

SUSE Manager 4.3

Retail Guide

January 18 2024



Table of Contents

| | |
|---|----|
| Retail Guide Overview | 1 |
| 1. Components | 2 |
| 1.1. The SUSE Manager Server | 2 |
| 1.2. Build Hosts | 2 |
| 1.3. Branch Servers | 2 |
| 1.4. Point-of-Service Terminals | 2 |
| 1.5. Fitting It All Together | 3 |
| 1.5.1. Hardware Types | 3 |
| 1.5.2. Branch System Groups | 3 |
| 1.5.3. Salt Formulas | 3 |
| 1.5.4. Saltboot | 4 |
| 2. Retail Requirements | 5 |
| 2.1. Server Requirements | 5 |
| 2.2. Branch Server Requirements | 5 |
| 2.3. Build Host Requirements | 6 |
| 2.4. Network Requirements | 6 |
| 2.5. POS Terminal Requirements | 6 |
| 2.5.1. UEFI Secure Booting Requirements | 7 |
| 3. Network Architecture | 8 |
| 3.1. Branch Server Network Configuration | 8 |
| 3.1.1. Dedicated Network Architecture | 9 |
| 3.1.2. External Network Architecture | 9 |
| 3.1.3. Shared Network Architecture | 10 |
| 4. Retail Installation | 11 |
| 4.1. Install with the Unified Installer | 11 |
| 4.1.1. Install SUSE Manager for Retail | 11 |
| 4.1.2. Install SUSE Manager for Retail Branch Server | 13 |
| 4.1.3. Install SUSE Manager for Retail Build Host | 16 |
| 4.2. Set Up the SUSE Manager for Retail Environment | 16 |
| 4.2.1. Prepare and Build Terminal Images | 17 |
| 4.2.2. Branch identification and architecture topology | 18 |
| 4.2.3. Required System Groups | 18 |
| 4.2.4. Configure Services for Saltboot | 19 |
| 4.2.5. Synchronize Images to the Branch Server | 22 |
| 5. Deploying Terminals | 23 |
| 5.1. Deployment basics | 23 |
| 5.1.1. Create a Hardware Type Group | 24 |
| 5.1.2. Assign and Configure the Saltboot Formula for Each Hardware Type Group | 24 |
| 5.1.3. Synchronize Images to the Branch Server | 25 |

| | |
|--|-----------|
| 5.1.4. Deploy Images to the Terminals | 26 |
| 5.1.5. Customize the Terminal Image Download Process | 26 |
| 5.2. Deploy Terminals – Other Methods..... | 27 |
| 5.2.1. Deploy Terminals with a Removable USB Device | 28 |
| 5.2.2. Deploy Terminals over a Wireless Network | 29 |
| 5.3. Deploy Terminals and Auto-Accept Keys | 32 |
| 5.3.1. Configure Saltboot to Send Auto-Signed Grain Once..... | 33 |
| 5.3.2. Configure Saltboot to Keep Auto-Signed Grains..... | 33 |
| 5.3.3. Configure Saltboot to Auto-Sign During PXE Boot | 34 |
| 5.3.4. Configure the Server to Auto-Accept..... | 35 |
| 5.4. Forced Saltboot image redeployment..... | 35 |
| 5.4.1. Force Saltboot redeployment using Salt grains | 36 |
| 5.4.2. Force Saltboot redeployment using custom info values..... | 36 |
| 5.4.3. Force Saltboot redeployment using Saltboot API call | 38 |
| 5.4.4. Force Saltboot redeployment by custom pillar | 39 |
| 5.5. Terminal Boot Process (Saltboot Diagram) | 40 |
| 5.6. Terminal Names | 42 |
| 5.6.1. Naming by HWType | 42 |
| 5.6.2. Naming by Hostname | 42 |
| 5.6.3. Naming by FQDN | 43 |
| 5.6.4. Assign Hostnames to Terminals | 44 |
| 5.7. Offline Use | 45 |
| 5.7.1. Offline Terminal Reboot..... | 45 |
| 5.7.2. Cached Terminal Updates..... | 46 |
| 5.8. Rate Limiting Terminals | 47 |
| 5.8.1. Troubleshooting | 47 |
| 6. Introduction to Retail Formulas | 48 |
| 6.1. Branch Server Formulas | 48 |
| 6.2. Partitioning and Image Deployment Formula | 49 |
| 7. Image Pillars | 50 |
| 8. Administration | 53 |
| 8.1. Mass Configuration | 53 |
| 8.1.1. Branch Server Mass Configuration | 53 |
| 8.1.2. Terminal Mass Configuration | 54 |
| 8.1.3. Export Configuration to Mass Configuration File..... | 55 |
| 8.2. Example YAML File for Mass Configuration | 55 |
| 8.3. Delta Images | 58 |
| 8.3.1. Building Delta Images | 58 |
| 8.3.2. Tuning Delta Generation | 59 |
| 8.3.3. Image Synchronizing to the Branch Server..... | 59 |
| 8.4. Network Administration..... | 59 |
| 9. Retail Migration | 61 |
| 9.1. Before You Migrate | 61 |

