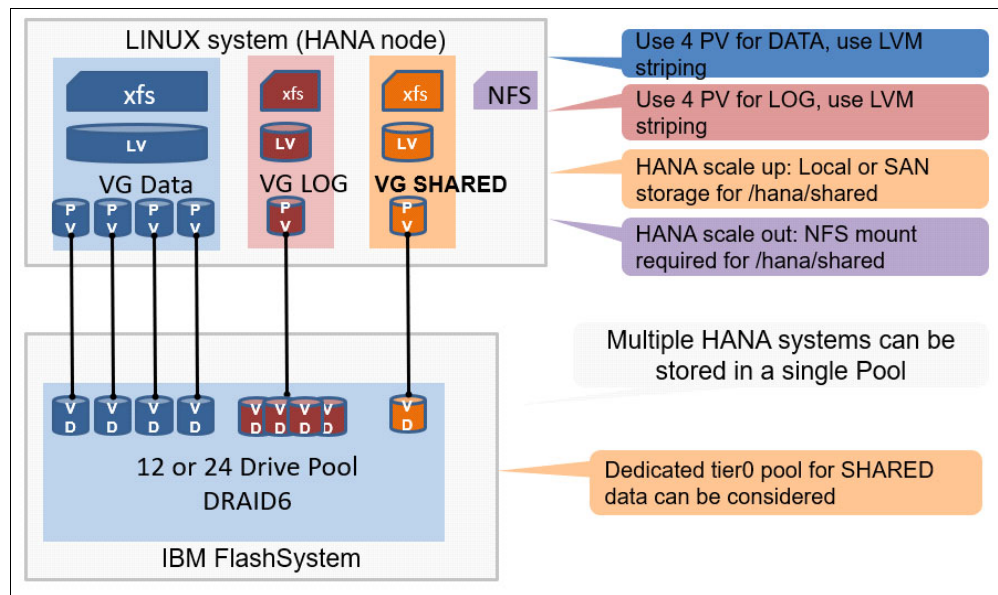
Figure 3   Logical storage configuration

The following central schema is used to set up the infrastructure for SAP HANA TDI:

1. Verify that you met all SAP-provided infrastructure requirements.

2. Configure the IBM storage systems.

3. Create VMs or logical partitions (LPARs) for the HANA systems.

4. Create SAN zones between storage systems and VMs or LPARS.

5. On the storage systems, map the storage volumes by logical unit number (LUN) to the target SAP HANA nodes.

6. Install the Linux operating system on VMs.

7. Perform Linux setup and tuning, including multipath and input/output (I/O) scheduler configuration.

8. On each Linux host, create the file systems for SAP HANA data, log, and shared.

9. Set up Network File System (NFS) services for the SAP HANA shared directory if you plan to scale out your HANA deployment by using a solution, such as SAP Business Warehouse (BW).

Applying SAP requirements

Consider the following primary SAP Notes for an IBM landscape:

• 2378874: Install SAP Solutions on Linux on IBM Power Systems (little endian)

• 2055470: SAP HANA on IBM Power Systems and IBM System Storage Planning and Installation Specifics: Central Note

SAN zoning between IBM SAN Volume Controller and IBM FlashSystem

The SAN administrator must create four SAN zones. For more information, see the section "Edge-core-edge topology" of Chapter 1, Storage Area Network, of *IBM System Storage SAN Volume Controller, and IBM FlashSystem 7200 Best Practices and Performance Guidelines*.

Configuring IBM FlashSystem

All configuration can be completed by using the graphical user interface (GUI). This GUI eases the setup process compared to the use of the command-line interface (CLI). The IBM GUI guides you through the initial setup, including configuration of a distributed array and distributed RAID (DRAID).

Complete the following steps to configure the IBM FlashSystem to be used by IBM SAN Volume Controller:

1. Create a single pool, called `Pool1`.

2. Include all SSDs in one distributed array (by using DRAID 6) and map that array to `Pool1`; turn on IBM Easy Tier®; turn off compression.

3. For small or medium configurations, create 4 VDs of equal size, by using 80% of the available `Pool1` capacity.

4. For large configurations, create 8 VDs of equal size, by using 80% of the available `Pool1` capacity.

5. Map the VDs to IBM SAN Volume Controller.

Configuring IBM SAN Volume Controller

To set up IBM SAN Volume Controller to provide an enhanced stretched cluster, follow the procedures up to Chapter 4 in *IBM Spectrum Virtualize and SAN Volume Controller Enhanced Stretched Cluster with VMware*.

Create two pools, site "A" and site "B," with the MD of IBM FlashSystem of "A" or "B."

Installing Linux

Install the Linux operating system according to the SAP Product Availability Matrix (PAM) for the chosen SAP HANA version. For more information, see SAP Note 1944799.

Creation of VMs (or LPARs)

Create VMs by using IBM Power Virtualization Center (IBM PowerVC). This setup is not part of this document.

During this setup process, IBM PowerVC also configures the appropriate SAN zones.

Creating volumes for SAP HANA VM

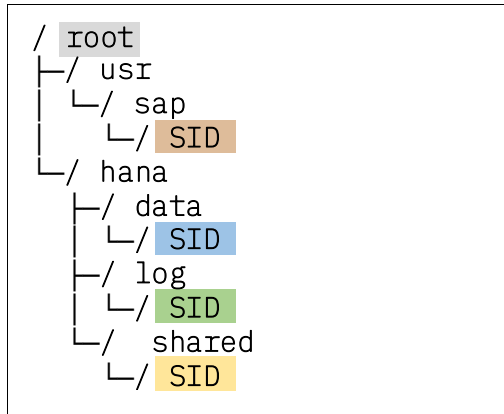Figure 4 shows the required file system layout. Each colored box indicates one mount point with one LV and one VG.

Figure 4   File system layout for SAP HANA VM

Consider the following sizing guidelines:

- Use one 100 GB physical disk (PD) for root VG

- Use one 50 GB PD for `/usr/sap`/SID VG

- Use one 100 GB PD for `/hana/shared/SID` VG

- Use four PDs for `/hana/data/SID` VGs, sized to anticipate growth, according to HANA sizing guidelines

- Use four PDs (maximum size 500 GB) for `/hana/log/SID` VGs, sized according to HANA sizing guidelines

Creating the file systems for SAP HANA

Complete the following steps to create the file system for SAP HANA:

1. Use four PVs for `/hana/data/log` and four for `/hana/log/SID`; use one for each `/usr/sap/SID` and one for `/hana/shared/SID`. `ANA (ana)` is used as SID in this example.

2. Create a VG with an alignment size of 1 MB:

   ```
   vgcreate -s 1M --dataalignment 1M analogvg /dev/mapper/WWID_1 dev/mapper/WWID_2
   ...
   ```

3. Create an LV with a stripe size of 256 KB:

   ```
   lvcreate -i 4 -I 256K -l 100%VG -n analoglv analogvg
   ```

4. Create the file system:

   ```
   mkfs.xfs -b size=4096 -s size=4096 /dev/analogvg/analoglv
   ```

5. Add file systems into `/etc/fstab` (see Figure 5).

```
/dev/anasharedvg/anasharedlv /hana/shared/ANA xfs defaults   0 0
/dev/anadatavg/anadatalv     /hana/data/ANA   xfs defaults   0 0
/dev/analogvg/analoglv       /hana/log/ANA    xfs defaults   0 0
/dev/anausrsapvg/anausrsaplv /usr/sap/ANA     xfs defaults   0 0
```

Figure 5   File systems

6. Mount the file system:

    `mount /hana/log/ANA`

    `mount ...`

Required Linux tuning

Tune the Linux multipath settings and I/O scheduler selection.

Multipath

For more information about multipath configuration settings for IBM FlashSystem, see IBM Knowledge Center and search for your IBM FlashSystem.

Using user-friendly names

Optionally, you can use user-friendly names, as shown in Figure 6, which makes managing and identifying volumes easier.

```
defaults {
  user_friendly_names          yes
}

multipaths {
  multipath {
    wwid    360060480000028350131253594d303030
    alias   HANA_ANA_DATA_1
  }
  multipath {
    wwid    360060480000028350131253594d303041
    alias   HANA_ANA_LOG_1
}
```

Figure 6   User friendly naming

I/O scheduler (elevator kernel parameter)

The default Completely Fair Queuing (CFQ) scheduler delivers low I/O performance for multiple parallel streaming read processes. The NOOP scheduler is recommended for setups with devices that perform I/O scheduling, such as intelligent storage devices, or in multi-path environments. The Deadline scheduler can provide a superior throughput over the CFQ I/O scheduler in cases where several threads read and write, and fairness is not an issue (for example, for several parallel readers from a SAN).

We recommend the use of the NOOP scheduler for SAP HANA.

For more information, see *SUSE System Analysis and Tuning Guide*.

HANA configuration for IBM SAN Volume Controller

The following `fileio` parameters must be set for IBM storage systems by using the SAP HANA tool `hdbparam`:

```
async_write_submit_active :  on
async_write_submit_blocks : all
async_read_submit         : on
```

For more information about the SAP HANA tool, see this web page (log in required).

Data protection: Backup and recovery

SAP HANA backups are required to protect against disk failures and enable earlier database resets. To perform online backups of SAP HANA, you can use the following IBM software tools:

- IBM Spectrum Protect for Enterprise Resource Planning: Data Protection for SAP HANA
- IBM FlashCopy backup for SAP HANA (FCTool)
- IBM Spectrum Copy Data Management (CDM)

Figure 7 shows the solution architecture of the virtualized SAP HANA environment in combination with Data Protection for SAP HANA, the FlashCopy backup solution, and IBM Spectrum Virtualize, which serves as the virtualization layer for storage.
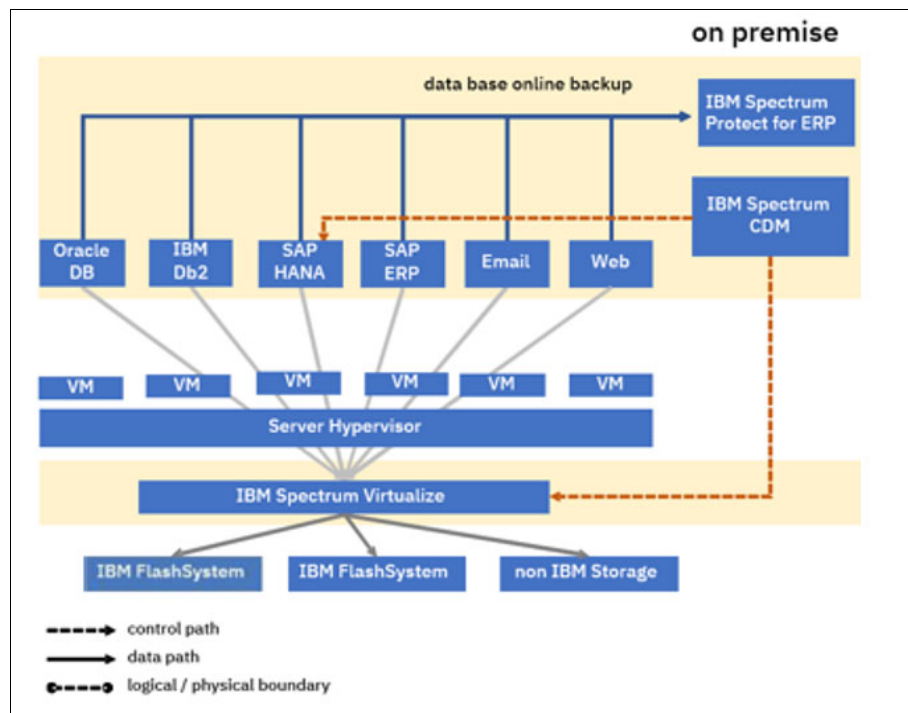
Figure 7   Solution architecture

If you experience a total loss of the SAP HANA appliance, use the SAP HANA recovery procedure to recover the system. For more information about hardware recovery, see the suitable SAP documentation.

The following sections show where to find detailed information about the Data Protection and FlashCopy backup solutions.

SAP HANA backup, IBM Spectrum Protect for Enterprise Resource Planning

IBM Spectrum Protect, formerly known as IBM Tivoli® Storage Manager for Enterprise Resource Planning, includes the package Data Protection for SAP HANA.

For more information about how to install, configure, and run this integrated solution, see the *IBM Spectrum Protect for Enterprise Resource Planning* documentation.

IBM FlashCopy backup for SAP HANA

To minimize backup and restore time, IBM offers a FlashCopy solution that combines SAP HANA snapshot technology with IBM Spectrum Virtualize. For more information, see *IBM FlashCopy backup and copy solution for SAP HANA TDI, V4.1*.

FlashCopy backup of SAP HANA with IBM Spectrum Copy Data Management

To minimize the backup and restore time IBM offers a solution that combines SAP HANA SNAPSHOT technology with IBM Spectrum Virtualize FlashCopy. IBM Spectrum Copy Data Management (CDM) is able to backup SAP HANA databases using SAP HANA Snapshot and IBM FlashCopy technology.

IBM Spectrum Copy Data Management offers support for SAP HANA 2.0. Single tenant configuration can now be automatically protected using storage snapshot.

CDM supports IBM Spectrum Virtualize, which includes the IBM FlashSystem family, IBM Spectrum Accelerate, and some foreign storage systems. The supported applications and use cases of CDM are shown in Figure 8. CDM provides you with centralized management for the SAP HANA backups across data centers both local and remote.
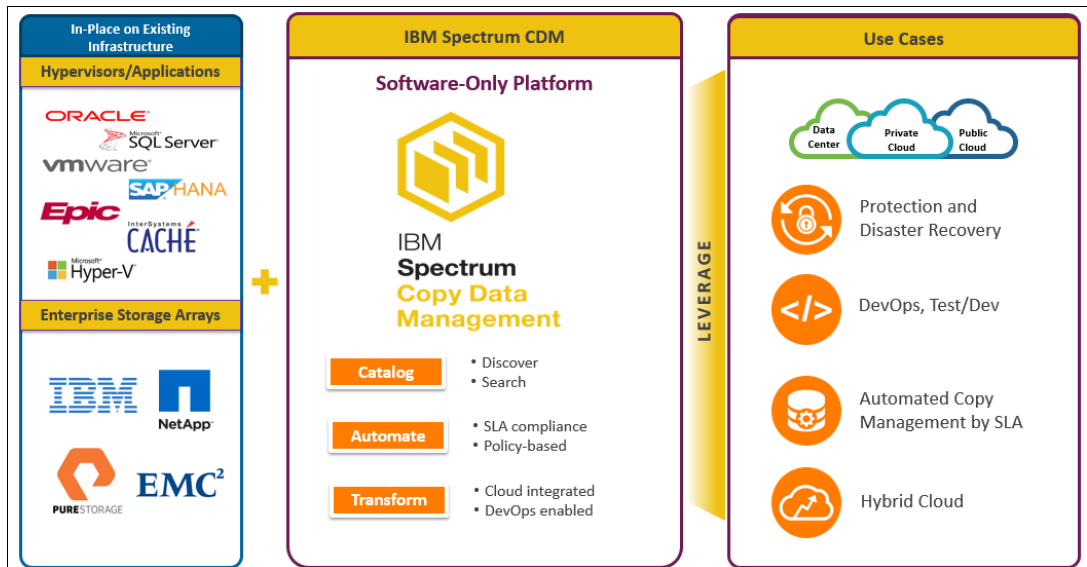
Figure 8   Spectrum copy Data Management platform

Since CDM version 2.2.11, SAP HANA transaction logs also can be backed up and restored. This feature gives users the ability to perform backups and point-in-time (PIT) restores of SAP HANA databases. For more information about CDM, see IBM Knowledge Center.

For more information about how to set up and use a combined CDM/Spectrum Protect for ERP environment to protect SAP HANA databases, see *Practical Guide: Protecting SAP HANA with IBM Spectrum Protect and IBM Spectrum Copy Data Management*, WP102813.

Data placement: SAP data aging

For many years, SAP classified internal data based on its age and access pattern. For more information, see the following SAP documentation:

- https://blogs.sap.com/2016/10/03/how-to-perform-data-aging-in-s4hana/
- http://www.sapsa.se/wp-content/uploads/2015/11/Wasserman.pdf
- https://www.slideshare.net/SAPTechnology/sap-hana-sps-09-dynamic-tiering

Dynamic tiering

Dynamic tiering is available for SAP BW on IBM DB2® and Oracle databases. Dynamic tiering distinguishes between:

- Hot: Loaded in HANA random access memory (RAM) during database startup

- Warm: Loaded into RAM on demand

- Cold: Stored outside the primary HANA database, but in a secondary anyDB database or SAP IQ by way of the nearline storage (NLS) interface

- Frozen: Stored outside any SAP database, or was archived and then deleted from the SAP database

SAP dynamic tiering retains Warm data in a HANA persistence layer and the data is brought into memory only when needed. Infrequently accessed data is retained in the persistence layer; if accessed frequently, it overwhelms the system.

Dynamic tiering can be installed on its own server, but not on its own HANA node. It is provided as an add-on option for native SAP HANA applications. At the time of this writing, SAP is evaluating dynamic tiering as an option for S/4HANA.

Data tiering optimization

Data lifecycle management strategies can be implemented with the data tiering optimization option. Data tiering optimization (DTO) is available for SAP BW on HANA or SAP BW/4HANA only and distinguishes between:

- Hot: Processed and stored on the primary HANA node

- Warm: Processed and stored on extension nodes

- Cold: Processed and stored outside the primary and extension nodes on SAP IQ (not anyDB), and accessed by way of an NLS interface

DTO improves data management by shifting data to the performance, storage, and processing tier with the cost and performance characteristics that are best suited for that data. Data tiering options from SAP provide a trusted approach for intelligent enterprises to manage data growth effectively and economically.

Data tiering is similar to conventional BW on HANA scale-out, but also provides the benefits of lower CPU/RAM requirements on the extension nodes.

Data aging

Data aging is available for S/4HANA only, and distinguishes between:

- Hot: Loaded in HANA RAM during startup
- Warm: Loaded into RAM on demand

Impact to storage

In all three cases, the hot and warm data is on a file system and stored on a storage system.

Especially with S/4HANA data aging, HANA acts as a classical database, such as IBM DB2. Data is loaded from the file system into RAM on demand, instead of entirely at start, because the database size (in the file system) is almost always much larger than the available RAM.

This IBM Blueprint can be used to set up all three SAP aging scenarios if you decide to set up data aging for your SAP instance.

Data protection: Disaster recovery

IBM Spectrum Protect is the eminent IBM enterprise backup and restore (EBR) tool. A key differentiator for IBM Spectrum Protect is that it performs a progressive incremental forever backup, which means that only a single full backup is required. All subsequent backups contain only incremental changes. This approach means that IBM Spectrum Protect solutions require less storage, compute, and network infrastructure when compared to other EBR solutions.

Through integration between the solutions, IBM Spectrum Protect and IBM Spectrum Virtualize manage and protect the huge volumes of data for big data analytics and new cognitive workloads.