| | |
|---|-----------|
| 9.1.1. Prepare to Migrate from SUSE Linux Enterprise Point of Service | 61 |
| 9.2. Migrate SUSE Linux Enterprise Point of Service 11 to SUSE Manager for Retail. | 62 |
| 9.2.1. Migration with Complete Data Dump | 62 |
| 9.2.2. Migration with Branch by Branch Data Dump | 63 |
| 9.2.3. Converting XML to YAML | 65 |
| 9.3. Upgrade SUSE Manager for Retail Branch Server | 66 |
| 10. Example configurations | 67 |
| 10.1. Set Up the SUSE Manager for Retail Environment with dedicated network for terminal | 67 |
| 10.1.1. Assumptions | 67 |
| 10.1.2. Create required system groups | 68 |
| 10.1.3. Assign and configure branch server formulas | 68 |
| 10.1.4. Setup partitioning | 72 |
| 10.1.5. Synchronize images | 73 |
| 10.2. Set Up the SUSE Manager for Retail Environment with dedicated network for terminal using bundled scripts | 73 |
| 10.2.1. Assumptions | 73 |
| 10.2.2. Quick set up of SUSE Manager for Retail using for Retail tools | 74 |
| 10.2.3. Setup partitioning | 75 |
| 10.2.4. Synchronize images | 77 |
| 10.3. Set Up the SUSE Manager for Retail Environment with shared network between terminals, branch servers and SUSE Manager | 77 |
| 10.3.1. Assumptions | 77 |
| 10.3.2. Create required system groups | 78 |
| 10.3.3. Assign and configure formulas | 78 |
| 10.3.4. Synchronize images | 81 |
| 10.4. Set Up the SUSE Manager for Retail Environment using containerized proxy ... | 81 |
| 10.4.1. Assumptions | 82 |
| 10.4.2. Create required system groups | 82 |
| 10.4.3. Saltboot group | 82 |
| 10.4.4. Comparing containerized and non-containerized workflows | 84 |
| 10.4.5. Validating Saltboot group configuration | 85 |
| 11. Best practices | 87 |
| 11.1. Deploying new images | 87 |
| 11.1.1. Controlling image versions | 88 |
| 12. What Next? | 90 |
| 12.1. More Documentation | 90 |
| 12.2. Support | 90 |
| 13. GNU Free Documentation License | 91 |

Retail Guide Overview

Updated: 2024-01-18

SUSE Manager for Retail 4.3 is an open source infrastructure management solution, optimized and tailored specifically for the retail industry. It uses the same technology as SUSE Manager, but is customized to address the needs of retail organizations.

SUSE Manager for Retail is designed for use in retail situations where customers can use point-of-service terminals to purchase or exchange goods, take part in promotions, or collect loyalty points. In addition to retail installations, it can also be used for other purposes, such as maintaining student computers in an educational environment, or self-service kiosks in banks or hospitals.

SUSE Manager for Retail is intended for use in installations that include servers, workstations, point-of-service terminals, and other devices. It allows administrators to install, configure, and update the software on their servers, and manage the deployment and provisioning of point-of-service machines.

This guide provides an overview of SUSE Manager for Retail, and its initial installation and setup.

It should be read in conjunction with the SUSE Manager documentation suite, available from <https://documentation.suse.com/suma>.

For more information about managing your SUSE Manager for Retail environment, or to find out where to get help, see [Retail](#) › [Retail-next](#).

Chapter 1. Components

SUSE Manager for Retail is made up of various components. For more on how these components work together, see [Retail › Retail-network-arch](#).

1.1. The SUSE Manager Server

The SUSE Manager server contains information about infrastructure, network topology, and everything required to automate image deployment and perform day-to-day operations on branches and terminals. This can include database entries of registered systems, Salt pillar data for images, image assignments, partitioning, network setup, network services, and more.

1.2. Build Hosts

Build hosts can be arbitrary servers or virtual machines. They are used to securely build operating system images.

For more information on build hosts, see [Administration › Image-management](#).

1.3. Branch Servers

Branch servers should be physically located close to point-of-service terminals, such as in an individual store or branch office. Branch servers provide services for PXE boot, and act as an image cache, Salt broker, and proxy for software components (RPM packages). The branch server can also manage local networking, and provide DHCP and DNS services.



- For monitoring, you can install Prometheus server on the Branch servers.
- For visualization and analysis, you can install Grafana with multiple data sources on a dedicated host. For more information about Prometheus and Grafana, see [Administration › Monitoring](#).

1.4. Point-of-Service Terminals

Point-of-Service (POS) terminals can come in many different formats, such as point-of-sale terminals, kiosks, digital scales, self-service systems, and reverse-vending systems. Every terminal, however, is provided by a vendor, who set basic information about the device in the firmware. SUSE Manager for Retail accesses this vendor information to determine how best to

work with the terminal in use.

In most cases, different terminals will require a different operating system (OS) image to ensure they work correctly. For example, an information kiosk has a high-resolution touchscreen, where a cashier terminal might only have a very basic display. While both of these terminals require similar processing and network functionality, they will require different OS images. The OS images ensure that the different display mechanisms work correctly.

SUSE Manager for Retail supports POS terminals that boot using both BIOS and UEFI. For UEFI booting terminals, you will need to configure the EFI partition in the Saltboot formula. For more information on EFI in the Saltboot formula, see [Specialized-guides › Salt](#).

1.5. Fitting It All Together

SUSE Manager for Retail uses the same technology as SUSE Manager, but is customized to address the needs of retail organizations.

1.5.1. Hardware Types

Because every environment is different, and can contain many different configurations of many different terminals, SUSE Manager for Retail uses hardware types to simplify device management.

Hardware types allow you to group devices according to manufacturer and device name. Then all devices of a particular type can be managed as one.

1.5.2. Branch System Groups

SUSE Manager for Retail uses system groups to organize the various devices in your environment.

Each branch requires a system group, containing a single branch server, and the POS terminals associated with that server. Each system group is identified with a Branch ID. The Branch ID is used in formulas and scripts to automatically update the entire group.

1.5.3. Salt Formulas

SUSE Manager for Retail uses Salt formulas to help simplify configuration. Formulas are pre-written Salt states, that are used to configure your installation.

You can use formulas to apply configuration patterns to your hardware groups. SUSE Manager for

Retail uses the Saltboot formula, which defines partitioning and OS images for terminals.

You can use default settings for formulas, or edit them to make them more specific to your environment.

For more information about formulas, see [Retail › Retail-formulas-intro](#).

1.5.4. Saltboot

Saltboot is a collection of tools and processes that are used to bootstrap, deploy and validate SUSE Manager for Retail terminals.

Saltboot consists of:

- Initialization:

The Saltboot **initrd** is created during image building and is required for bootstrapping terminals.

- Saltboot state:

The Salt state that contains the logic for the entire Saltboot process.

- Partitioning pillars:

The Salt pillar structure that describes how terminals are partitioned and what image is deployed on each terminal.

- Images and boot images pillars:

When the image building feature in SUSE Manager successfully builds an image that contains the Saltboot **initrd**, the image and boot image Salt pillars are created.

The Saltboot process involves the SUSE Manager Server, a terminal running the saltboot **initrd**, and the branch server providing the Saltboot services to the terminal.

For a detailed diagram explaining how the Saltboot boot process works, see [Retail › Retail-saltboot-diagram](#).

Chapter 2. Retail Requirements

Before you install SUSE Manager for Retail, ensure your environment meets the minimum requirements. This section lists the requirements for the SUSE Manager for Retail installation. These requirements are in addition to the SUSE Manager requirements listed at [Installation-and-upgrade › General-requirements](#).



SUSE Manager for Retail is only supported on the x86-64 architecture.

2.1. Server Requirements

Table 1. Hardware Requirements for SUSE Manager Server

| Hardware | Recommended |
|-------------|--|
| CPU | Minimum 4 dedicated 64-bit CPU cores |
| RAM: | Test Server Minimum 8 GB |
| | Base Installation Minimum 16 GB |
| | Production Server Minimum 32 GB |
| Disk Space: | / (root) 24 GB |
| | /var/lib/pgsql Minimum 50 GB |
| | /srv Minimum 50 GB |
| | /var/spacewalk Minimum 50 GB per SUSE product and 360 GB per Red Hat product |

2.2. Branch Server Requirements

Table 2. Hardware Requirements for Branch Server

| Hardware | Recommended |
|----------|--------------------------------------|
| CPU | Minimum 2 dedicated 64-bit CPU cores |
| RAM: | Test Server Minimum 2 GB |

| Hardware | Recommended |
|-------------|--------------------------------|
| | Production Server Minimum 8 GB |
| Disk Space: | / (root) Minimum 24 GB |
| | /srv Minimum 100 GB |
| | /var/cache Minimum 100 GB |

2.3. Build Host Requirements

Table 3. Hardware Requirements for Build Host

| Hardware | Recommended |
|-------------|--------------------------------|
| CPU | Multi-core 64-bit CPU |
| RAM: | Test Server Minimum 2 GB |
| | Production Server Minimum 4 GB |
| Disk Space: | / (root) Minimum 24 GB |
| | /var/lib/Kiwi Minimum 10 GB |

2.4. Network Requirements

- The SUSE Manager Server requires a reliable and fast WAN connection.
- The branch server requires a reliable WAN connection, to reach the SUSE Manager Server.
- If you are using a dedicated network, the branch server requires at least two network interfaces: one connected to the WAN with a reachable SUSE Manager Server, and one connected to the internal branch LAN.
- Terminals require a LAN connection to the branch server network.

2.5. POS Terminal Requirements

Table 4. Hardware Requirements for Terminals

| Hardware | Recommended |
|-------------|---|
| RAM: | Minimum 1 GB for hosts that need to run OS images built with Kiwi (PXE booted or not) |
| Disk Space: | Disk space depends on size of the OS image |

For more information, see the documentation of the underlying system (in this case: SUSE Linux Enterprise Server 15).

For more information on SUSE Manager for Retail POS terminals, see documentation on SUSE Manager Salt clients ([Client-configuration](#) › [Client-config-overview](#)).

2.5.1. UEFI Secure Booting Requirements


Secure boot from the network using UEFI PXE or UEFI HTTP is supported on both SUSE Linux Enterprise Server 12 and SUSE Linux Enterprise Server 15. Booting from a hard disk using UEFI Secure Boot is fully supported on SUSE Linux Enterprise Server 15 images only.