Figure 9 on page 17 shows the architecture for a possible disaster-recovery scenario. In this use case, we have two on-premises data centers.
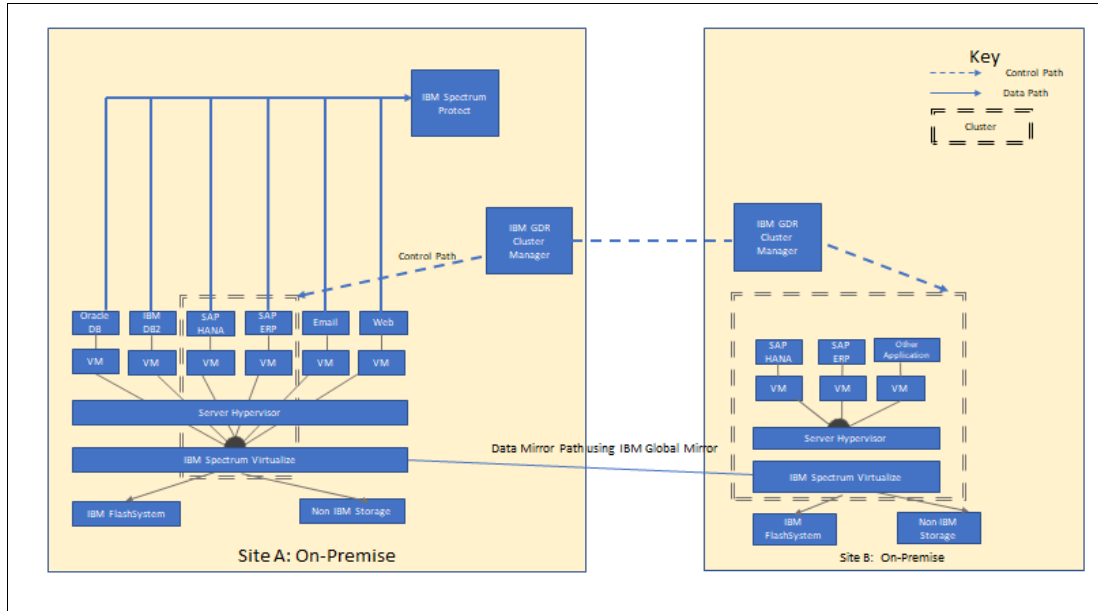
Figure 9   Possible disaster-recovery scenario

Site A is the main site. Under normal operations, IBM Spectrum Virtualize manages the storage resources that are needed by the VMs and their associated database applications, such as SAP S/4HANA. IBM Spectrum Protect is providing backup, archive, recovery, space management and, if necessary, bare-metal system recovery.

Site B is the failover site. It also is considered to be on-premises. IBM Global Mirror is mirroring the data on the IBM SAN Volume Controller cluster from Site A to Site B. IBM Geographically Dispersed Resiliency for Power Systems serves as the cluster manager that pairs the VMs across the two sites. This pairing ensures that each site has the resources available to run the host if a planned or unplanned move occurs. IBM Geographically Dispersed Resiliency for Power Systems is installed on both sites.

In case of a physical disaster on Site A, HANA can be restarted on Site B. If a logical disaster occurs, the data can be restored from the IBM Spectrum Protect instance on Site A. Some downtime occurs, but a business can successfully recover from either scenario, whether a physical or a logical disaster.

Disaster recovery with geographically dispersed resiliency

IBM Geographically Dispersed Resiliency for Power Systems is a set of software components that provide a disaster-recovery mechanism for VMs that are running on an IBM POWER® processor-based server.

The main purpose of disaster recovery is to have a defined (and possibly automated) procedure for recovery from a major business impact, such as an outage of the entire data center as the result of an earthquake, flood, or storm.

Recovery time objective (RTO) and recovery point objective (RPO) for disaster recovery are normally different from the RTO and RPO values for availability.

17

IBM Geographically Dispersed Resiliency for Power Systems offers the following benefits:

- Simplified disaster-recovery management
- Cost savings that are based on elimination of the need for hardware and software resources at the backup site
- Reduced licensing costs
- Reduced administrative costs
- Ease of deployment because no high-availability clustering is required
- The ability to test the disaster-recovery environment without affecting the production environment

IBM Geographically Dispersed Resiliency for Power Systems uses VM technology to restart VMs at a backup site if a disaster or planned system maintenance by the administrator occurs. All of this is controlled by a single management system that is separate from the production systems. Replication is performed at the storage layer, which provides storage-based mirroring with automated failover.

IBM Geographically Dispersed Resiliency for Power Systems consists of several necessary components (see Figure 9 on page 17) that work together to provide a highly available environment for systems across two sites.

Figure 10 shows the principal setup. For more information about the various components, subsystems, and tasks that are associated with IBM Geographically Dispersed Resiliency for Power Systems, see *IBM Geographically Dispersed Resiliency for Power Systems*.
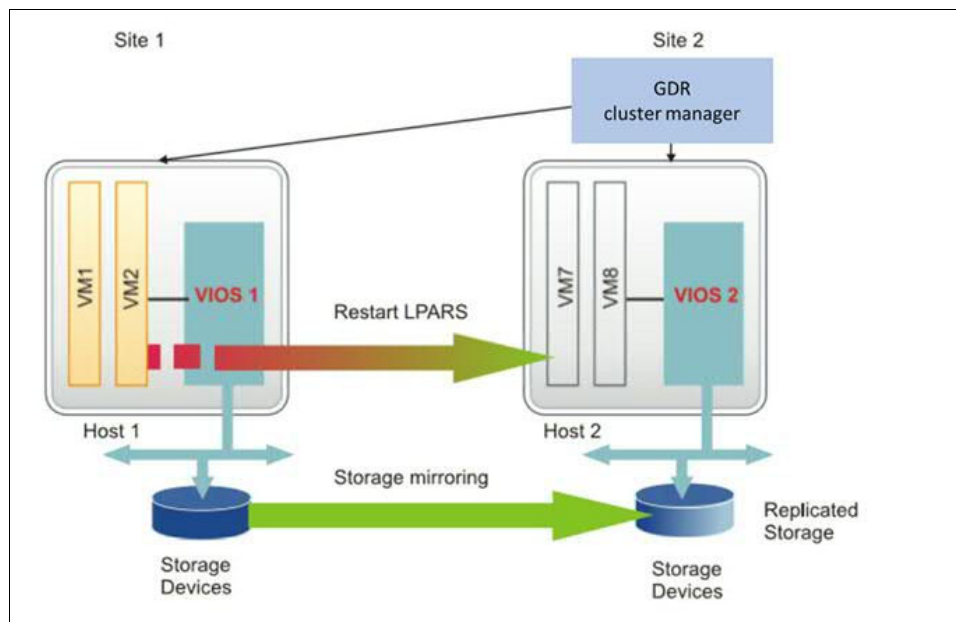
Figure 10   Components of IBM Geographically Dispersed Resiliency for Power Systems

SAP Landscape Management

This section covers SAP Landscape Management.

SAP HANA infrastructure

This section covers the technical overview and SAP Landscape view.

Technical overview and SAP Landscape overview

SAP Landscape Management (LaMa) provides functionalities that customers can use to streamline their IT operation for SAP environments, such as:

• Monitoring

• Start, stop, and relocate SAP systems (including mass operations)

• Provisioning of SAP systems, operating systems, virtual machines, and storage volumes (LUNs)

• SAP system copy, cloning, and refresh

Some of these functionalities are provided by SAP, others integrate with SAP partner solutions. The IBM integrated solutions are: IBM PowerVC, IBM Spectrum Protect Snapshot (FlashCopy Manager) (see Figure 11).
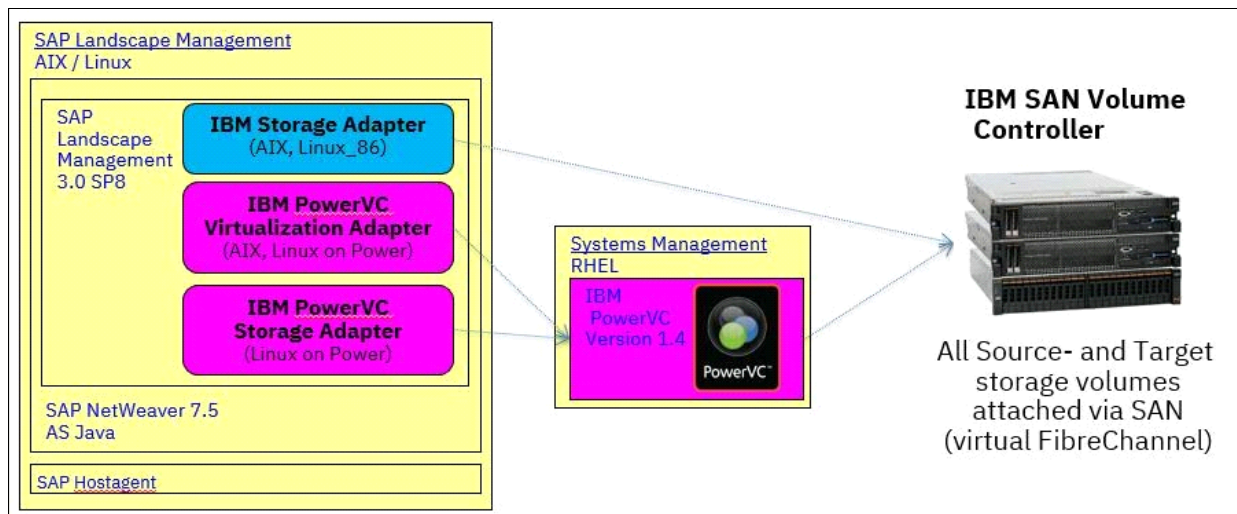
Figure 11   SAP LaMa infrastructure example

Most relevant for IBM Storage is the system copy (cloning) and the provisioning parts of SAP Landscape Management. These are the use cases that we are addressing with this Blueprint.

SAP Landscape Management uses IBM PowerVC or IBM Spectrum Protect Snapshot to start FlashCopies or snapshots on IBM storage. IBM PowerVC is also used for the provisioning of LPARs (VMs), storage volumes (LUNs), and so on (see Figure 12).
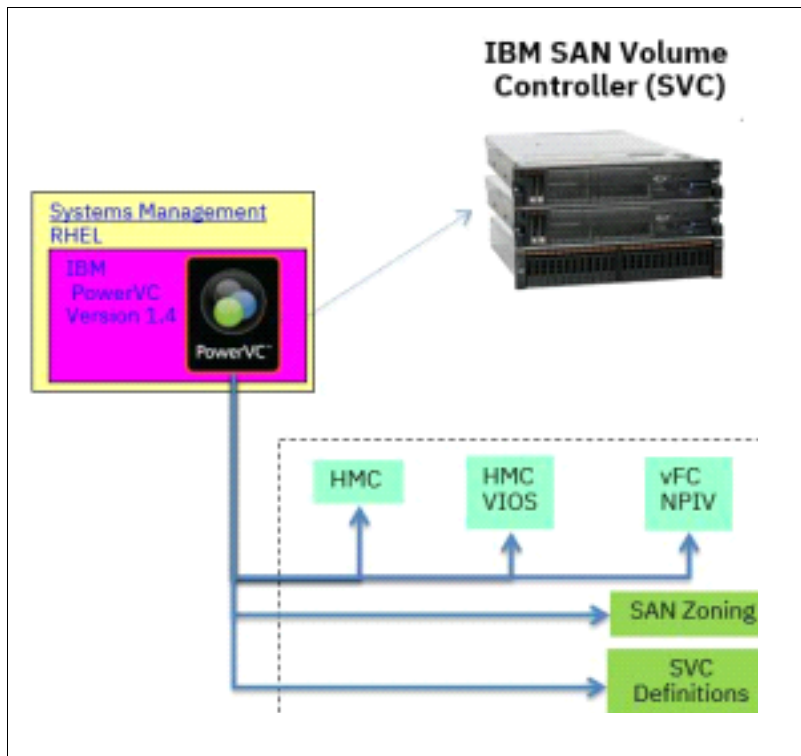
Figure 12   IBM Power Systems and storage management with IBM PowerVC

In combination with Post-Copy Automation (PCA, delivered by SAP), customers can reduce the time that is required to set up a QA, test, or education system (as copy from a production system) from days to a few hours. In this case, storage-based snapshots replace time-consuming backups and restores of SAP systems and PCA handles application-specific tasks that are required after a system copy.

The users of LaMa are SAP administrators who are interested in SAP systems, databases, and sufficient file systems capacity. IT operations (such as LUN provisioning and operating system installation) are started by LaMa, and can be transparent to the SAP administrator, but not to the user who configures LaMa. You can compare the roles to the car driver and the mechanic.

Figure 13 shows the minimum landscape to set up and test the LaMa functions in the IBM lab.
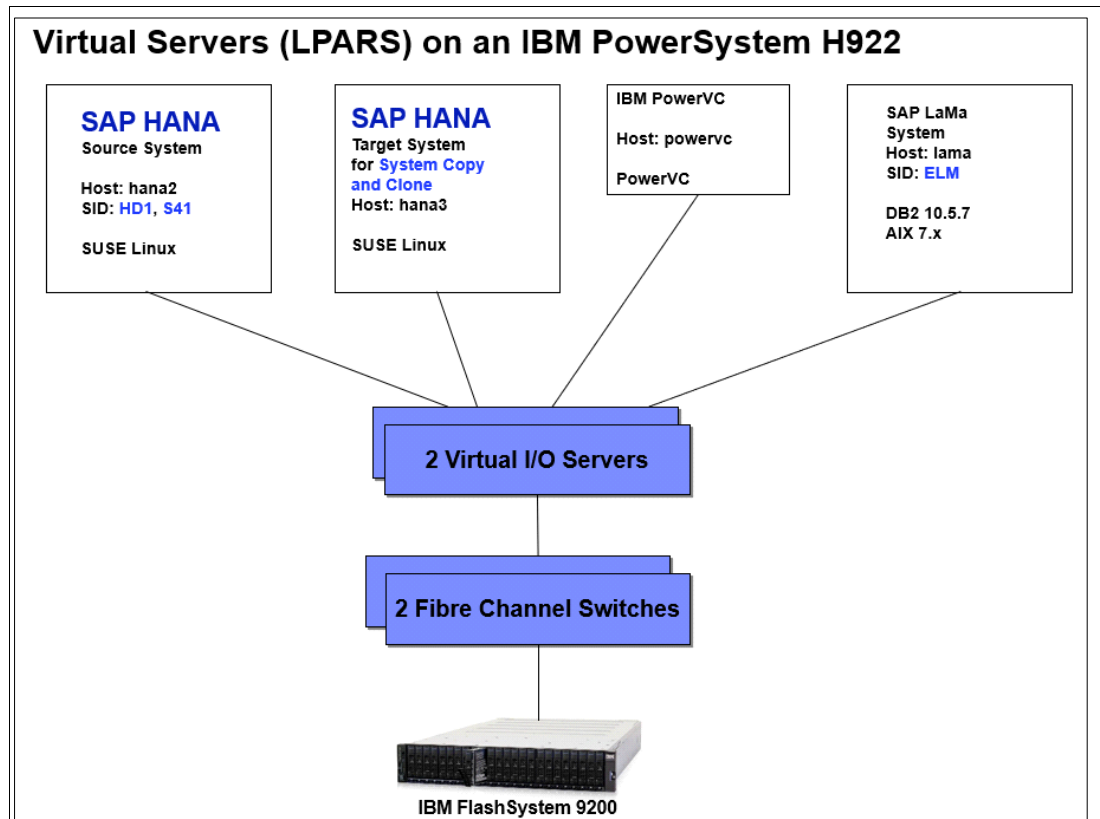
Figure 13   SAP Landscape Management 3.0 setup

Note: The SAP note 2050537, "Support for SAP HANA in SAP Landscape Management", provides a consolidated overview of the SAP Landscape Management release history and associated functionalities.

SAP Netweaver installation (and upgrade to latest SP)

An SAP Landscape Management installation requires the following major tasks:

• Install SAP NetWeaver Application Server 7.5 for Java.

• Download and install the Landscape Management software packages on the NetWeaver Application Server.

SAP NetWeaver runs on multiple operating systems. We chose IBM AIX® and set up the following software:

• IBM AIX 7.2 Technology Level 03 Service Pack 2 Built 1845
• SAP NetWeaver 7.50
• SAP Landscape Management 3.0 SP10

For more information about the prerequisites and installation process for SAP NetWeaver 7.5, see the SAP manual *Installation of SAP Systems Based on the Application Server Java of SAP NetWeaver 7.1 to 7.5 on UNIX: IBM DB2 for Linux, UNIX, and Windows (Database: IBM DB2 for Linux, UNIX, and Windows, Operating System: UNIX/Linux).*

21

For more information, see the following SAP Notes:

- 1972803: SAP on AIX: Recommendations
- 1707361: Inst. Systems Based on NW 7.1 and Higher: UNIX DB2 for Linux, UNIX, and Windows

These manuals and SAP notes describe the hardware and software requirements and the file system layout that is required for SAP NetWeaver with a specific database. We chose IBM DB2 database. Table 6 shows a sample file system layout.

Table 6 Sample file system layout

Filesystem	GB blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/db2bigSIDlv	19.88	19.71	1%	75	1%	/db2/ELM
/dev/db2sidlv	5.31	1.70	69%	7256	2%	/db2/db2elm
/dev/logdirlv	23.88	19.20	20%	67	1%	/db2/ELM/log_dir
/dev/sapdatalv	49.88	44.38	12%	10	1%	/db2/ELM/sapdata1
/dev/saptemp1lv	9.88	9.83	1%	4	1%	/db2/ELM/saptemp1
/dev/sapmntlv	11.88	9.19	23%	1233	1%	/sapmnt
/dev/usrsaplv	59.38	42.48	29%	550624	6%	/usr/sap

The SAP NetWeaver installation is described in the SAP manual, *SAP Landscape Management 3.0, Enterprise Edition, Installation Guide*, which is available on the SAP Help Portal.

The installation of SAP NetWeaver requires the following SAP tools:

- Software Provisioning Manager (SAP installer, sapinst) to install SAP components
- Software Update Manager 1.0 to perform release upgrades, update single components, and install additional packages

We also recommend to use the SAP Maintenance Planner to find and download appropriate and compatible software packages.

To work with the SAP Maintenance Planner, see the SAP Help Portal and search for "Maintenance Planner". Then, click the button of the corresponding application.

> **Note:** For more information about SAP Landscape Management 3.0, enterprise edition, as an add-on to SAP NetWeaver Application Server for Java, see SAP note 2350235, "SAP Landscape Management 3.0 - Enterprise edition".

SAP Landscape Management installation and upgrade to latest SP

SAP provides the Landscape Management software called SAP Component Archive (SCA) files, which can be installed and updated by using the Software Update Manager (SUM).

SUM and the SCA files are available for download from the SAP Support Portal. At the portal, select **Download Software**.

Log on to the LaMa server as `<sid>adm` and unpack the Software Update Manager package by using the following command:

```
SAPCAR –xf <download directory>/<path>/<Archive>.SAR -R /usr/sap/<sid>
```

For example:

```
./sapcar -xf /instimages/SUM/SUM10SP23_1-20006540.SAR -R /usr/sap/ELM
```

To start the Software Update Manager, run the following command as the root user:

```
/usr/sap/<SID>/SUM/STARTUP <SID>
```

Then, open a web browser and access the Software Update Manager by using the following URL:

```
http://<hostname>:1128/lmsl/sumjava/<SID>/index.html
```

The Software Update Manager guides you through the installation process (see Figure 14).
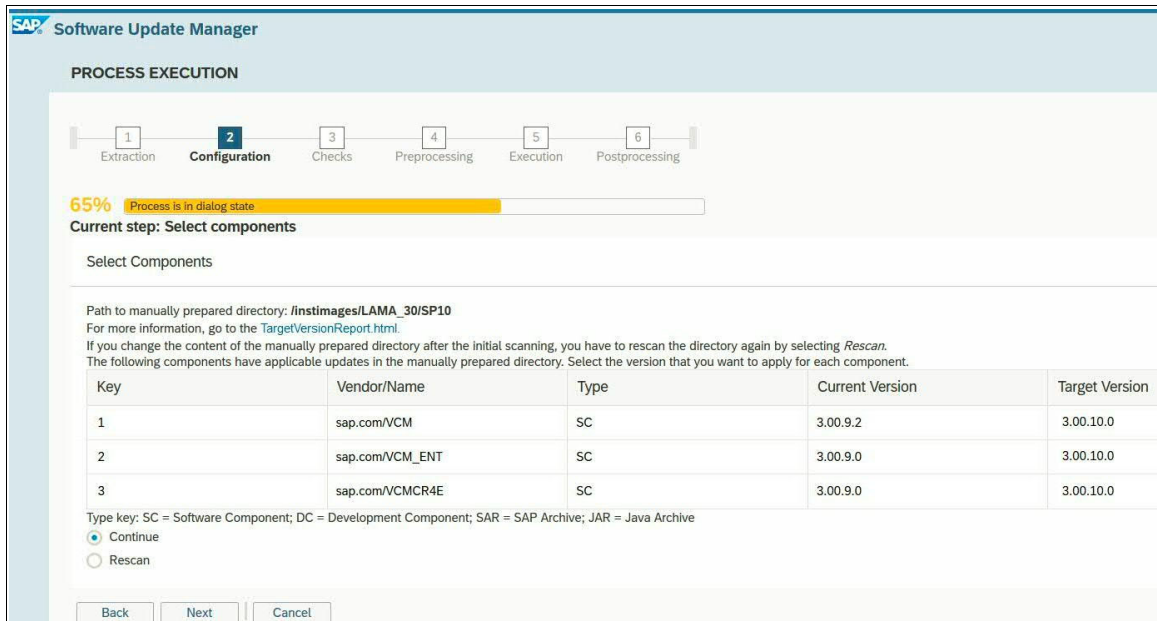
Figure 14   SAP LaMa upgrade with Software Update Manager

Three SCA packages are required for the SAP LaMa 3.0 Enterprise Edition: VCM, VCM_ENT, and VCMCR4E.

For more information about SUM and the LaMa installation, see the following SAP manuals:

• *Updating SAP Java Systems on UNIX and Linux: IBM DB2 for Linux, UNIX, and Windows (Upgrade Guide -Software Update Manager 1.0)*

• *SAP Landscape Management 3.0, Enterprise Edition, Installation Guide*

Initial SAP Landscape Management setup

Now that SAP LaMa is installed, log on to the SAP NetWeaver Administrator (NWA) application and complete the following steps:

1. Open `http://lamahost:50000/nwa` URL in a web browser.

2. In the SAP NetWeaver Administrator application, select **Configuration** → **Scenarios** → **Configuration Wizard**. Then, select I**nitial Setup for SAP Landscape Management**, (see Figure 15 on page 24).

Figure 15   Initial setup for SAP LaMa with the Configuration Wizard in SAP NWA

3. Access the SAP Landscape Management application by using the URL and perform the basic configuration steps for your SAP landscape.

More information

The *SAP Landscape Management 3.0, Enterprise Edition User Guide* is available in the SAP Help Portal and describes the initial setup and the configuration steps for the LaMa use cases. Access the SAP Help Portal and search for "Landscape Management 3.0" to download this manual.

For more information, see SAP note 2039615: *Managing system landscapes with SAP Landscape Management Enterprise Edition*.

For more information about the configuration steps that are required for this IBM Blueprint sample setup are described in "SAP LaMa basic setup and configuration" on page 29.

IBM System integration to SAP LaMa

This section describes the integration of SAP LaMa into an IBM Power Systems environment. For the integration of virtualized IBM PowerLinux systems and IBM storage, the IBM Power Virtualization Center (PowerVC) is required, as shown in Figure 16 on page 25. If the SAP system was installed on Linux_x86, the PowerVC counterpart IBM Spectrum Protect Snapshot is required.
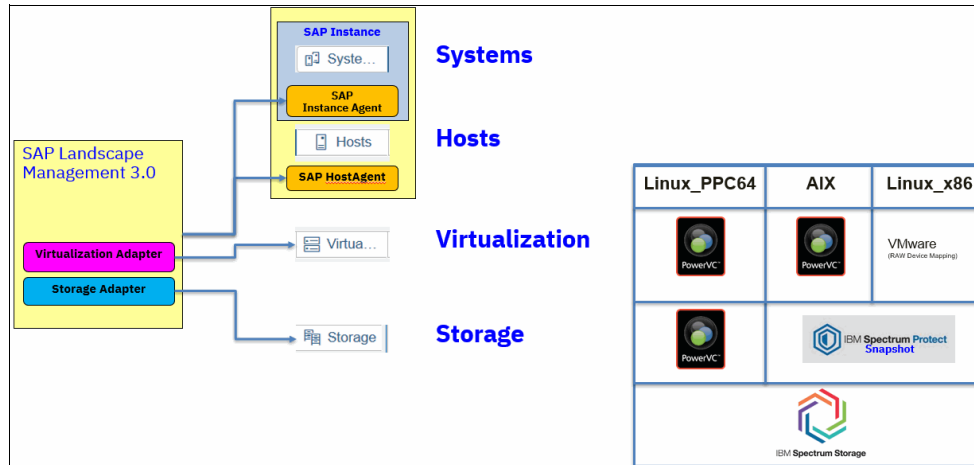
Figure 16   Integration of SAP LaMa with IBM Power Systems

The following IBM components must be integrated into SAP LaMa:

- IBM Power Virtualization Center (IBM PowerVC) for virtualization management
- IBM FlashSystem 9200 storage system
- IBM PowerSystem H922
- SAP Adaptive Extensions with IBM libraries

For the host communication with SAP LaMa, the SAP Host Agent must be installed on every managed SAP host. The SAP Host Agent is the most important agent in SAP LaMa. Ensure that you keep the Host Agent updated with the latest patch level on each host in the landscape. The SAP Host Agent provides an update feature that helps you with that task.

The SAP Host Agent enables SAP LaMa to start and stop instances, and to monitor the status of computer systems.

For more information, see SAP note 1031096: *Installing Package SAPHOSTAGENT*.

IBM PowerVC Setup

IBM PowerVC Virtualization Center is an advanced virtualization and cloud management offering, which is built on OpenStack, that provides simplified virtualization management and cloud deployments for IBM AIX, IBM i, and Linux VMs running on IBM Power Systems. PowerVC is designed to improve administrator productivity and simplify the cloud management of VMs on Power Systems servers.

To manage SAP HANA environments on PowerLinux with SAP LaMa, PowerVC must be integrated into SAP LaMa. PowerVC manages the IBM Power Systems and IBM storage system for the SAP LaMa tasks.

> **Note:** The minimal supported IBM PowerVC version is 1.4.0.

To install PowerVC, see Installing IBM Power Virtualization Center.

The pre-installation checklist for PowerVC (`PVC_checklist.pdf` that is attached to SAP Note 2372251) provides more information about the environment that is required for supporting the IBM PowerVC Storage Adapter for SAP LaMa 3.0. It also specifies the supported versions of the hardware and software components that compose the system landscape for SAP LaMa Tasks.

IBM PowerVC includes the following tasks:

- An IBM PowerVC user with Administrator privileges is required. SAP LaMa uses this user for creating the LPARs, creating and deleting LUNs, and mapping them to the target LPAR.

- Communication between SAP LaMa and IBM PowerVC by using "server certificates" provides the most secure level. The IBM PowerVC certificate is used, and the host name within the certificate is verified against the host name in SAP LaMa. In this case, the IBM PowerVC server certificate must be imported into the keystore of the SAP NetWeaver engine on which SAP LaMa is running.

- PowerVC must be defined in LaMa as Virtualization manager and as Storage manager as described in the following steps:

  a. To register PowerVC as Virtualization manager, browse to C**onfiguration** → **Infrastructure** → **Virtualization Managers**, as shown in Figure 17.
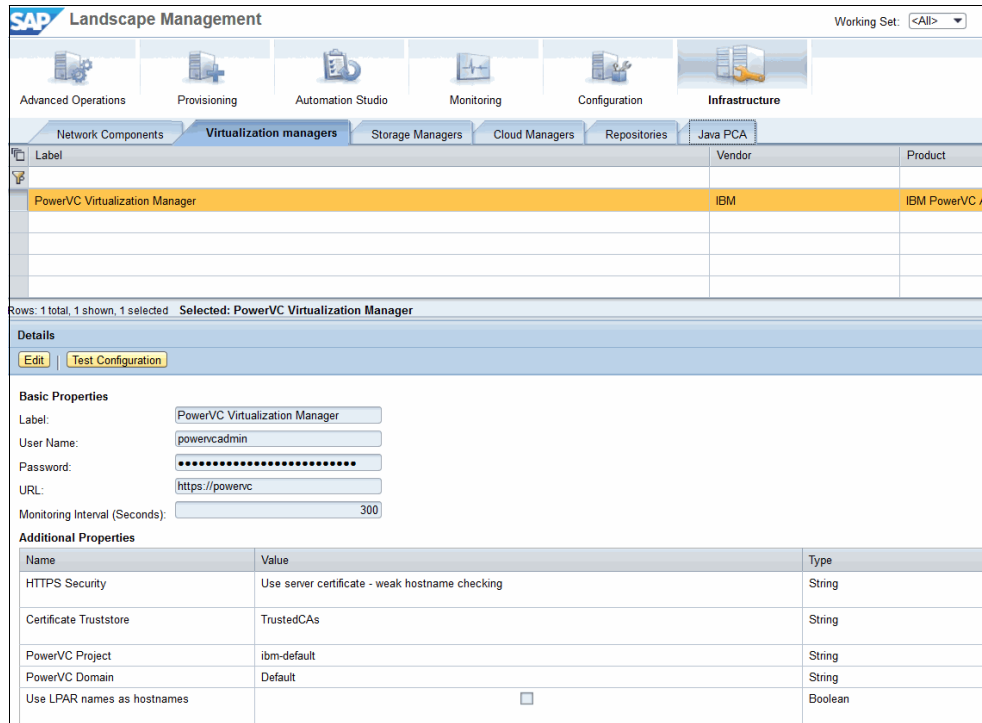
Figure 17   PowerVC integration as Virtualization manager

  b. To register PowerVC as Storage manager, browse to **Configuration** → **Infrastructure** → **Storage Managers**, as shown in Figure 18 on page 27.
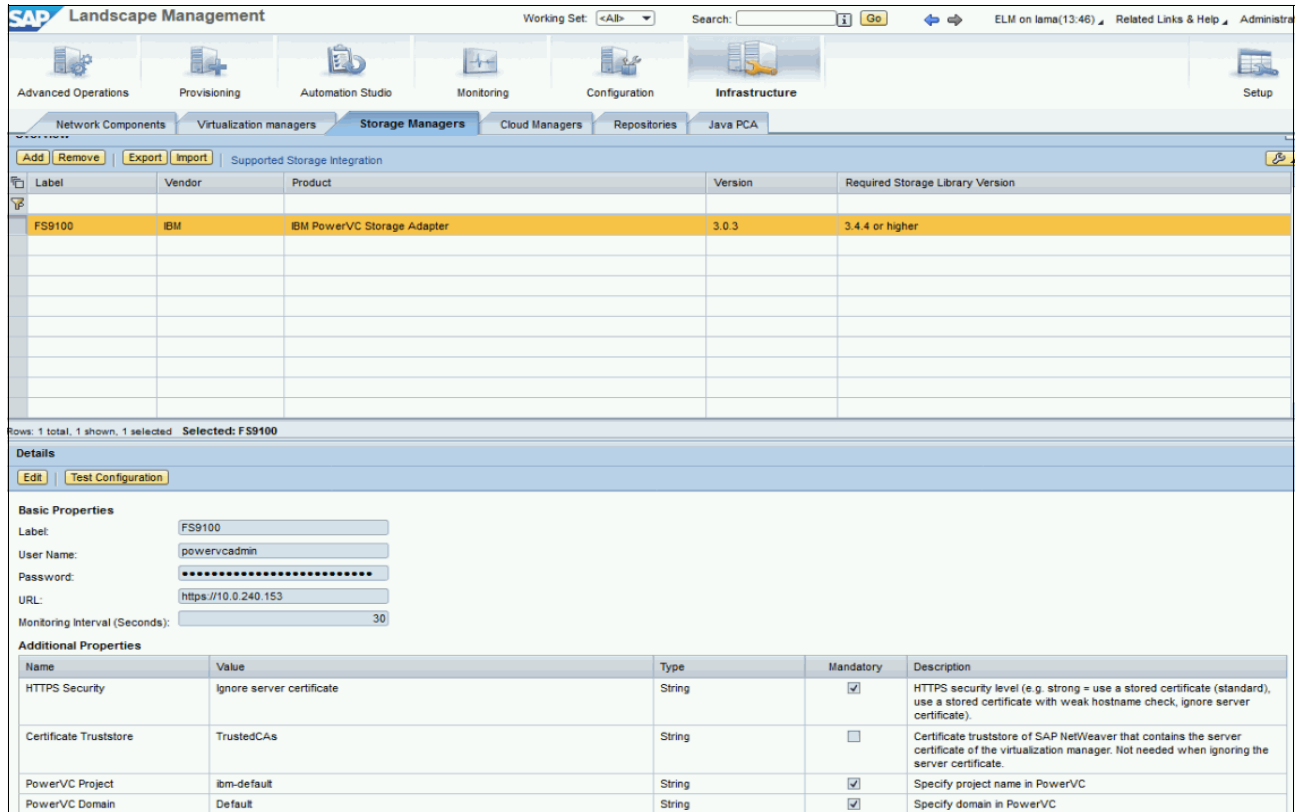
Figure 18   PowerVC integration as Storage manager

During the execution of the use cases, SAP LaMa calls:

- Storage commands by way of PowerVC to retrieve information about the storage volumes that contain SAP systems.

- Storage commands by way of IBM PowerVC to configure and run the cloning and copying process of SAP systems.

For more information about IBM PowerVC and SAP LaMa, see the following SAP notes:

- 2372251: *Integration of IBM Components into SAP LaMa for Power*
- 2372252: *Configuration of IBM PowerVC virtualization adapter for SAP LaMa*
- 2075173: *Integration of IBM Components into SAP LVM 2.1 for IBM Power*
- 2372935: *Configuration of IBM PowerVC Storage Adapter for SAP LaMa*

IBM FlashSystem family

IBM FlashSystem 9200 systems are virtualized, NVMe-based, all-flash, software-defined storage solutions that unlock the power of your data by combining the latest IBM Spectrum Virtualize, IBM FlashCore®, IBM Storage Insights, and IBM FlashSystem 9200 multicloud software technologies.

The IBM FlashSystem 9200 storage system, as shown in Figure 19 on page 28, runs IBM Spectrum Virtualize software and introduces the following remarkable new features in comparison to the predecessor models:

- NVMe drive support: NVMe is a new logical device interface standard from 2011 for accessing non-volatile storage media that is attached by way of a PCI Express bus.

- Lower latencies through RDMA, direct memory access from the memory of one node into that of another without involving either operating system.

- Clustering option with the older V7000 systems.

Because traditional storage systems can provide capacity for established backup environments, the new IBM FlashSystem 9200 systems are the storage foundation for building the next generation Data Protection and Reuse solutions:

- The IBM Spectrum Virtualize code stack allows you to create snapshot copies of VMs quickly and efficiently
- IBM FlashSystem 9200 Multi-Cloud solution for data reuse, protection, and efficiency

It is recommended to use fully equipped 9200 units only, with 24 NVMe Flash Core Modules.

Figure 19   IBM FlashSystem 9200 NVMe

For more information about the IBM FlashSystem, see *IBM FlashSystem 9200 and 9100 Best Practices and Performance Guidelines*.

For more information about which IBM storage models are certified for SAP HANA, see the Certified and Supported SAP HANA Hardware Directory.

IBM Power System H922

IBM Performance Optimization With Enhanced RISC (POWER) IBM POWER9™ servers were launched in 2018 and are based on high-performance processors, the ninth generation of its kind. POWER9 servers are one of the most powerful, secure, and scalable servers in the world. The top-of-the-line POWER9 server model is E980, which can support up to 192 cores and 64 TB of memory. POWER9 servers also are available in entry-level and midrange models H922 and H924. POWER9 servers support IBM AIX, Linux, and IBM i operating systems.

Built and optimized for SAP HANA, the IBM Power System H922 (see Figure 20) and H924 boast a large memory footprint of up to 4 TB in a dense form factor for gaining insights faster while delivering high reliability and performance. For this Blueprint setup, a H9223-22H Power System was used and the H922 featured following configuration:

- CPU: 20
- Memory: 512 GB
- Fibre Channel adapter: 2 x PCIe2 16 Gb 2-port Fibre Channel adapter
- Network: PCIe3 2-port 25/10Gb NIC@ROCE SFP 28 Adapter

Figure 20   IBM Power System H922

The communication between SAP LaMa and the H922 is done by way of the PowerVC server. In SAP LaMa, no configuration is required for the PowerSystem.

For more information about the IBM H922, see *IBM Power Systems H922 and H924Technical Overview and Introduction*.

SAP Adaptive Extensions

The SAP Adaptive Extensions (SAPACEXT*.SAR) package contains the IBM storage and platform libraries for SAP LaMa. The libraries must be installed in addition to SAP Host Agent for performing various operations by using SAP Landscape Management software. You install SAP Adaptive Extensions on every managed host where SAP HANA is installed.

Download the latest patch for your platform of SAPACEXT.SAR from the SAP Help Portal. (At the portal, click **Support Packages and Patches** → **Browse Download Catalog** → **SAP Technology Components** → **SAP ADAPTIVE EXTENSIONS**.

This package contains the following main components:

- `libsaposprep.o`

  This library includes specific functions for activating and deactivating virtual IP addresses and handling of IPSec firewall.

- `libsapacosprep_ibm.o`

  This library includes specific functions for mounting and unmounting NFS/GPFS file systems.

For more information about installing the SAP Adaptive Extensions, see "Basic host configuration" on page 30.

For more information about the SAP adaptive extensions, see the SAP Help Portal and SAP Note 1759181.

SAP LaMa basic setup and configuration

This section describes how to set up SAP LaMa before any landscape entities are configured. The setup can be accessed after logging in by clicking **Setup** in the main menu.

The earlier style Web Dynpro view is opened for the setup GUI (see Figure 21).
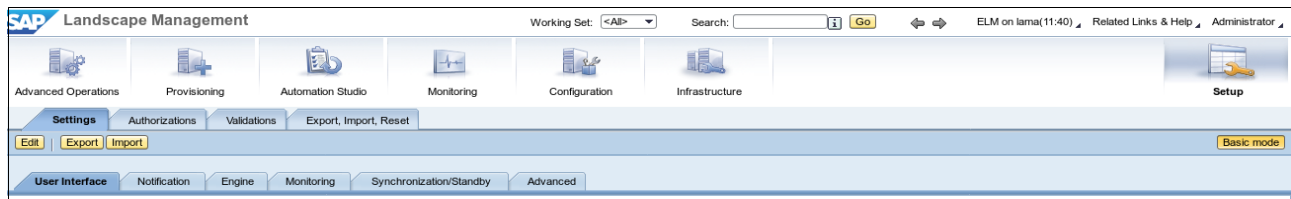
Figure 21   Web DynPro main window

Established GUI

SAP LaMa 3.0 SP10 offers a new and advanced Web GUI, but the Web Dynpro interface is still available because it is required for some of the landscape configuration tasks. Some of the views of these established applications are marked as deprecated, but it makes sense to re-enable them for a few reasons. For example, much of the documentation (SAP notes and external docs) refers to the Web Dynpro GUI; therefore, it makes life easier if it is available.