You cannot boot SUSE Linux Enterprise Server 12 images using UEFI secure boot from a hard disk. This is due to limitations with the legacy Kiwi service. You need to either disable UEFI secure boot, or upgrade your terminals to SUSE Linux Enterprise Server 15.

Chapter 3. Network Architecture

SUSE Manager for Retail uses a layered architecture:

- The first layer contains the SUSE Manager Server.
- The second layer contains one or more branch servers to provide local network and boot services. It also contains one or more build hosts.
- The final layer contains any number of deployed point-of-service terminals.



Branch servers should be physically located close to point-of-service terminals, such as in an individual store or branch office. We recommend you have a fast network connection between the branch server and its terminals. Branch servers provide services for PXE boot, and act as an image cache, Salt broker, and proxy for software components (RPM packages). The branch servers can also manage local networking, and provide DHCP and DNS services.


SUSE Manager for Retail Branch Servers are implemented as enhanced SUSE Manager Proxies. For technical background information on SUSE Manager Proxies, see [Installation-and-upgrade › Install-proxy-unified](#).

3.1. Branch Server Network Configuration

You can use branch servers in different network configurations, depending on your installation requirements.

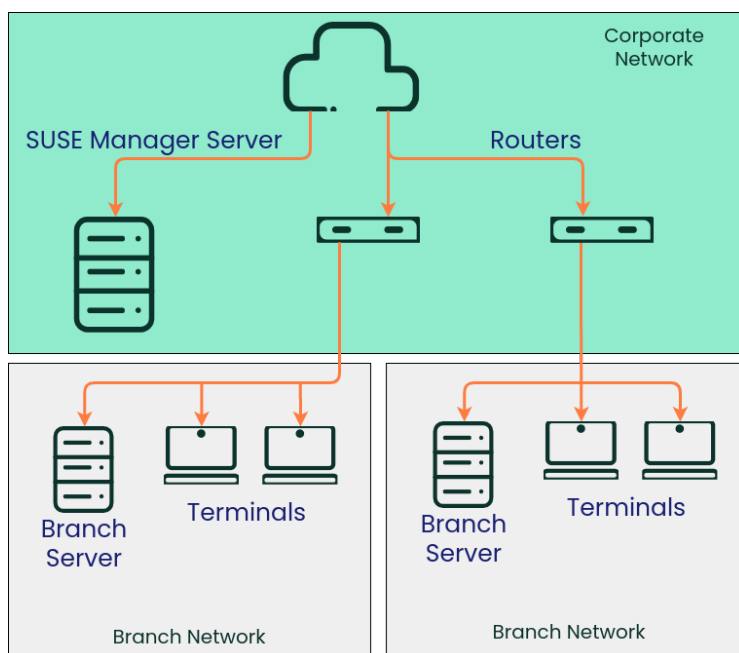
3.1.1. Dedicated Network Architecture

The branch servers are in the same network as the SUSE Manager Server, and terminals use an isolated branch network. In this configuration, the branch servers are in the corporate network, and provide all DHCP, DNS, PXE, FTP, and TFTP services to the terminals in the branch networks.




3.1.2. External Network Architecture

The branch servers are in separate branch networks, along with the terminals they manage. In this configuration, external routers provide DHCP and DNS services to the branch servers and the terminals, and the branch server provides PXE, FTP, and TFTP services to the terminals in their branch network.



3.1.3. Shared Network Architecture

The branch server and the terminals are connected to the same network as the SUSE Manager Server. In this configuration, external routers provide DHCP and DNS services to the branch servers and the terminals, and the branch server provides PXE, FTP, and TFTP services to the terminals in their branch network.



For more information about network administration on SUSE Manager for Retail, see [Retail › Retail-admin-network](#).

Chapter 4. Retail Installation

SUSE Manager for Retail and SUSE Manager for Retail Branch Server are installed using the SUSE Linux Enterprise Server Unified Installer.

4.1. Install with the Unified Installer

SUSE Manager for Retail is a SUSE base product. This section describes how to install SUSE Manager for Retail from SUSE Linux Enterprise Server installation media with the Unified Installer.




Before installing SUSE Manager, ensure your physical or virtual machine has enough disk space and RAM by checking the requirements at [Retail › Retail-requirements](#).