To enable the established views, browse in the Setup sheet to **Settings** → **User Interface**. Ensure that **Activate Deprecated Web Dynpro Views** and **Display Individual Instance Operations** are enabled by selecting them.

SAP LaMa Engine Settings

The Engine Settings are used to control the communication between SAP LaMa and the landscape components. It is important to adjust these settings before starting with the integration of any component so they fit your needs.

To access the SAP LaMa Engine Settings, browse in the Setup sheet by clicking **Settings** → **Engine.** Ensure that **Expert Mode** is enabled for this view by selecting the **Expert Mode** option (see Figure 22).
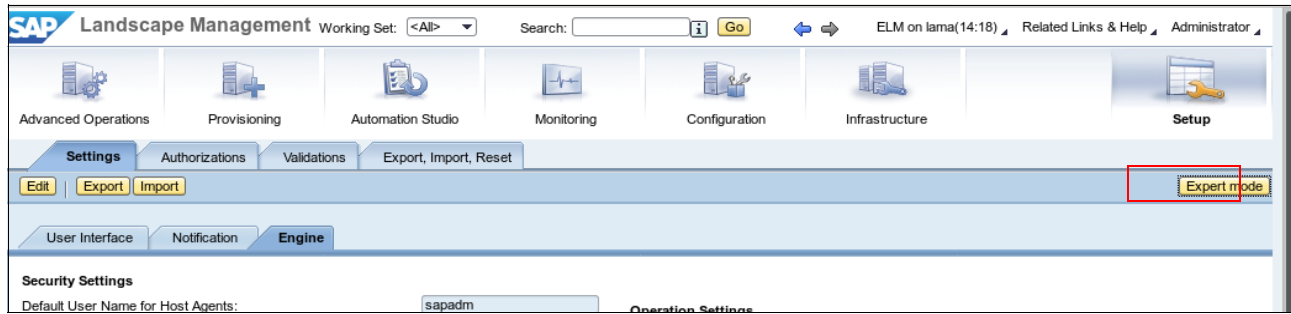
Figure 22   Enable Expert Mode Setting

Security settings

You must define common or individual credentials for the SAP Host Agents of the landscape. It is acceptable practice to define a common user and password for all Host Agents. However, common users might not be allowed for security reasons, or password authentication is forbidden at all, especially in larger environments. In such cases, select the **Allow Configurations or Credentials Individually per Host Agent** option. By using this option, you can specify individual users and passwords or an X.509 certificate per Host Agent.

In the demonstration environment, a common user name and password are used. In the same way, the settings for SAP Instance Agents can be adjusted to meet a higher level of security requirements.

Operation settings

Adjust these settings to enable the execution and management of multiple instances on the same host. To run multiple instances (not databases) select the **Allow multiple instances on the same Host** option. This option is required if you plan to run a system database with a central instance on the same host, for example.

Select the **Allow Multiple Database Instances on Same Host** option if you plan to run multiple database instances on the same host. For more information, see SAP Note 919135.

Select the **Allow MountPoint Creation** option if you want to copy or clone the S/4HANA database. During the provisioning process, missing mount points are created on the target host.

Basic host configuration

This section describes how to integrate a Linux or AIX host into the SAP landscape. SAP LaMa uses an agent for communication called SAP Host Agent. This document describes the manual installation of the SAP Host Agent, which works in the same way for Linux and AIX.

The following prerequisites must be met:

- If the host runs a firewall, ensure that TCP ports 1128 and 1129 are accessible from the LaMa system.

- Obtain SAP Software Provisioning Manager (SWPM) and Host Agent for your operating system and architecture from SAP.

For more information, search the SAP Help Portal for "SAP Host Agent".

The SAP Host Agent can be installed manually or by using the SAP Software Provisioning Manager, which is required for other SAP installations.

Because the Host Agent software is shipped as an SAP Archive, the SAP proprietary SAPCAR.EXE application must be downloaded to unpack the archive.

Complete the following steps:

1. Unpack the agent archive into a temporary directory, as shown in the following example:

```
hana1: cd /tmp/host_agent_install_dir
hana1: /tmp/SWPM/SAPCAR -xvf /tmp/SAPHOSTAGENT40-40-70002261.SAR -manifest
SIGNATURE.SMF
```

It is important to specify the -manifest option because the SAP signature file that is required for validating the software is not extracted otherwise.

2. Install the software, as shown in the following example:

```
hana1: cd /tmp/host_agent_install_dir
hana1: ./saphostexec -install -verify
```

The SAP Host Agent is installed into the /usr/sap/hostctrl directory and the sapadm user is created if it does not exist.

3. Set a password for sapadm, as shown in the following example:

```
hana1: passwd sapadm
```

If the SAP LaMa Engine was set up to use a common password for all Host Agents, the sapadm password must be set.

4. Ensure that the agent was installed correctly and is working, as shown in the following example:

```
hana1:  /usr/sap/hostctrl/exe/saphostexec -version
************************* Component ********************
…
********************************************************
-------------------
SAPHOSTAGENT information
-------------------
kernel release           721
kernel make variant      721_REL
compiled on              Linux GNU SLES-12 ppc64le cc4.8.5  for
linuxppc64le
compiled for             64 BIT
compilation mode         Non-Unicode
compile time             Oct 20 2018 04:32:59
```

For Linux, SAP offers the SAP Host Agent software alternatively as an RPM package, which makes the installation process less manual.

Installing SAP Adaptive Extension

SAP Adaptive Extensions are required by LaMa to manage SAP HANA systems by way of the Host Agent. The software must be downloaded from SAP as a SAR archive.

Install the downloaded archive by using the following command:

```
hana1: /usr/sap/hostctrl/exe/sapacosprep -a InstallAcExt -m
/tmp/SAPACEXT_44-20010482.SAR
```

SAP LaMa basic configuration

This section covers the SAP LaMA basic configuration.

Pool creation

SAP landscapes often contain many entities, which might become confusing quickly. Therefore, SAP Landscape is organized into *pools*, which are a logical grouping of coherent components. Pools can be grouped in a higher level that is called *containers*, but this process is out of the scope of this Blueprint because we do not work with containers in the demonstration environment.

Creating at least one pool is required because all components that are added later require a pool to which they are assigned.

Complete the following steps to create a pool:

1. Click the **Configuration** icon in Web Dynpro.

2. In the Pools tab, click **Add Pool**.

3. Define a freely chosen name and select **Save** to save the pool (see Figure 23 on page 33).
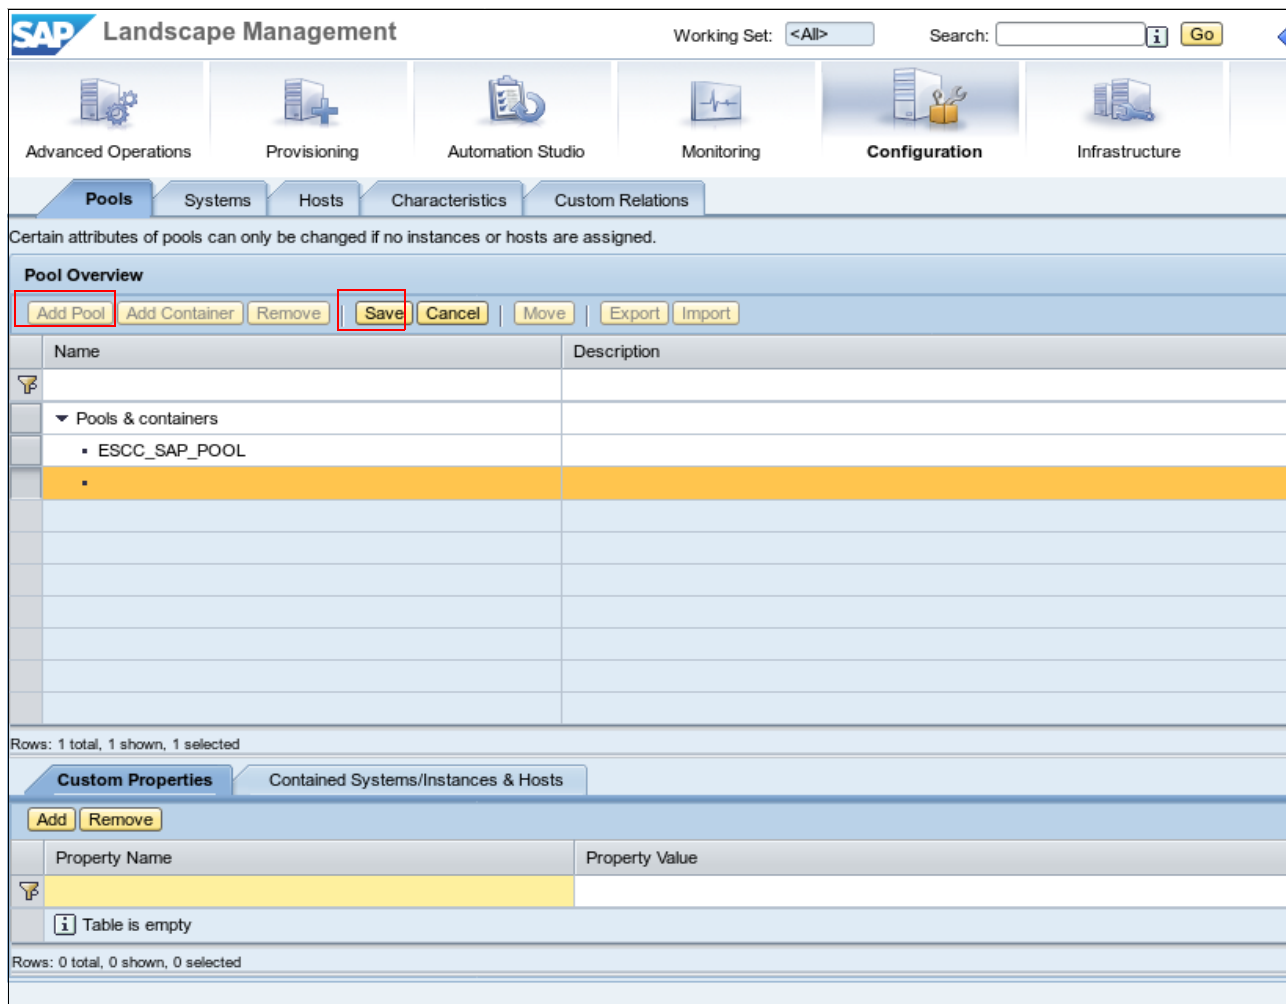
Figure 23   Adding Pools

SAP landscape infrastructure configuration

Landscape infrastructure components are required for most of the landscape management tasks.

Although hosts and SAP instances can be configured without any infrastructure components, it is good practice to define the required infrastructure before the SAP systems are configured for SAP LaMa.

The demonstration environment includes configured network components, software repositories, virtualization managers, and storage managers.

The following sections discuss the purpose of these components.

Network configuration

At least one network must be defined in an SAP LaMa environment. A network definition is a IPv4 network that contains the IPv4 broadcast address and the netmask of this network (see Figure 24 on page 34).

In an advanced configuration step (not covered here), an automated DNS update can be established. For more information, see SAP Note 1572841.
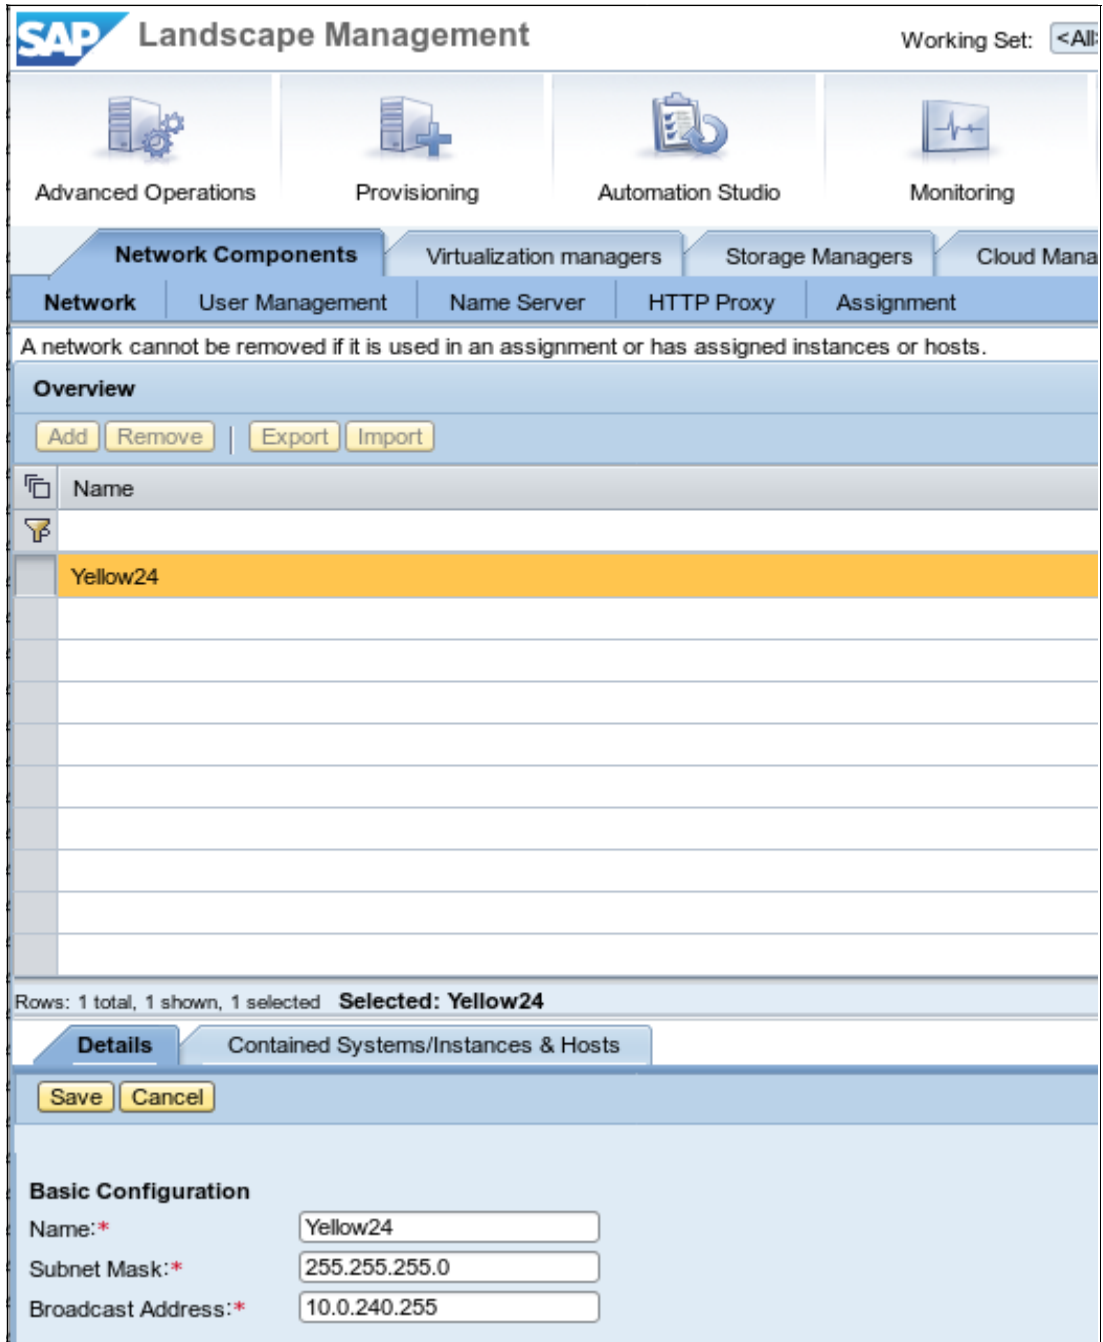
Figure 24   Adding network components

Storage Adapter configuration

Storage Adapters can be used for obtaining information about the attached storage environment. SAP LaMa offers different adapter types, which are all specific to the vendor and product.

In the demonstration environment, the IBM PowerVC Storage Adapter is used. For more information about the configuration and integration of this component, see "IBM System integration to SAP LaMa" on page 24.

Virtualization Manager configuration

Virtualization Managers often are deployed to manage a data center's virtualization environment. SAP LaMa makes use of some of those Virtualization Managers. LaMa manages the SAP relevant parts of a virtual environment through the Virtualization Manager.

The demonstration environment integrated the IBM PowerVC Virtualization manager. For more information about the configuration and integration of this component, see "IBM System integration to SAP LaMa" on page 24.

SAP software repository configuration

SAP software repositories are also defined in SAP LaMa. They are used for system updates and tasks, such as cloning or deploying to a new system. Repositories are defined as paths, where the required software components must exist. It is the responsibility of the administrator to ensure that a repository is complete.

SAP LaMa uses some predefined provisioning types for a repository (see Figure 25).

Figure 25   Software Repositories

SAP instance configuration

SAP LaMa communicates with SAP instances by way of agents, SAP Host Agents, and SAP Instance Agents. Host Agents are required for performing landscape management tasks that interact with a host, such as preparing and unpreparing SAP instances and databases. SAP LaMa managements tasks that interact with the instances require the SAP Instance Agents.

Other types of communication are done by way of the storage and virtualization adapters because not all of the required information for all possible management tasks can be gathered from hosts and instances only.

The more provisioning tasks that are not directly related to hosts and instances require those external tools. Consider that cloning an SAP instance requires setting up the storage on the target host, including items, such as network and SAN configuration.

Discover Hosts

The first step of instance integration is to define hosts. Because SAP Host Agents are running on each of the hosts, we can now configure LaMa to connect to these hosts and add them as SAP hosts. *SAP host* means here that the host can run SAP instances or databases.

To discover hosts in LaMa, click the **Configuration** icon in Web Dynpro. In the **Hosts** tab, select **Discover** (see Figure 26).
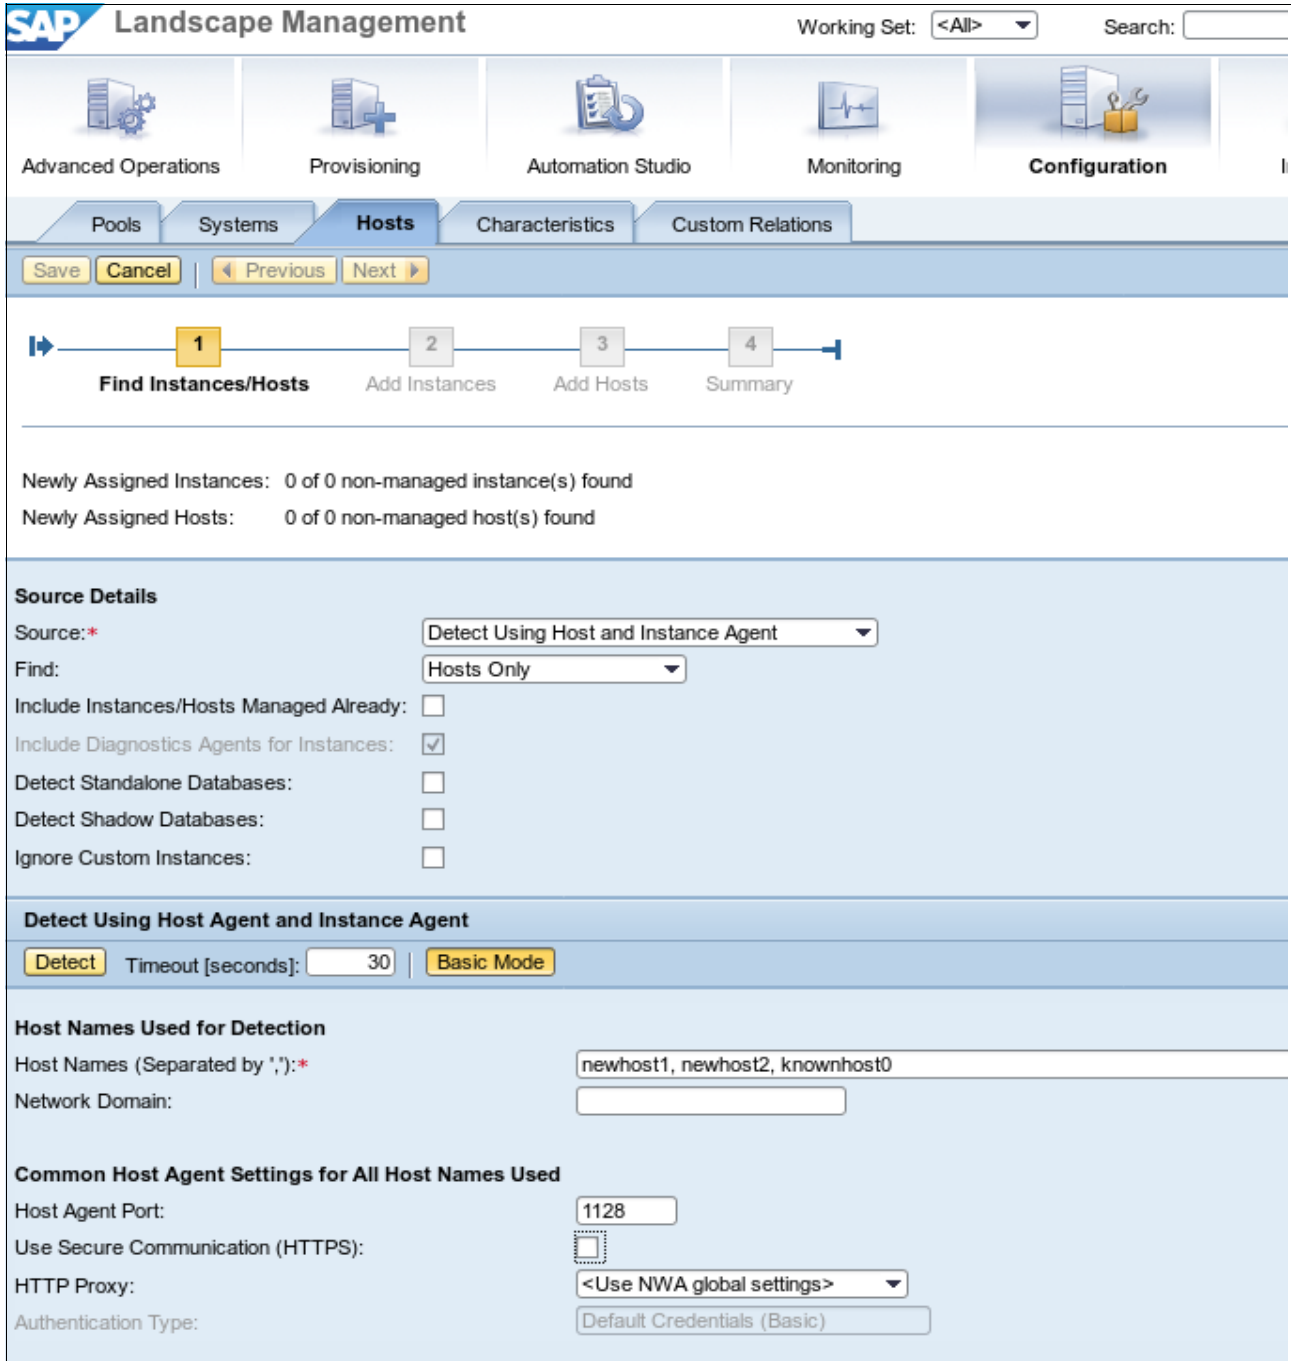
Figure 26   Search for Hosts

The detection source must be specified first because we want to detect hosts. Therefore, **Detect Using Host and Instance Agent** must be selected. Only hosts should be detected by this step; therefore, **Hosts only** must be found.

The resolvable names of the hosts and the network port of the Host Agent must be specified. This port usually is port 1128, which establishes a plain HTTP connection. If a higher level is required, a HTTPs connection can be chosen by using port 1129.

The HTTPS connection enforces at least TLS version 1.0. For more information about enabling TLS 1.1 and 1.2. in NetWeaver, see SAP note 2417205.

The Authentication Type is fixed to use the default credentials as defined by the SAP LaMa Engine setup.

The hosts should now appear in the Host Overview table, and each can now be edited for setting more features.

A wizard guides through the following configuration process (see Figure 27):

1. Change or verify the configuration. The version of the SAP Adaptive Extensions also can be queried here.
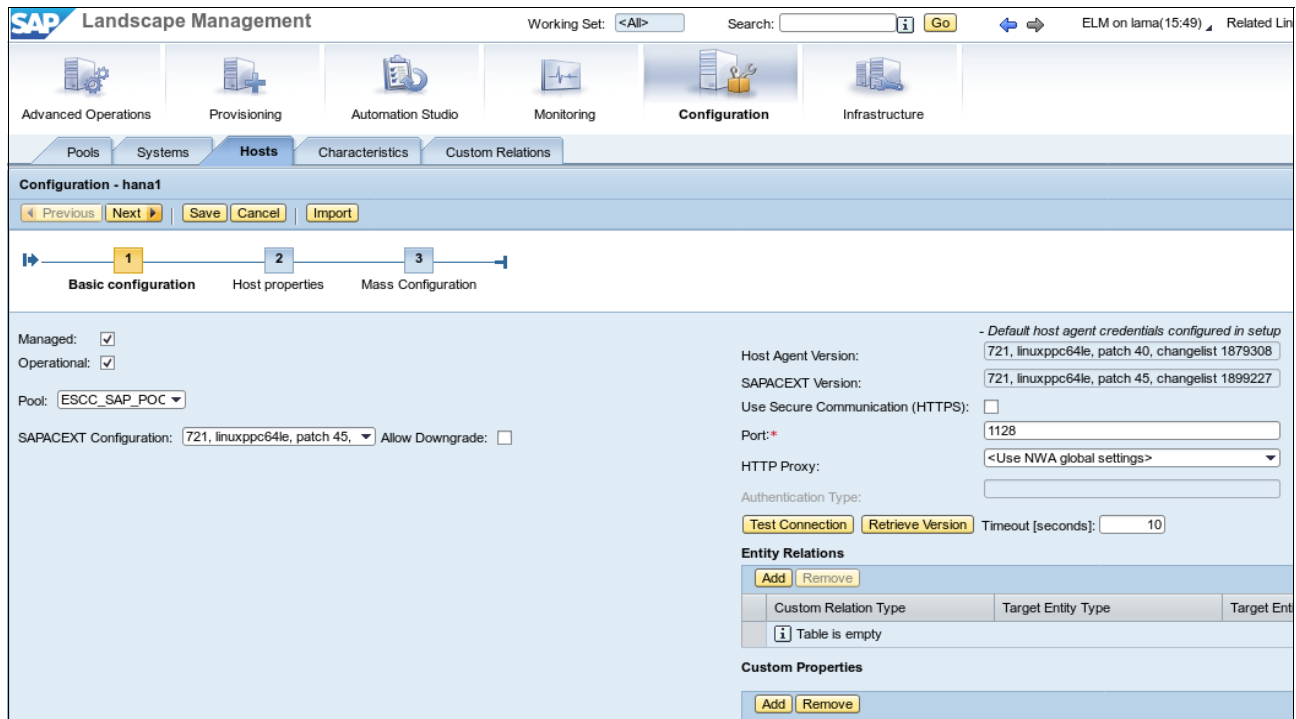
Figure 27   Configuring a detected host

2. Configure more host properties (see Figure 28 on page 38).

   Settings must be defined whether that host is enabled for Adaptive Extension and for system isolation.

   The formerly defined network must be assigned to the host so that SAP LaMa knows how to handle the network connection to that host when it performs system provisioning.

   Host-specific storage properties (such as SAN WWPNs) are also specified here. These values can be retrieved directly from host by using the Adaptive Extensions.
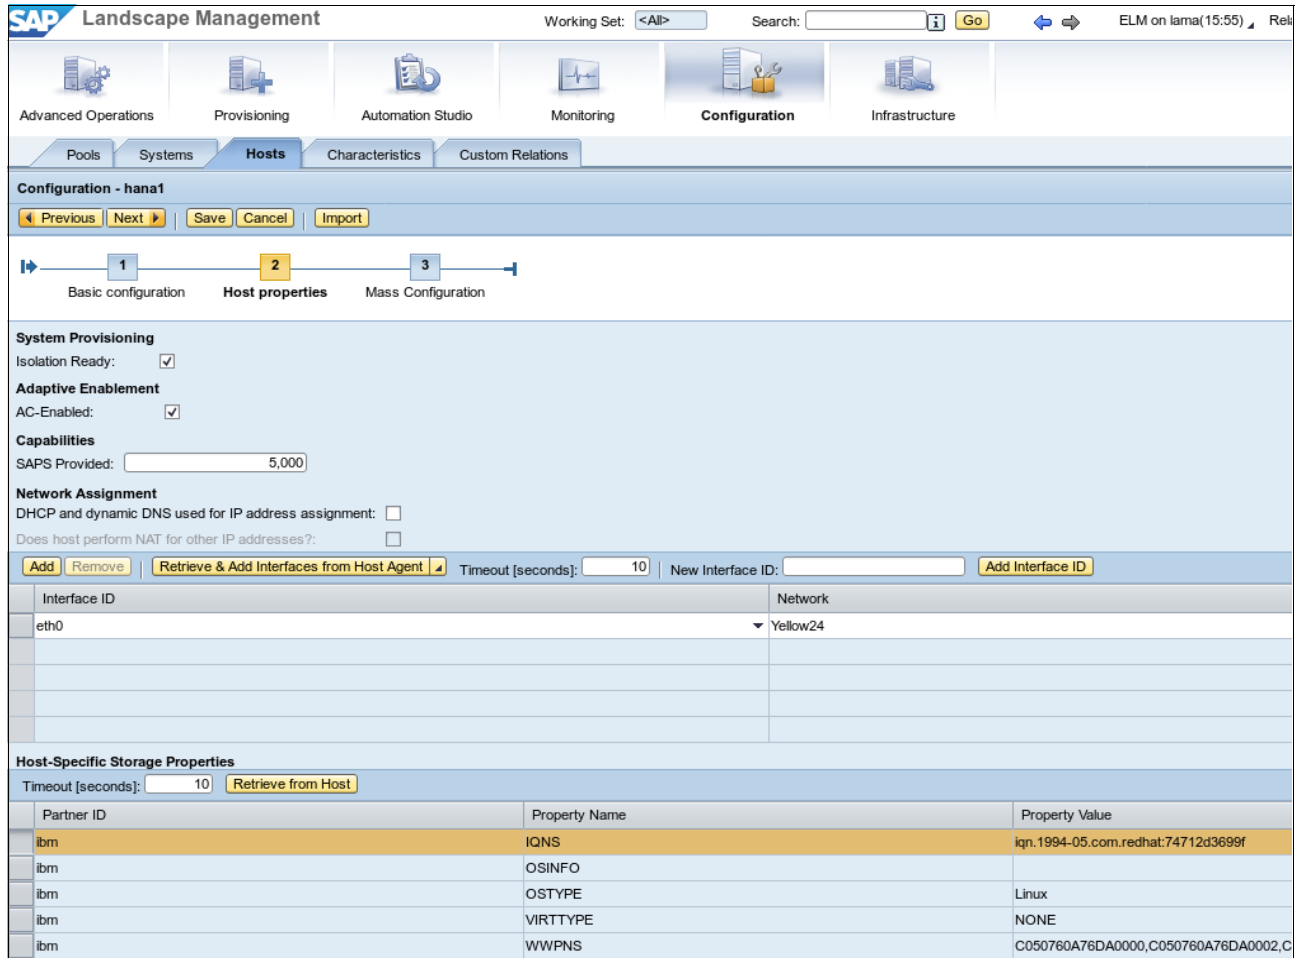
Figure 28   Configuring more host properties

Discover SAP Instances

SAP Instances must be discovered by SAP LaMa before they can be integrated into the Landscape. LaMa offers different discovery sources, so there are many ways to detect and integrate systems.

Systems that are at a host can be detected by using the SAP Host Agent. However, if the current location of a system is unknown, the Landscape Scanner is used, which scans the local network for systems.

Regardless from which source a system is discovered, it is accessed by using the instance agent, which is a function or piece of software that is inside the system. Instance agents work similar to host agents. They use predefined network ports and are configured for several authentication mechanisms.

For more information about configuring Instance Agents, see SAP Note 1474866.

In the test environment, we use the basic authentication feature. However, because the databases were not yet installed by using a common user name, we must specify individual credentials for each system.

In this scenario, complete the following steps to detect systems by way of existing Host Agents:

1. In the web GUI, click **Configuration** → **Systems**. The Web Dynpro view opens (see Figure 29).
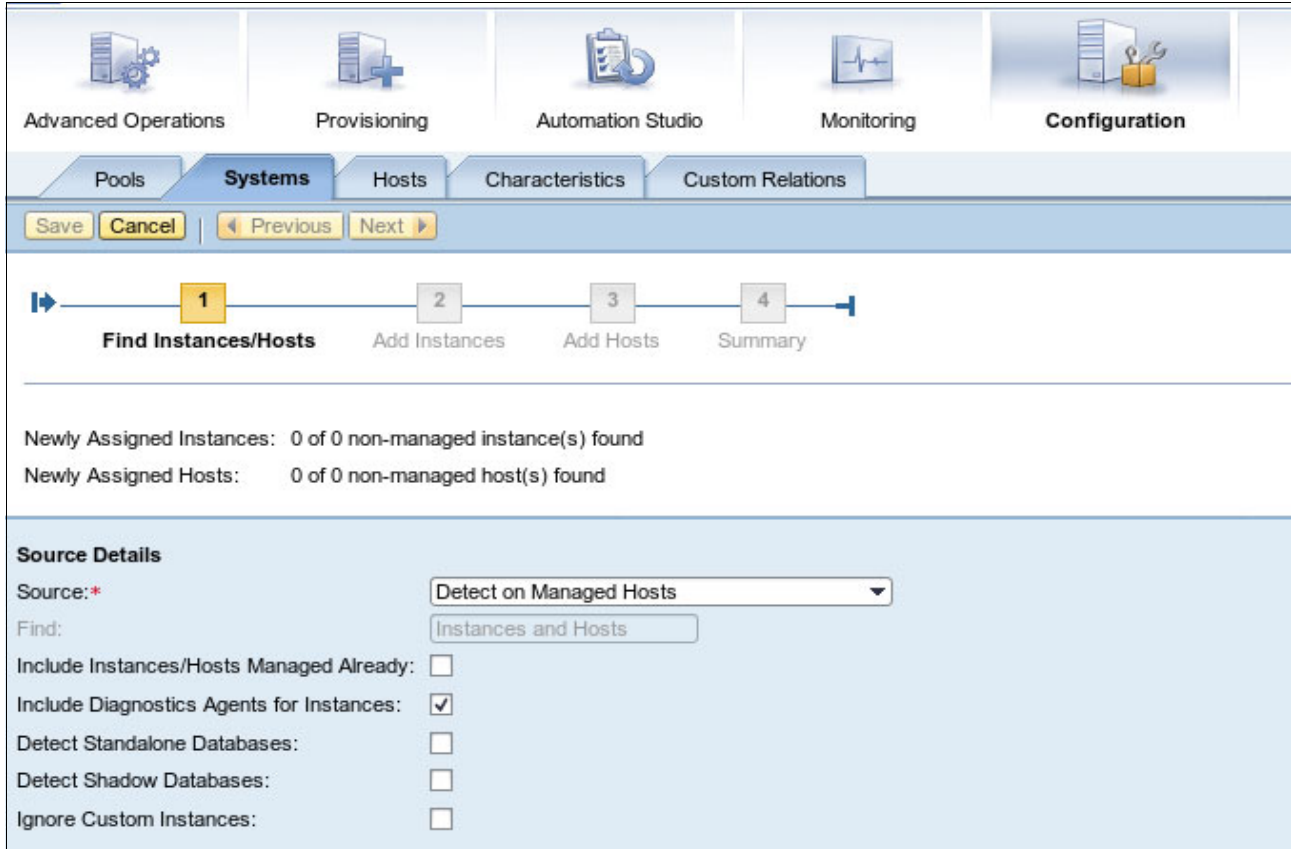
Figure 29   Discover SAP instances using SAP Host Agent

2. Click **Discover** to start the configuration mask for running a discovery.

3. Choose the hosts on which you want to detect systems.

At the bottom of this window are the settings for the instance agent that must be specified (see Figure 30).

Figure 30   Instance Agent credential settings

Consider the following points:

- It is important to select **User Name/Password (Basic)** as the Authentication Type for the Instance Agent. Other types do not work in the demonstration environment because the instances were not set up for those authentication types.

- The window shows a second set of credentials to be defined, which are Common Database Administrator Settings. These settings can be ignored for now because they work only if all eligible databases were created with the same DB admin credentials.

- The DB admin user is not required for connection to the instance agent.

- The DB admin is a user who is defined inside the database; the instance agent user is an operating system user.

- If we do not access the database, we do not need any DB users to know. However, the DB admin is used later by SAP LaMa for several operations.

Configure SAP instances

More configuration steps might be required after an instance is detected by LaMa. Because not all information that is accessible by the instance agent is gathered automatically, check the details of each detected system or instance (see Figure 31).

Figure 31   Detected instances overview

After selecting **Edit**, the extra information can be retrieved and is completed automatically.

Some extra setup is required for database instances. We did not specify any database user until now, so we must do that here.

Depending on the installation settings, a common user and password or individual users must be configured. Both Administrator and Manager users are required, and those users must be specified for each database instance. This information is used in configuring a minimum of four users:

• SystemDB Administrator
• SystemDB Manager
• Instance Administrator
• Instance Manager

The database credentials are shown in Figure 32.

Database Configuration

Database Name:	LM1
Type:	HDB
SystemDB Manager User Name:	SYSTEM
SystemDB Manager Password:	••••••••••••••••••••••••••
SystemDB Administrator User Name:	SYSTEM
SystemDB Administrator Password:	••••••••••••••••••••••••••
Manager User Name:	SYSTEM
Manager Password:	••••••••••••••••••••••••••
Administrator User Name:	SYSTEM
Administrator Password:	••••••••••••••••••••••••••
Instance Name:	HDB01
Database Role Active:	MASTER
Database Role Configured:	MASTER
Configuration Directory:	/hana/shared/LM1/LM1/sapdbctrl-config
SAP Database Administration Installed:	☐

Figure 32   Database credentials

When a common user was specified, all user and password combinations are the same.