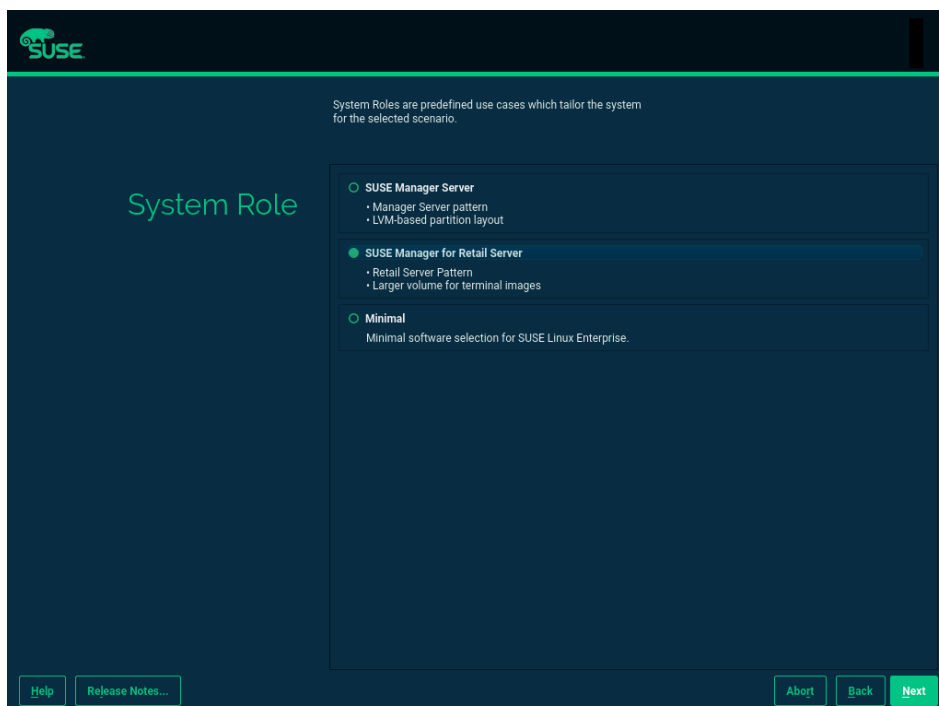
4.1.1. Install SUSE Manager for Retail

Procedure: Installing SUSE Manager for Retail Server from a DVD Image

1. Boot your server from the installation image. If booting fails you might need to adjust the boot order in the BIOS.
2. When prompted, select **Installation**.
3. In the **Language, Keyboard and Product Selection** screen, check the **SUSE Manager Server** checkbox, and click [**Next**].



4. Read and agree to the End User Licence Agreement, and click [**Next**].
5. In the **Registration** screen, check the **Register System via scc.suse.com** checkbox, enter your SUSE Customer Center credentials, and click [**Next**].
6. OPTIONAL: In the **Add On Product** screen, select any additional or add-on products you require, and click [**Next**].
7. In the **System Role** screen, check the **SUSE Manager for Retail Server** checkbox, and click [**Next**].



8. In the **Suggested Partitioning** screen, accept the default values, or use the [**Guided Setup**] or [**Expert Partitioner**] options to customize your partitioning model, and click [**Next**].
9. In the **Clock and Time Zone** screen, enter your region and timezone, and click [**Next**].
10. In the **Local Users** screen, create a new user, and click [**Next**].
11. In the **System Administrator "root"** screen, create the "root" user, and click [**Next**].
12. Review the settings on the **Installation Settings** screen. Ensure that SSH access is open.
13. On the **Installation Settings** screen click [**Install**].

When the installation procedure has finished, you can check that you have all the required modules by using the **SUSEConnect --status-text** command at a command prompt. For SUSE Manager for Retail Server, the expected modules are:

- SUSE Linux Enterprise Server Basesystem Module
- Server Applications Module
- Web and Scripting Module
- SUSE Manager Server Module

Procedure: Running the Setup Script on the SUSE Manager for Retail Server

1. On the SUSE Manager for Retail Server, at the command prompt, as root, run the setup script:

```
yast2 susemanager_setup
```

2. Follow the prompts to set up your account. Take note of the passwords you set, you will need them later on.


Continue with general SUSE Manager configuration and channel synchronization at **Installation-and-upgrade › Server-setup**.

4.1.2. Install SUSE Manager for Retail Branch Server


Procedure: Installing the Branch Server from a DVD Image

1. Boot your server from the installation image. In case of trouble, you might need to adjust the boot order in the BIOS.

2. When prompted, select **Installation**.
3. In the **Language, Keyboard and Product Selection** screen, check the **SUSE Manager Retail Branch Server** checkbox, and click **[Next]**.



4. Read and agree to the End User Licence Agreement, and click **[Next]**.
5. In the **Registration** screen, check the **Register System via scc.suse.com** checkbox, enter your SUSE Customer Center credentials, and click **[Next]**.
6. OPTIONAL: In the **Add On Product** screen, select any additional or add-on products you require, and click **[Next]**.
7. In the **System Role** screen, check the **SUSE Manager Proxy** checkbox, and click **[Next]**.



8. In the **Suggested Partitioning** screen, accept the default values, or use the [**Guided Setup**] or [**Expert Partitioner**] options to customize your partitioning model, and click [**Next**].
9. In the **Clock and Time Zone** screen, enter your region and timezone, and click [**Next**].
10. In the **Local Users** screen, create a new user, and click [**Next**].
11. In the **System Administrator "root"** screen, create the "root" user, and click [**Next**].
12. Review the settings on the **Installation Settings** screen. Ensure that SSH access is open.
13. On the **Installation Settings** screen click [**Install**].

When the installation procedure has finished, you can check that you have all the required modules by using the **SUSEConnect --status-text** command at a command prompt. For Branch Server, the expected modules are:

- SUSE Linux Enterprise Server Basesystem Module
- Server Applications Module
- Web and Scripting Module
- SUSE Manager Proxy Module
- SUSE Manager Retail Branch Server Module

Procedure: Configuring and Registering the Branch Server

1. Create an activation key based on the **SLE-Product-SUSE-Manager-Retail-Branch-Server-4.3-Pool** base channel. For more information about activation keys, see [Client-configuration › Activation-keys](#).
2. In the **Child Channels** listing, select the recommended channels by clicking the **include recommended** icon:
 - SLE-Module-Basesystem15-SP5-Pool for x86_64 SMRBS 4.3
 - SLE-Module-Basesystem15-SP5-Updates for x86_64 SMRBS 4.3
 - SLE-Module-Server-Applications15-SP5-Pool for x86_64 SMRBS 4.3
 - SLE-Module-Server-Applications15-SP5-Updates for x86_64 SMRBS 4.3
 - SLE-Product-SUSE-Manager-Retail-Branch-Server-4.3-Updates for x86_64
3. Use this activation key in SUSE Manager Proxy registration at [Installation-and-upgrade › Proxy-registration](#).
4. Configure SUSE Manager Proxy. For more information on how to do this, see [Installation-and-upgrade › Proxy-setup](#).



The branch server must be configured as a Salt managed proxy.



Cobbler TFTP is not supported on SUSE Manager for Retail. Do not configure the **susemanager-tftpsync-recv** tool on the SUSE Manager for Retail Branch Server.

4.1.3. Install SUSE Manager for Retail Build Host

Build hosts are regular SUSE Linux Enterprise Server installations registered to SUSE Manager as Salt clients. For more information on how to install and register Salt clients to SUSE Manager, see [Client-configuration › Registration-overview](#).

To prepare a build host from an already registered Salt client, see [administration:image-management.pdf](#).

4.2. Set Up the SUSE Manager for Retail Environment

To set up the SUSE Manager for Retail environment, you will need to have already installed and configured:

- SUSE Manager for Retail Server

- one or more SUSE Manager for Retail branch server proxies, or containerized proxy
- one or more SUSE Manager build hosts

This section covers how to configure your SUSE Manager for Retail environment, including:

- Prepare POS images
- Prepare system groups
- Configure services for Saltboot
- Synchronize POS images to the branch servers

The very first time you set up the SUSE Manager for Retail environment, you will need to perform all configuration steps. You will need to revisit some of these steps later on as you are working with SUSE Manager for Retail.

For example, the first time you configure the branch server, you will need to have images prepared for synchronization. If you are configuring more than one branch server, you can use the same images across different branch servers.

If you have an existing environment, and need to build new images, you do not need to re-initialize the branches. You will need to synchronize the images, and can skip setting up the services on the branch server.

Usually, POS images are rebuild when updated packages are available, and synchronized to the branch servers before the update window opens.

4.2.1. Prepare and Build Terminal Images

For information about SUSE Manager image building, see **Administration › Image-management**.

SUSE Manager for Retail POS images are images specifically tailored for SUSE Manager for Retail environment and designed to be deployed using PXE booting mechanism.

4.2.1.1. POS Image Templates

As starting point, SUSE provides basic templates at <https://github.com/SUSE/manager-build-profiles/tree/master/OSImage>. These templates need to be adapted for specific usecases, for example by including specific applications, configuration settings, and users.



By default, POS templates do not include a system user. You will not be able to

login as a user to a system that has been installed with a SUSE provided template. However you can use Salt to manage clients without a system user. You can use Salt to install a system user after the terminal has been deployed.

4.2.1.2. SLES 11 SP 3 Terminals



SLES 11 reached end of life and is no longer supported and can stop working at any moment.

POS Terminals based on SUSE Linux Enterprise Server 11 SP 3 can be deployed in much the same way as other terminals, with a few differences.

- You must use the SLES 11 template
- SLES 11 images need to be activated with the **SLES11 SP3 i586** and **SLEPOS 11 SP3 i586** channels



Ensure that SLES 11 images are built on the SLES 11 build host. Building on the incorrect build host will cause your build to fail.

If you are building images for SLES 11 using profiles from an HTTPS git repository that uses TLS 1.0 or greater, it will fail. SLES 11 does not support later versions of TLS. You will need to clone the repository locally to use it for building.

4.2.2. Branch identification and architecture topology

Before you configure the branch server, ensure you have decided on networking topology and you choose **branch id**.

For information about the possible network topologies, see [Retail › Retail-network-arch](#).

As a **branch id** select any alphanumeric string.

4.2.3. Required System Groups

SUSE Manager for Retail requires:

- branch system group for every branch server proxy, using **branch id** as its name
- hardware type system group for every used hardware type, using **HWTYPE:** prefix in its name

For more information about hardware type groups, see [Retail › Retail-deploy-terminals](#).



Missing mandatory system group will cause terminal bootstrap to fail.

SUSE Manager for Retail also recognizes two optional groups for better overview:

- **TERMINALS**
- **SERVERS**

You can create system groups using the SUSE Manager Web UI. Navigate to **Systems › System Groups** and click [**Create System Group**].

For more information about system groups, see **Reference › Systems**.

During terminal bootstrap terminal automatically joins:

- branch system group based on received **branch_id**. This will make branch group formulas available to the terminal.
- HWType group based on SMBios information received from terminal. This will make Saltboot partitioning pillar available to the terminal.
- **TERMINALS** if this group exists.



SUSE Manager for Retail command line tools create required system groups and branch group automatically.