SAP LaMa use cases that are related to storage

This section covers the SAP LaMa use cases that are related to storage.

System Copy

With SAP Landscape Manager, you can create system copies or system clones of your SAP HANA environment at storage volume level. Although a system copy requires a different unique system ID, the system ID is not changed for the clone.

System clones are often used to create isolated testing, demonstration, or training systems while system copies are used to expand existing implementations or to create QA and Test systems.

As part of the cloning or copying, the newly created system can be isolated from the rest of the landscape to avoid conflicts with the original systems.

In this document, the storage-based System Copy of an SAP HANA system is described. The following steps are performed by SAP Landscape Management:

1. Clone SAP System.
2. Rename cloned SAP System.
3. Conduct post-copy activities.

To perform a system copy of the SAP HANA system, the following prerequisites must be met:

- Licenses, such as third-party and SAP licenses, are necessary.
- Suitable permissions are required, as described in the following web pages of the SAP Help Portal:
  – Standard Permissions
  – Standard Roles and Groups
- Other prerequisites:
  ◦ Valid Remote Function Call user
  ◦ No validation errors on the host, system, or instance, which are used for the system copy
  ◦ No alerts show up on the instance
- For a system copy of an SAP HANA system:
  ◦ The installation must be performed adaptively. For more information, see this web page of the SAP Help Portal.
  ◦ Software Provisioning Manager
  ◦ SAP HANA lifecycle management tool (`hdblcm`)
  ◦ SAP Host Agent must be able to write to the `operations.d` subfolder

  For more information about the prerequisites check, see this web page of the SAP Help Portal.

To create an end-to-end system copy of an SAP HANA System, browse to **Provisioning** in the Landscape Management system.

Select the SAP HANA system that you want to create the System Copy of in the **System and AS** tab and then, click **Copy System** (see Figure 33).
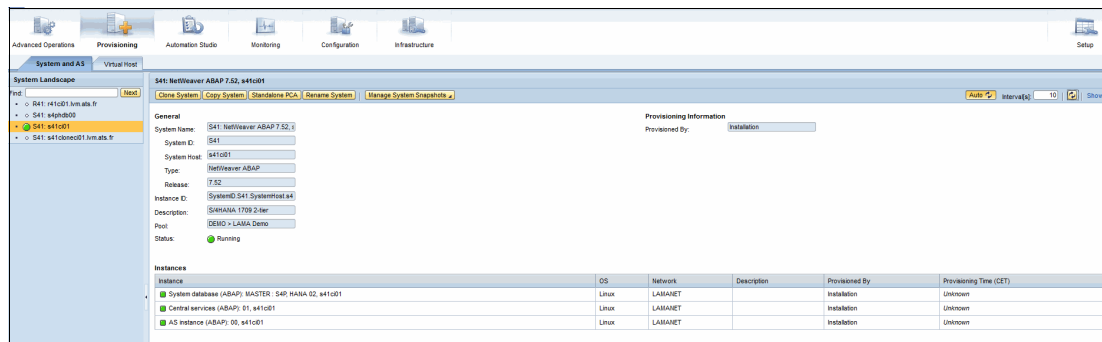
Figure 33   SAP HANA Source

43

Specify a new System ID and a new HANA SID and select the Pool from where the target system is provisioned. Enter a new master password for the target system, which is used as System password; for example, the database users and for the operating system (see Figure 34).
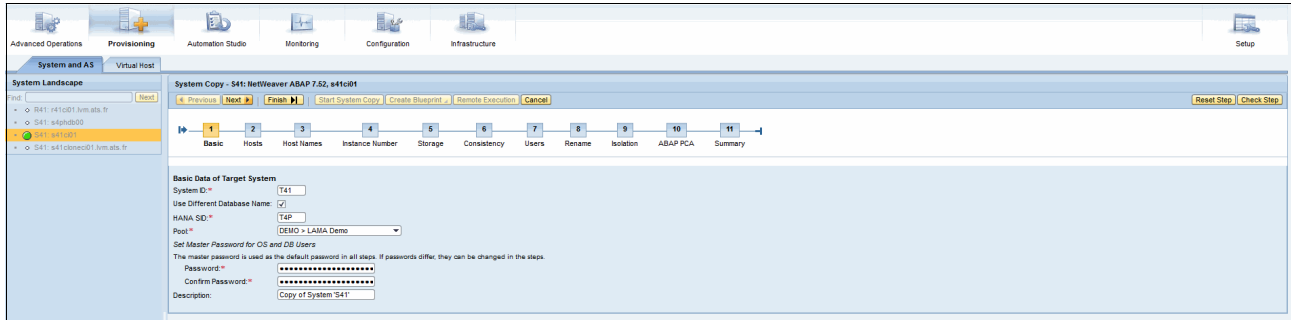
Figure 34   Basic Data of Target System

Next, select the host for the target system. You can place the system on one or several hosts.

You might want to select a host that does not have any other instances running to avoid interferences (see Figure 35, Figure 36, and Figure 37 on page 45).
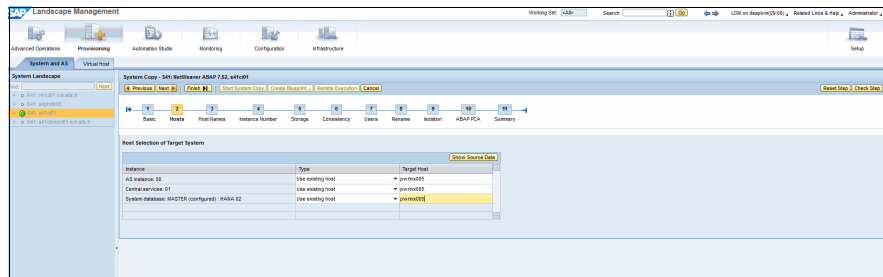
Figure 35   Instances where not many resources are required, all instances can be on one host
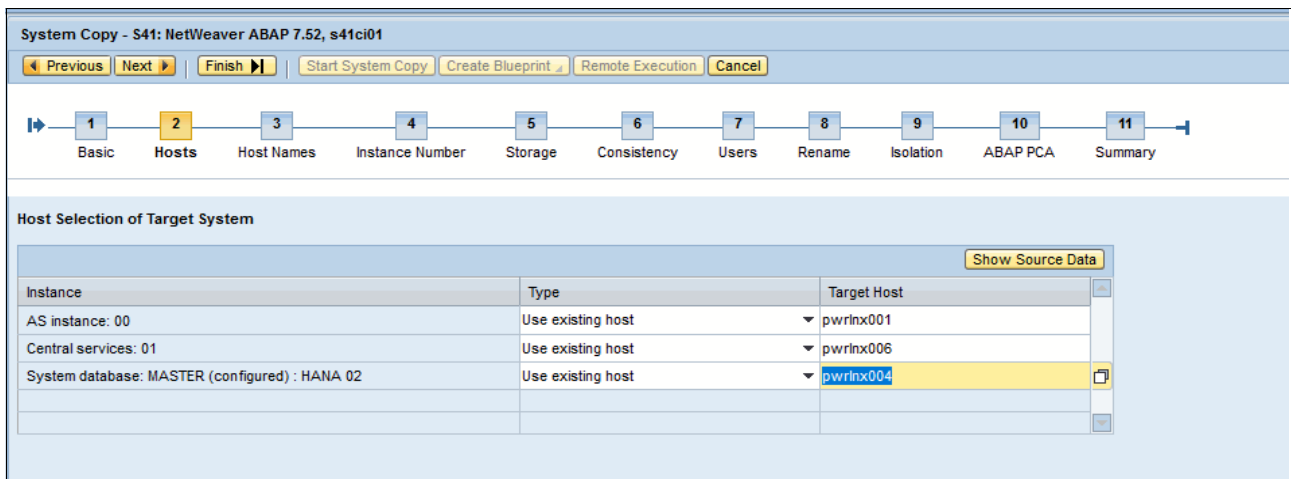
Figure 36   Application Server, central services, and database are running on individual servers for maximum performance
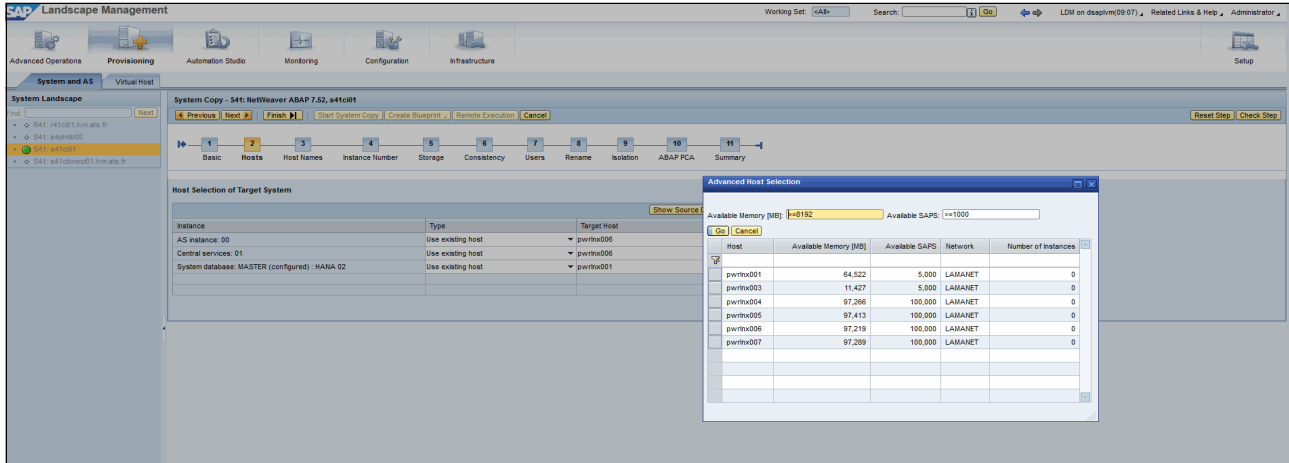
Figure 37   For high performing systems pick a host with 0 number of instances running on to avoid inferences

With SAP Landscape Management, an adaptive setup is recommended, which means that application servers, central instances, and databases get their own virtual host names to disconnect the applications from the physical host. It allows the use of virtual IP addresses when applications are moved from one host to another. SAP Landscape Management configures the necessary DNS entries. For systems that are not adaptively installed, it is not possible to change the host name (see Figure 38).
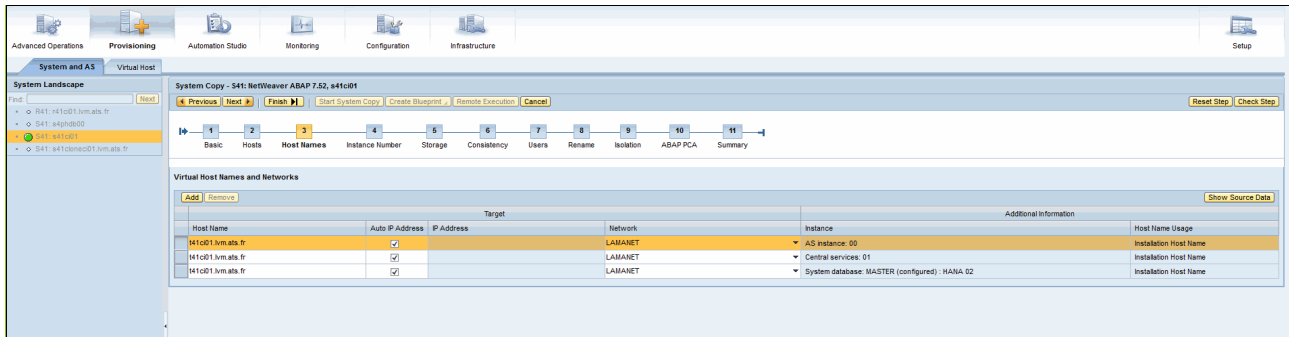
Figure 38   Virtual host names

Next, the instance number can be changed, which might be necessary when several QA or development systems are consolidated on a single host (see Figure 39).
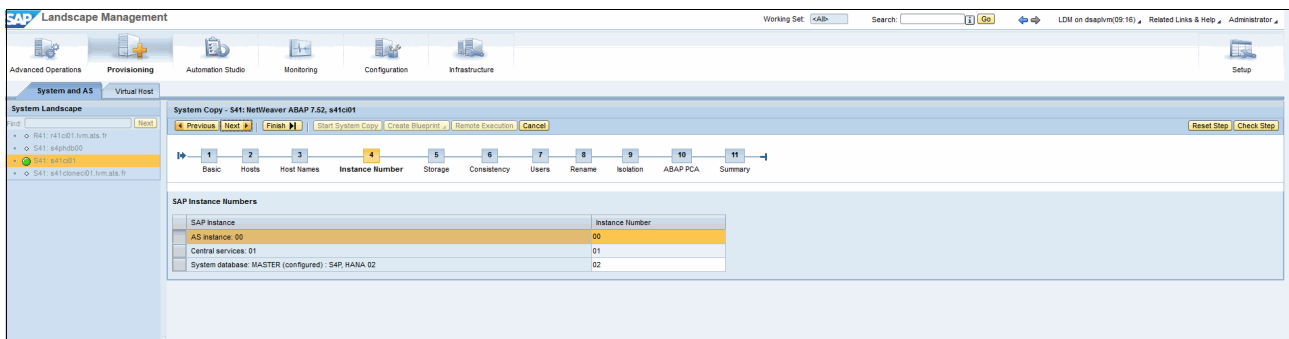
Figure 39   Changing the instance number

For a storage-based system copy, the complete data is copied from one volume to another volume by using the Storage adapters. The needed volumes can be selected (see Figure 40).
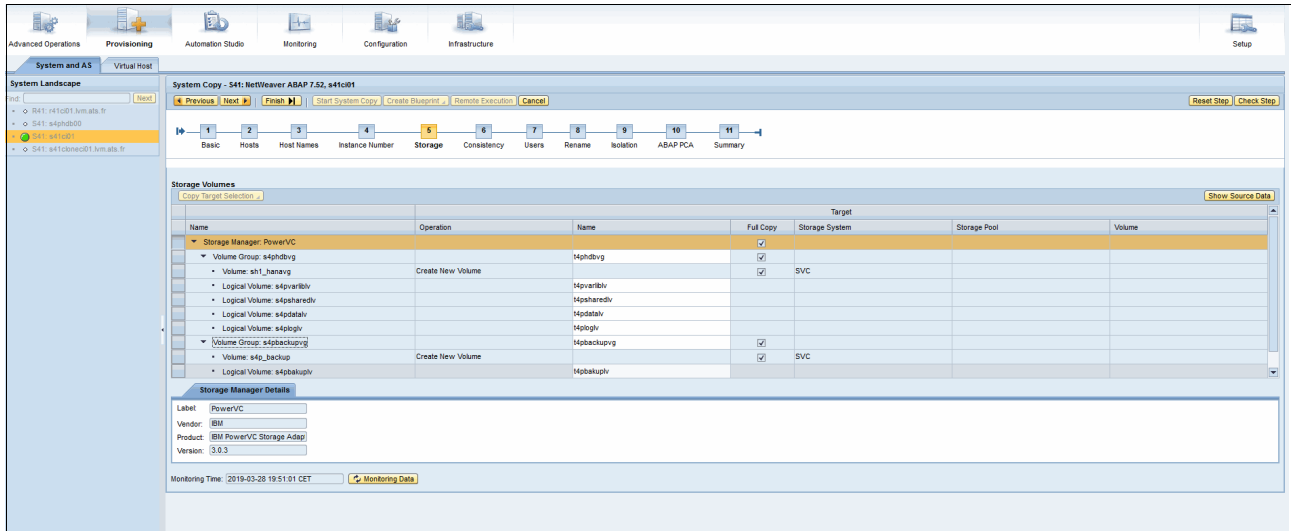
Figure 40   Storage Based System Copy

To create a consistent copy, the database must be prepared; for example, all data must be written to the disk first. This process can be done online or offline and the execution can be scheduled (see Figure 41).
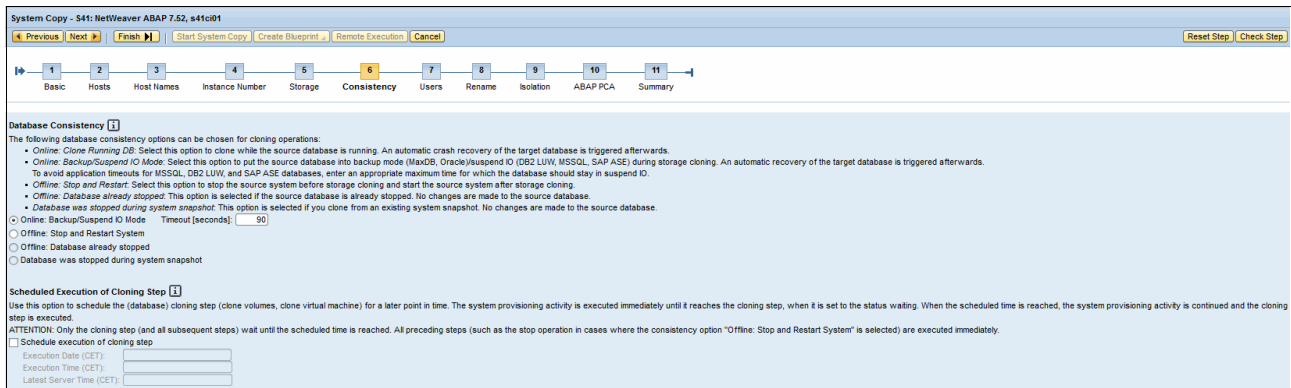
Figure 41   Options to create a consistent database

Figure 42 shows an example of how to schedule the copy creation.
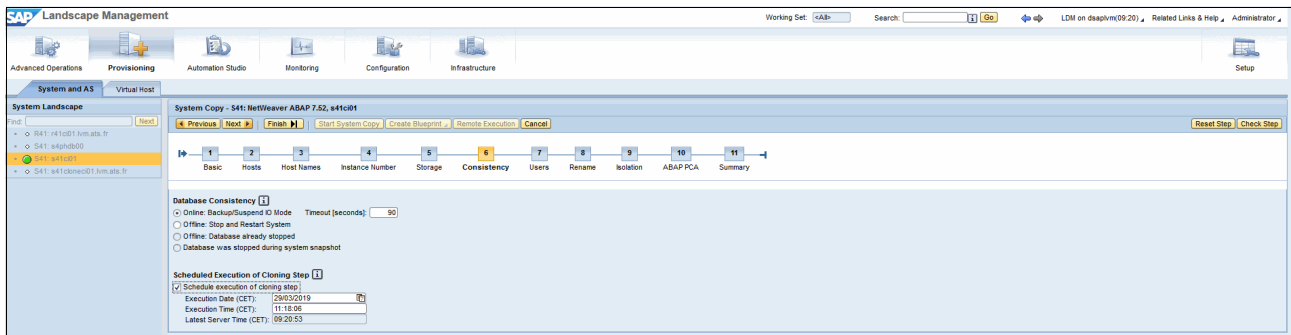
Figure 42   Schedule Database Copies