In case you plan to use the branch server as a monitoring server with Prometheus, be aware that Prometheus demands additional hardware resources. For more information about installing Prometheus, see **Administration › Monitoring**.



In case you plan to use the branch server with Ansible software, be aware that Ansible demands additional hardware resources. For more information about installing Ansible, see **Administration › Ansible-integration**.

4.2.4. Configure Services for Saltboot

Saltboot technology is used to deploy POS images to the terminals. Saltboot consists of Saltboot enabled initrd (build as part of POS images) and Saltboot Salt states.

This section covers general information about generic Saltboot requirements. For configuration examples, see **Retail › Example-configurations**.

4.2.4.1. Enable PXE network boot in the terminal network

Saltboot is usually used in network boot environment. For this to work **DHCP** service for the network terminal is connected to must have **PXE** or sometimes called **BOOTP** support enabled.

Example 1. Example of ISC DHCP server configuration with PXE booting enabled

```
if substring (option vendor-class-identifier, 0, 10) = "HTTPClient" {
    option vendor-class-identifier "HTTPClient";
    filename "<FQDN of branch server proxy>/saltboot/shim.efi";
}
else {
    if option arch = 00:07 {
        filename "boot/shim.efi";
        next-server <IP address of branch server proxy>;
    }
    else {
        filename "boot/pxelinux.0";
        next-server <IP address of branch server proxy>;
    }
}
```

Notice two important options, **next-server** which is set to the branch server IP address and **filename** set to the **boot/pxelinux.0** for BIOS based system and **boot/shim.efi** for UEFI systems with SecureBoot support.



Containerized branch proxy uses different **filename** then regular branch server.

For containerized branch proxy set **filename** to the **pxelinux.0** for BIOS based system and **grub/shim.efi** for UEFI systems with SecureBoot.

4.2.4.2. Saltboot service discovery

Saltboot requires some information where the Salt master is and from where to download the image. Saltboot tries multiple discoveries to obtain this information, described below.

For successful terminal deployment, both service discoveries must be successful. Depending on your architecture, choose what strategy works for you best.

4.2.4.2.1. Salt master discovery

During Saltboot `initrd` start, integrated Salt client needs to find branch server proxy to connect to. This discovery is trying following steps:

- **MASTER** kernel command line option is set, then this is used as Salt master
- resolve **salt** CNAME, if successful then resolved value is used as Salt master
- use **salt** as a Salt master

Once Salt master is determined, Salt client configuration is generated and started.



Using fully qualified domain name in **MASTER** or **salt** CNAME is important.

If used fully qualified domain name is different from fully qualified domain name of branch server proxy known to SUSE Manager, Saltboot may work correctly, however proxy detection of terminal will not work.

4.2.4.2.2. Download server discovery

Before POS image is downloaded to the terminal, download server discovery is done to find where to download image from:

- **saltboot_download_server** pillar is set for terminal, then its value is used
- **saltboot:download_server** pillar is set for terminal, then its value is used
- resolve **ftp** hostname

Value obtained by download server discovery is then used together with POS image pillar to fetch correct image from correct location.

4.2.4.3. Terminal partitioning and image selection

Last piece for Saltboot is to provide partitioning for terminal. This is done individually for each hardware type of terminals. For more information about hardware types, see [Retail › Retail-deploy-terminals](#).

Above mentioned steps are mandatory minimum for successful Saltboot deployment. For configuration examples, see [Retail › Example-configurations](#).

4.2.5. Synchronize Images to the Branch Server

The OS image you use on the SUSE Manager server must be synchronized for use to the branch server. You can do this with the Salt `image-sync` state, part of the **Image Synchronization Formula**.

Procedure: Synchronizing Images to the Branch Server

1. On the SUSE Manager server, run this command:

```
salt <branch_server_minion_id> state.apply image-sync
```

2. The image details will be transferred to `/srv/saltboot` on the branch server.

You can also set synchronization to run automatically on the branch server. Configure the image synchronization formula to apply the highstate regularly. For more information about **Image Synchronization Formula**, see [Specialized-guides › Salt](#).

Chapter 5. Deploying Terminals

This section covers how to integrate terminals into your SUSE Manager for Retail environment. You can prepare the SUSE Manager for Retail installation for image deployment. Finally, you can deploy terminals using network boot and other methods.

5.1. Deployment basics

When you have the SUSE Manager Server and Branch Server set up, you are ready to deploy point-of-service terminals by following these steps:

1. Create hardware type groups
2. Assign and configure the Saltboot formula for each hardware type group
3. Synchronize images to the branch server
4. Deploy images to the terminals

Each procedure is detailed in this section.

For other methods of booting terminals, including using a USB device, or booting over a wireless network, see **Retail › Retail-deploy-terminals-other**.

For SUSE Manager 4.2 and later, terminals can be either x86-64 or ARM64 architecture. For earlier versions, terminals must be x86 architecture only.

If you have many terminals, you can handle them with a script. For more information, see **Retail › Retail-mass-config**.

Before terminals can be deployed, ensure you have prepared a Saltboot-based operating system image. For more information about building OS images, see **Administration › Image-management**.