In the **Users** tab, you can change or create users for the new system. If you configured your user management system in SAP Landscape Management, SAP Landscape Management creates users for you and assigns them to the correct groups (see Figure 43).
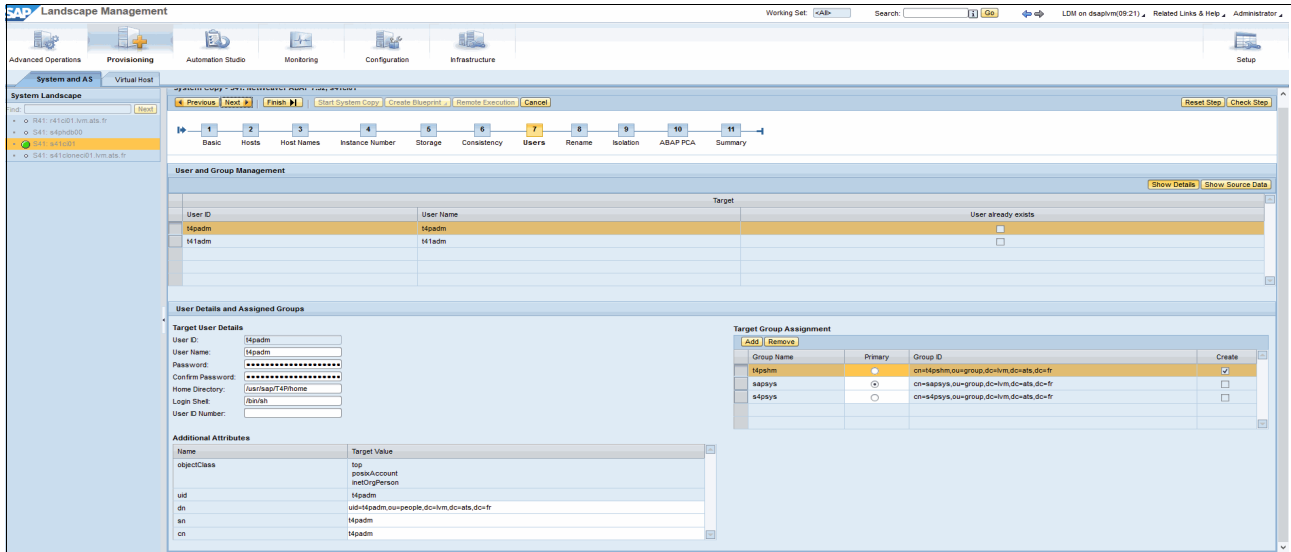
Figure 43   Configure Users

The rename functionality is using the Software Provisioning Manager to change users and passwords and file system names for the target system (see Figure 44).
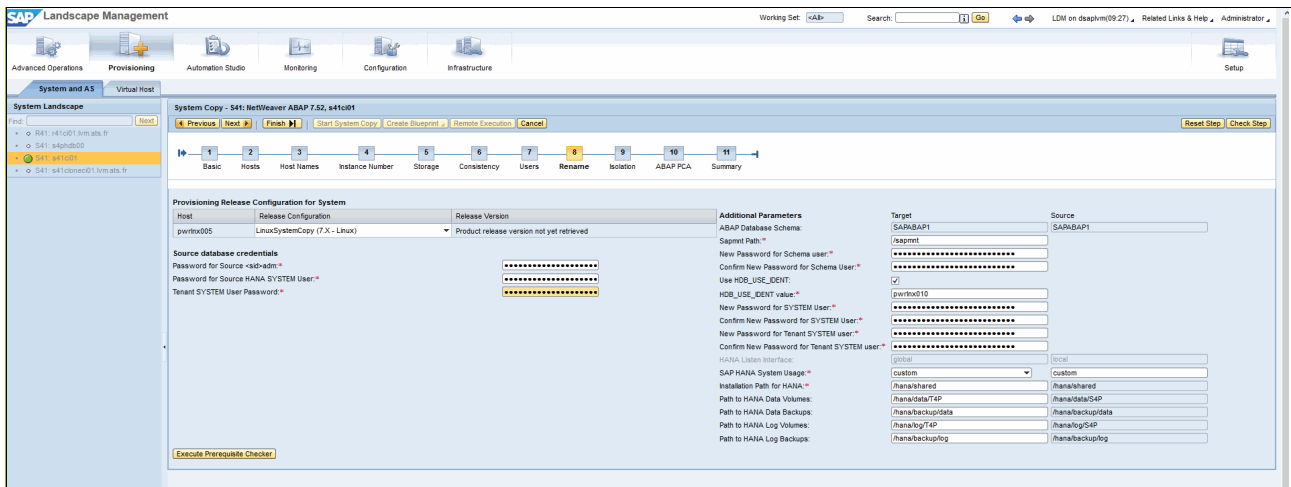
Figure 44   Rename tab

To see whether the system is ready for system copy, run the prerequisite checker, which checks all instances simultaneously.

To avoid interferences of the copied system with the environment in the SAP landscape, the target system must be shielded. Therefore, SAP Landscape Management uses the firewall of the operating system to avoid any unwanted communication. In the Isolation tab, you can select ports to open necessary communication with the new system, but block all outgoing communication (see Figure 45 on page 48).
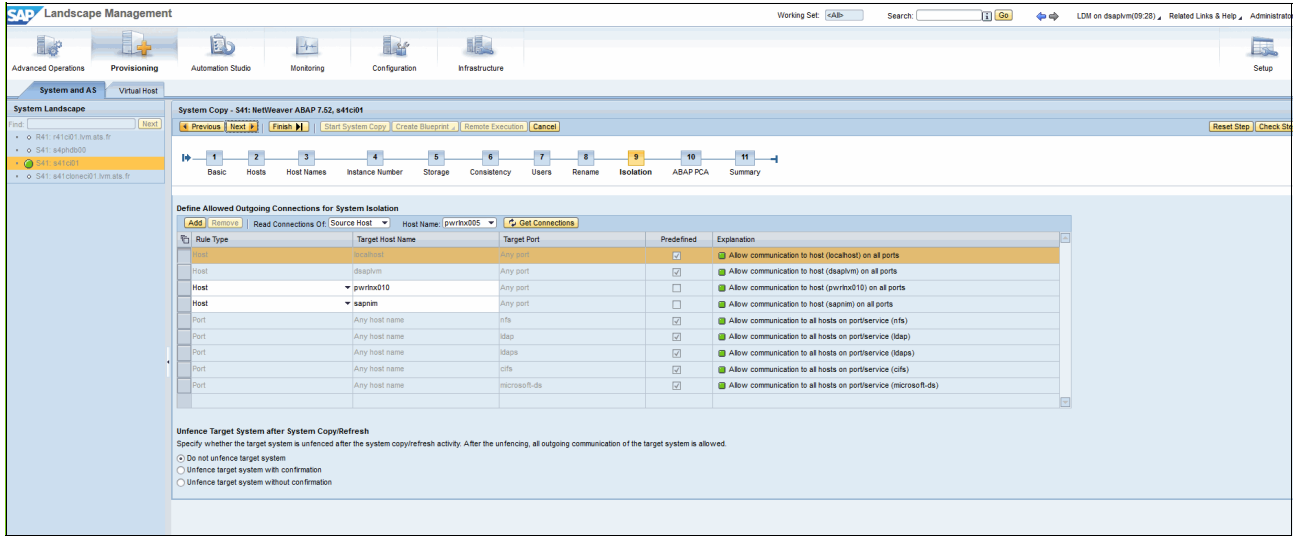
Figure 45   Configure Isolation

To make all the necessary changes inside the target system, a default target list is provided at the source system and is run at the target system after start in the Post Copy Automation (PCA). The task list can be modified to take care of environment-specific changes.

You also can specify that more than one task list is run sequentially in the order that they are listed. However, the default task list is run on client 000 first (see Figure 46, Figure 47, and Figure 48 on page 49).

Here, you also can change the passwords for all configured remote function call (RFC) destinations to the master password.
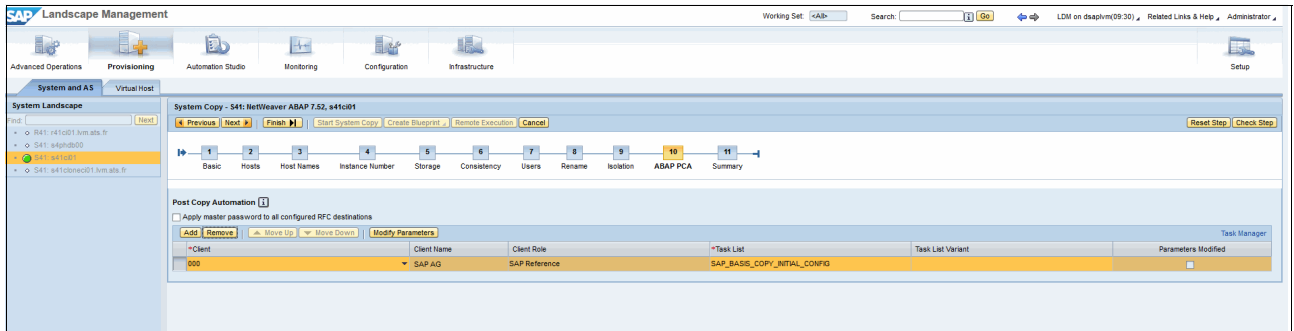
Figure 46   One PCA task list
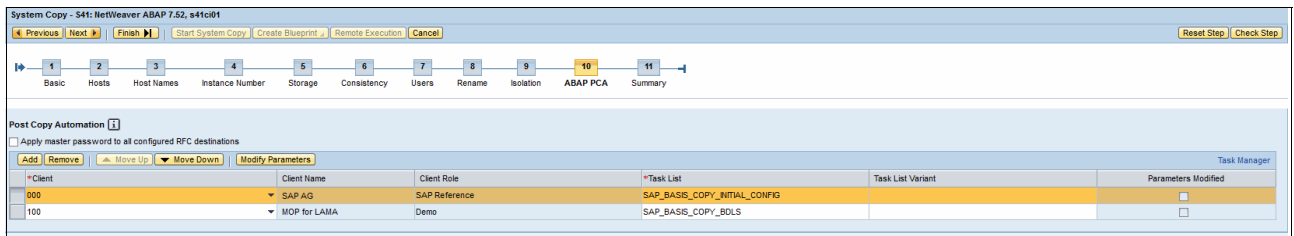
Figure 47   Multiple PCA task lists
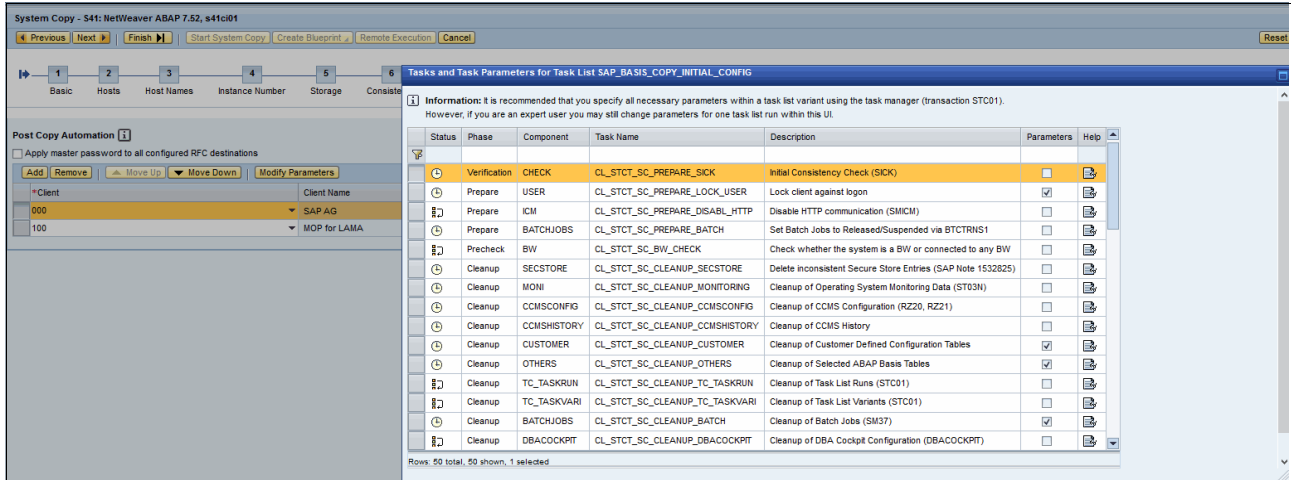
Figure 48   Individual tasks of PCA task list

In the Summary tab, all entered parameters are displayed; for example, host names and instance numbers. After you check the information, the System copy can be started (see Figure 49).
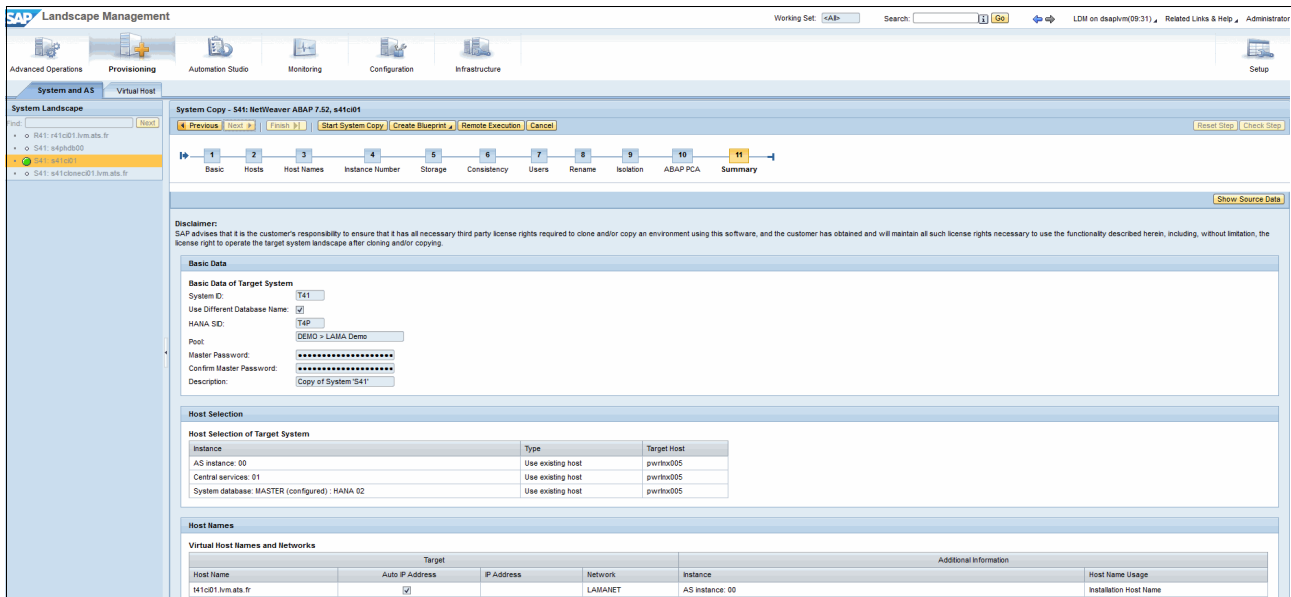
Figure 49   Summary Tab

The execution of all of the System Copy tasks with all individual steps can be monitored in the Monitoring view (see Figure 50 on page 50).

Figure 50   Monitoring View

Select the task in the top window to see the progress of the individual steps in the bottom window, as shown in Figure 51 and Figure 52.

Figure 51   Individual Steps of System Copy Task

Figure 52   Successful completion of the individual Steps

After all of the steps are completed successfully, the System Copy Process Status changes to Completed in the Overview window of the Monitoring tab (see Figure 53).

Figure 53   Successful Completion of the System Copy Process

System refresh

Refreshing a System copy (for example, QA or Test) with data means overwriting the target system with the latest data from a source system while maintaining the configuration.

For refreshing the system, the following tasks are performed by SAP Landscape Management (see Figure 54 on page 52):

• Export Configuration

• Stop

• Unprepare:
  – Unregister database
  – Detaches DB from host
  – Unbind virtual IP addresses
  – Unmounts file systems
  – Unprepare IP addresses

• Clone

• Prepare

• Activate Isolation

• System-rename

• Start

• Import configuration + PCA

• Deactivate Isolation

The Refresh option can be found by using the Provisioning tab of Lama.
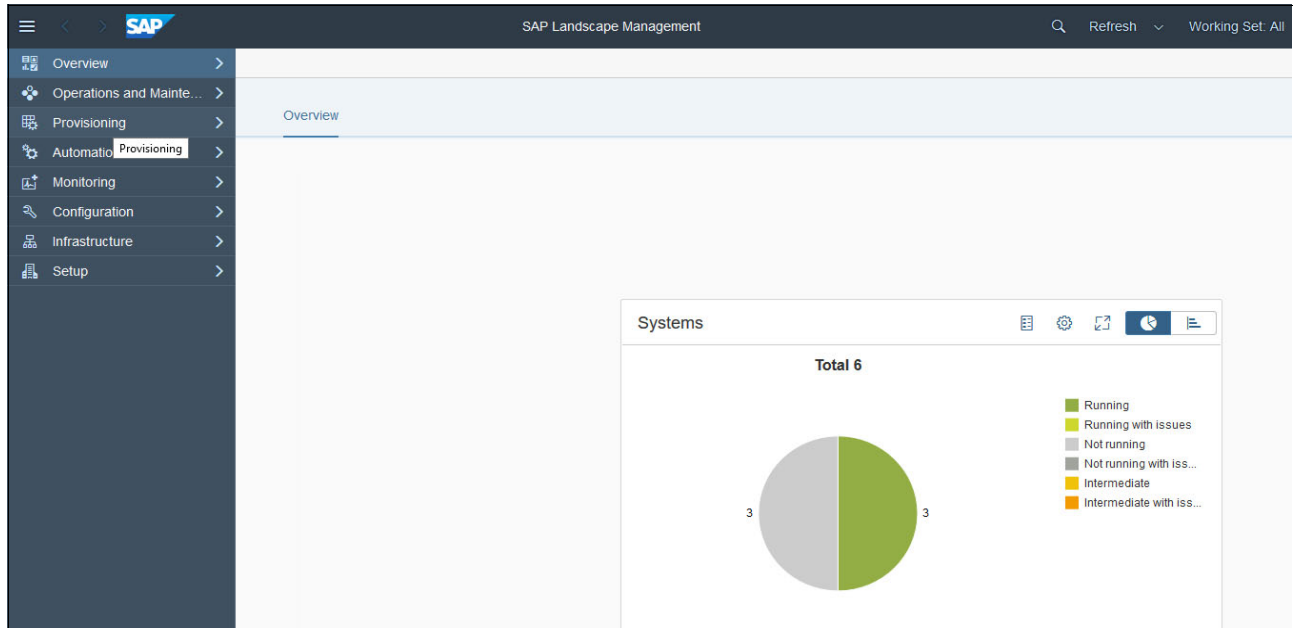
Figure 54   SAP Landscape Management Overview window

In the following example, T41 is a system copy of S41 and is used for QA. To update (for example) a previously copied or cloned system T41 with all the changes of S41, the refresh system functionality can be used (see Figure 55).
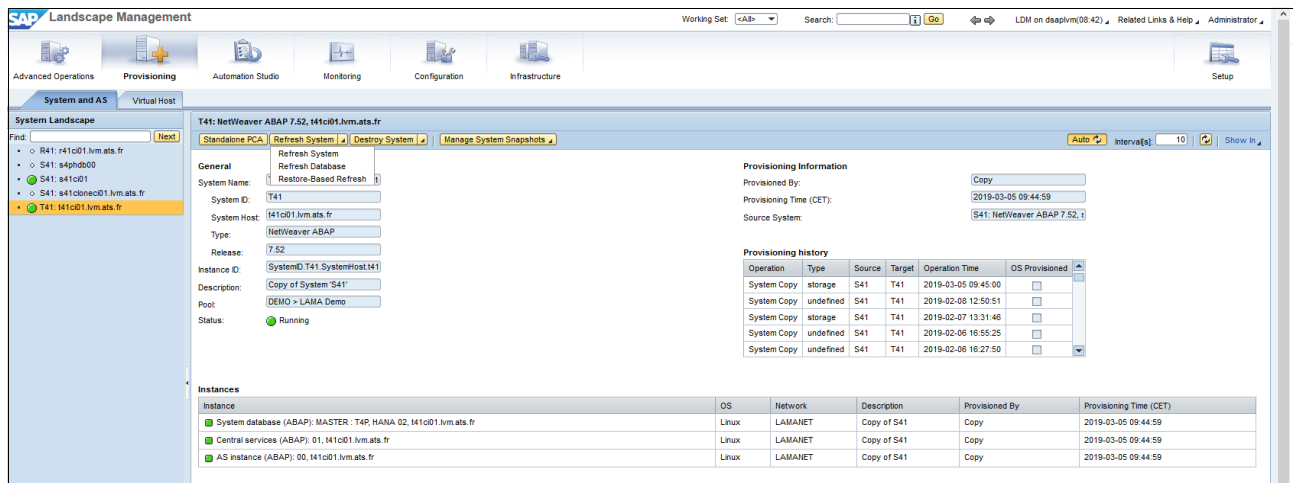

Figure 55   Configure System Refresh

The following options are available:

- Refresh System: Refreshes the complete system, including central instance, application server, and database.

- Refresh database: Refreshes only the database. Configuration settings of the application instance and central instance are kept.

- Restore-refresh Database: For SAP HANA databases only.

Refresh is done by restoring a backup of the source system (for example, the production system) into the target system (for example, the QA System). The configuration settings of application instance and central instance are kept.

Refresh System

During the Refresh System process, the target system is overwritten by the Source Target data.

To restore the target configuration settings after the system refresh, they must be exported first. If the configuration settings exist, they can be used by specifying the directory:

- Step 1: Basic (see Figure 56)

Figure 56   Figure 51:Basic Configuration

- Step 2: Hosts (see Figure 57)

    In the next step, the source and target hosts can be selected. Because this refresh is a system refresh, SAP Landscape Management completes the current target system information, but you can define a new target host that gets you a new refreshed system.

Figure 57   Host Configuration

- Step 3: Host names (see Figure 58 on page 54)

    The IP addresses and virtual host names are displayed in the Hostnames tab.
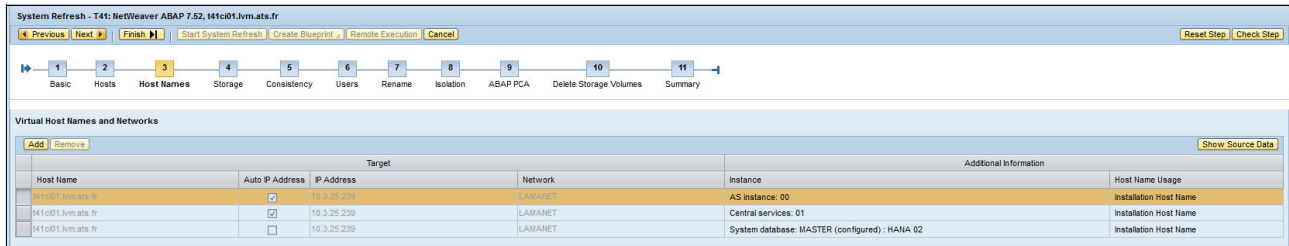
Figure 58   Virtual Hostnames Configuration

- Step 4: Storage

    In the Storage tab, the storage can be selected where the data is copied to and if a snapshot is sufficient or if a full copy is necessary (see Figure 59).
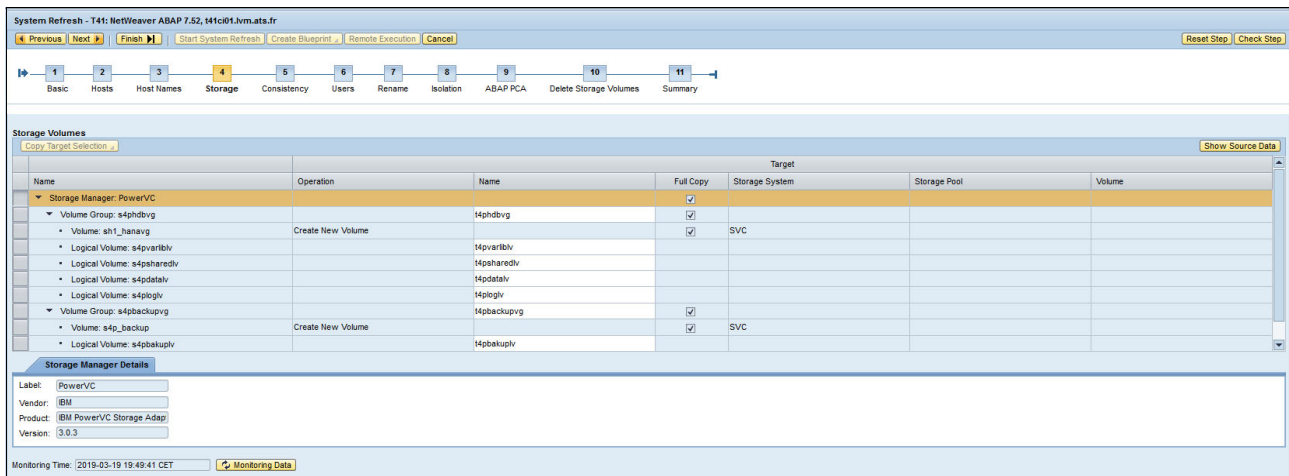
Figure 59   Storage Configuration

- Step 5: Consistency

    The consistency options can be selected in the Consistency tab. For production environments, an online copy might be necessary; however, for this step, the database must be put in a consistent state (Suspend I/O Mode), which ensures that all data is written to the disk (see Figure 60).

Figure 60   Consistency Configuration

- Step 6: Users

  Users can be selected in the User's tab and passwords can be changed (see Figure 61).

Figure 61   Users Configuration

- Step 7: Rename

  The system rename can be configured in the Rename tab; for example, the directories to use (see Figure 44 on page 47). If any warnings occur that indicate that directories exist, they can be ignored because this issue is expected for a refresh.
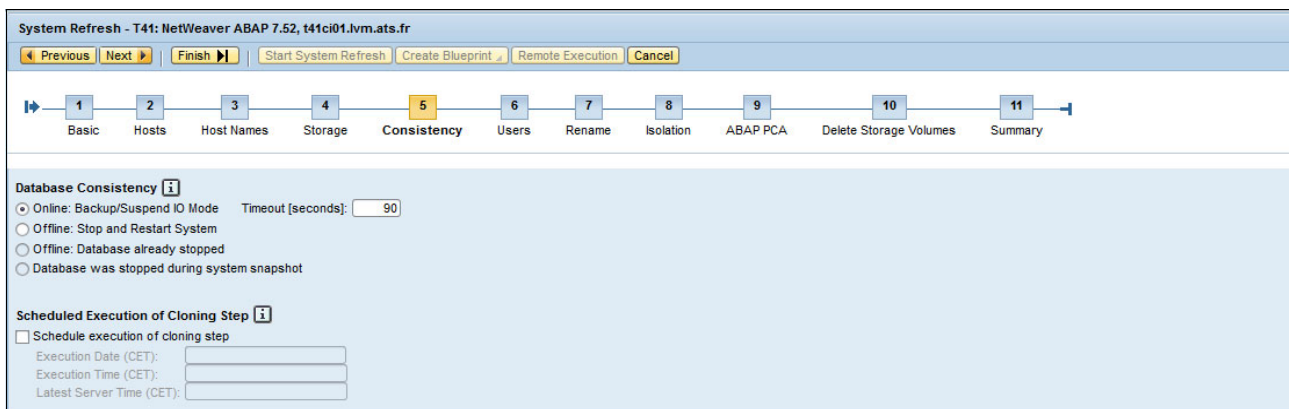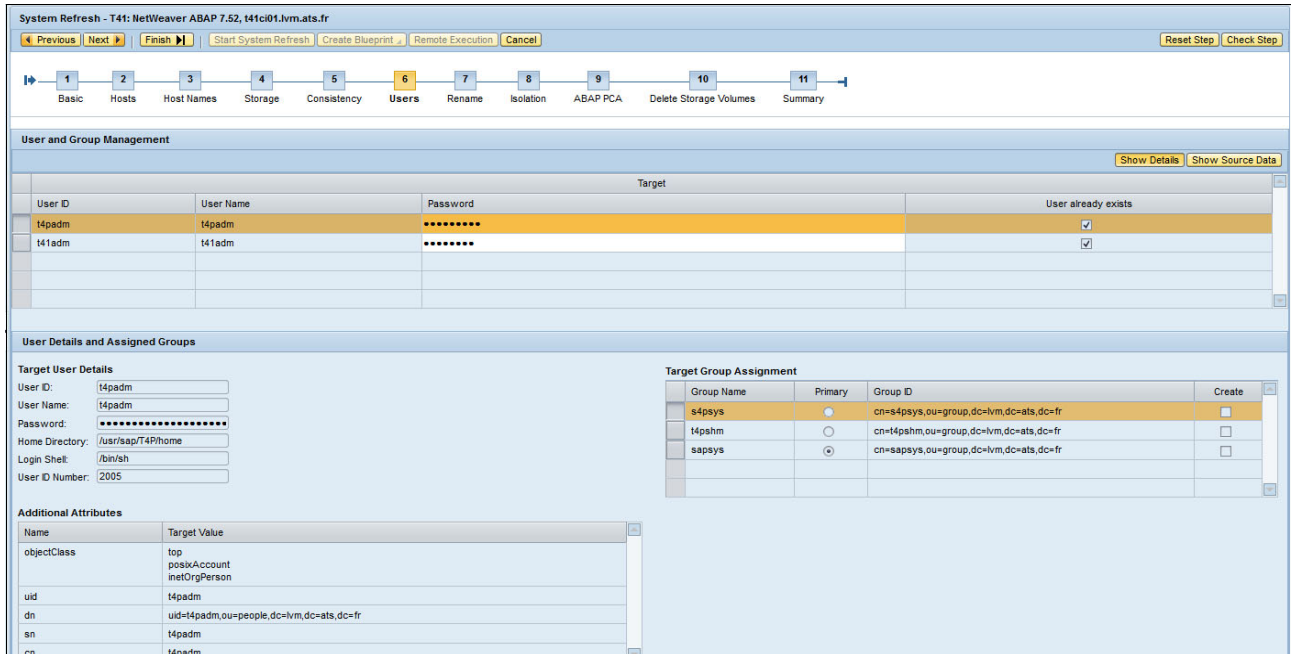
- Step 8: Isolation

  After the copy, the system must be isolated to not communicate with other systems in the landscape (see Figure 45 on page 48).

- Step 9: ABAP PCA

  The ABAP Post copy automation (PCA) process makes all necessary changes in the target system; for example, imports the exported configuration settings, again by way of a task list (see Figure 45 on page 48 or Figure 47 on page 48).

- Step 10: Delete Storage Volumes

  Configure if you want to delete the original storage volumes or to keep them.

- Step 11: Summary

  After checking the system refresh configuration, the Start System Refresh can be started.

In the Monitoring window, the System Refresh Task is listed with the individual tasks in the window.

Relocate

This section describes the relocation of an instance or a set of instances from one host to another physically or virtually. Instances can be relocated in running status or not running status. The operation retains the status of the instances. Not running instances are unprepared and prepared.

Prerequisite

If you relocate SAP instances that are connected to an SAP HANA database, ensure that you installed the `hdbclient` with the `-H <virtual-hostname>` option. If this is not the case, the `hdbuserstore` entry is generated only for the current physical host. Check the `<hdbclient-dir>/install/installation.ini` file. This file must contain an `HOSTNAME=<virtual-hostname>` entry.

For more information, see Secure User Store (hdbuserstore) in the SAP HANA Security Guide.

Notes

SAP Landscape Management cannot relocate SAP HANA tenant databases.

An SAP HANA System can be relocated only to a target system that runs on the same hardware platform as the source system. Complete the following steps:

1. From the operations and maintenance window, select the SAP system that you want to relocate and use the drop-down arrow that is next to Operations to select the relocate functionality (see Figure 62).
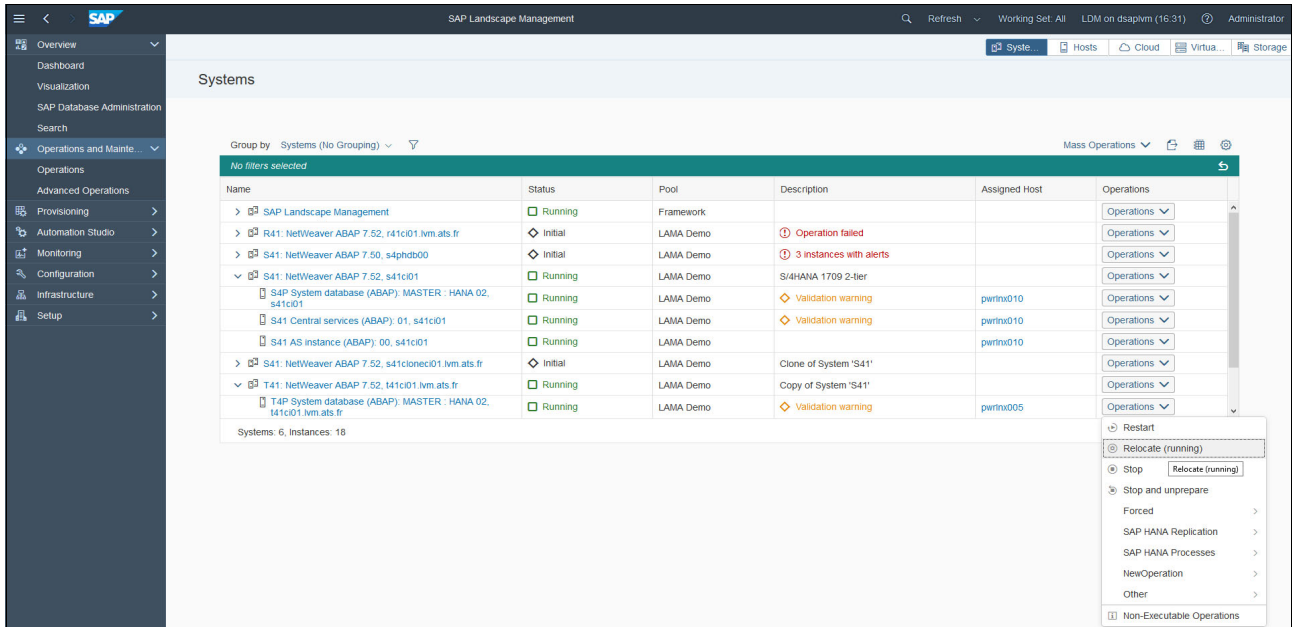
Figure 62   Operations and Maintenance window

2. In the Setup execution window, select the target host (see Figure 63).
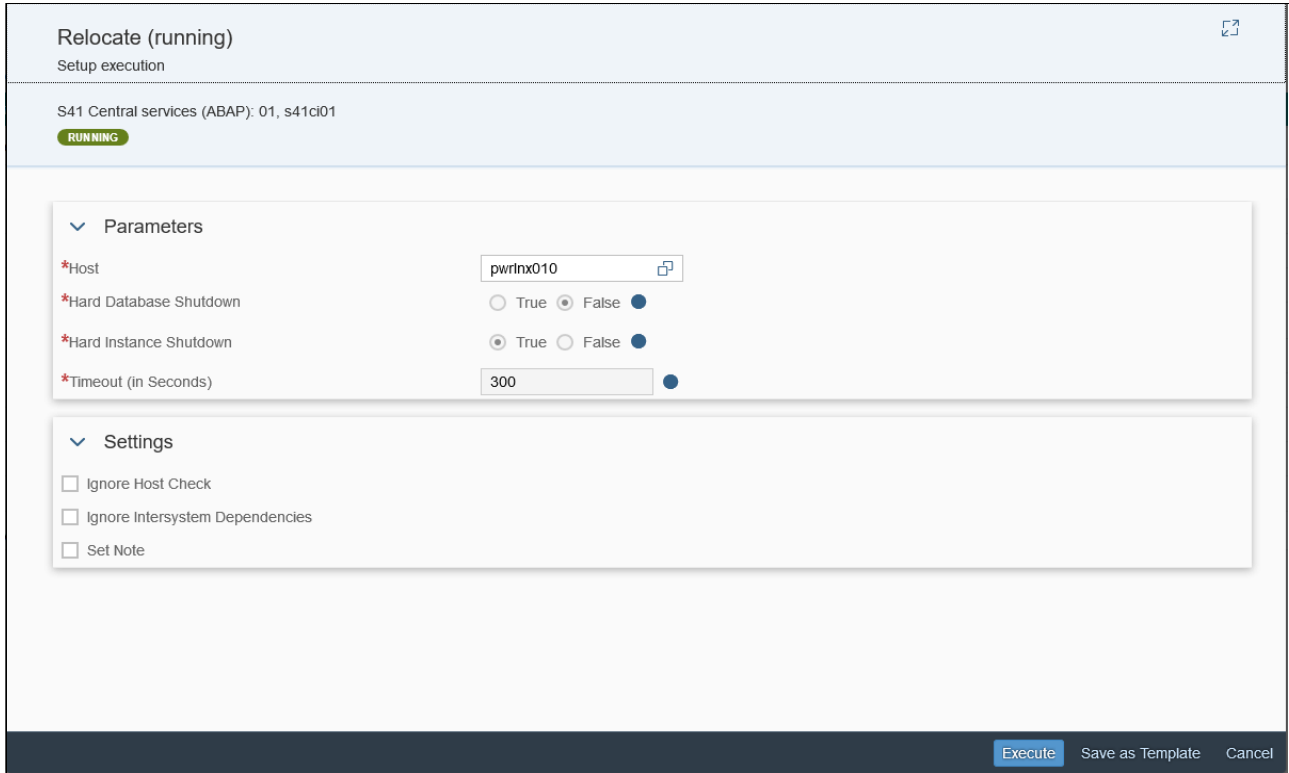
Figure 63   Configure Host

3. Hosts can be selected from the Allowed values window that opens. Click the icon that is next to the host (see Figure 64).

Figure 64   Available Hosts

If you select **Hard Database Shutdown** or **Hard Instance Shutdown**, the database or instance stops immediately, even if clients are connected.

Specifying Hard Instance Shutdown as "False" allows you to specify a timeout if the instance supports a soft shutdown mode.

Selecting **Ignore Host Check** or **Ignore Intersystem Dependencies** is not recommended and should be used with care.

4. Click **Execute** and check the progress of the individual tasks in the monitoring window.

For more information, see the User Guide SAP Landscape Management 3.0 Enterprise Edition on the SAP Help Portal.

Appendix

To ensure that SAP applications run smoothly, SAP offers the tailored data center integration (TDI) approach to hardware selection. This approach allows customers to maintain hardware flexibility with sensitivity to budget.

For more information about all storage systems that are certified for SAP HANA production, see the Certified and Supported SAP HANA Hardware web page (log in required).

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.

59

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 FlashCore®	POWER9™
DB2®	IBM FlashSystem®	Redbooks (logo) ®
Easy Tier®	IBM Spectrum®	System Storage™
FlashCopy®	IBM Spectrum Storage™	Tivoli®
IBM®	POWER®	

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.

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.

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

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

VMware, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

IBM®