After you have registered new terminals, always check the SUSE Manager Web UI to ensure your terminals have connected successfully to the branch server. The terminals must not have directly connected to the SUSE Manager Server by mistake.

5.1.1. Create a Hardware Type Group

Each terminal requires a specific hardware type, which contains information about the product name and terminal manufacturer. However, at the beginning, the SUSE Manager database does not have this information. To tell SUSE Manager what image to deploy on each terminal, you can set hardware type groups. After you have created a new hardware type group, you can apply the Saltboot formula to the group and configure it for your environment.

Hardware types allow you to group devices according to manufacturer and device name. Then, all devices of a particular type can be managed as one.

Empty profiles can be assigned to a hardware type group either before or after registration. If an empty profile is not assigned to a hardware type group before registration, it will be assigned to group that best matches the product information available to it.

For this procedure, you will require the system manufacturer name and product name for your terminal.

Procedure: Creating a Hardware Type Group

1. Determine the hardware type group name. Prefix the name with **HWTYPE;** followed by the system manufacturer name and product name, separated by a hyphen. For example:

```
HWTYPE:POSVendor-Terminal1
```

2. In the SUSE Manager Web UI, navigate to **Systems › System Groups**, and click the **[Create Group]** button.
3. In the **Create System Group** dialog, create a new system group, using the hardware type group name you determined in step one of this procedure.



Only use colons, hyphens, or underscores in hardware type group names. Spaces and other non-alphanumeric characters will be removed when the name is processed.

5.1.2. Assign and Configure the Saltboot Formula for Each Hardware Type Group

Each hardware type group must have the Saltboot formula applied.

Procedure: Assigning the Saltboot Formula

1. Open the details page for your new hardware type group, and navigate to the **Formulas** tab.
2. Select the Saltboot formula and click **[Save]**.
3. Navigate to the **Formulas › Saltboot** tab.
4. Configure the Saltboot formula. For more information about the Saltboot formula, see **Specialized-guides › Salt**.

5.1.3. Synchronize Images to the Branch Server

Procedure: Synchronizing Images to the Branch Server

1. On the SUSE Manager server, run this command:

```
salt <branch_server_salt_id> state.apply image-sync
```

5.1.3.1. Using a SUSE Linux Enterprise Server11 SP3 32-bit based images

If you have 32-bit machines included in your branch, then you must use a 32-bit boot image as a default boot image.



If a 32-bit boot image is not used as a default boot image, 32-bit terminals will be unable to boot and operate properly.

Check the available boot images and their architecture from the command line:


```
salt <branch_server_salt_id> pillar.item boot_images
```

Output:

```
POS_Image_JeOS6-6.0.0:
-----
arch:
  x86_64
...
legacy-6.0.0:
-----
arch:
  i686
```

In this example, the **legacy-6.0.0** boot image is 32-bit.

You can set the default boot image in the **Image Synchronization** formula on the branch server, by adding the chosen boot image name to the **Default boot image** field. For more information about **Image Synchronization** formula, see **Specialized-guides › Salt**.



5.1.4. Deploy Images to the Terminals

When you have your bootstrap image ready, you can deploy the image to the terminals.

Procedure: Deploying Images to the Terminals

1. Power on your POS terminals.
2. The branch server will bootstrap the terminals according to the data you have provided.

5.1.5. Customize the Terminal Image Download Process

You can change the terminal boot process using Salt pillars. Two Salt pillars allow you to change the protocol and server used to download the image.

- The **saltboot_download_protocol** pillar specifies which protocol should be used to download the image to the terminal. This overrides the default protocol specified in the image pillar. Allowed values are **http**, **https**, **ftp**, or **tftp**.
- The **saltboot_download_server** pillar specifies which server to use to download the image. This overrides the default hostname specified in the image pillar.

Example: Changing the Saltboot Image Download Protocol

This example changes the protocol used for all terminals.

Edit the **/srv/pillar/top.sls** file:

```
base:
  '*':
    - saltboot_proto
```

Edit the `/srv/pillar/saltboot_proto.sls` file:

```
saltboot_download_protocol: http
# can be http, https, ftp, tftp
```

Example: Changing the Saltboot Image Download Location

This example changes the download location for all terminals on a specified branch server.

Edit the `/srv/pillar/top.sls` file:

```
base:
  'minion_id_prefix:$branch_prefix':
    - match: grain
    - $branch_prefix
```

Edit the `/srv/pillar/$branch_prefix.sls` file:

```
saltboot_download_server: $download_server_fqdn
```



In this example, the download server must be prepared by the `image_sync` state before you begin.

5.2. Deploy Terminals – Other Methods

If you are not able to boot terminals from the network, you can create a live USB image and deploy terminals using a removable USB storage device. You can also bootstrap terminals across a wireless network.



Hardware type groups must be created and images must be synchronized before continuing. For more information, see [Retail › Retail-deploy-terminals](#).



After you have registered new terminals, always check the SUSE Manager Web UI to ensure your terminals have connected successfully to the branch server, and not directly to the SUSE Manager Server by mistake.

5.2.1. Deploy Terminals with a Removable USB Device

If you do not want to boot terminals from the network, you can create a live USB image and deploy terminals using a removable USB storage device. This is useful if you cannot boot your terminals from the network, or if you do not have a local SUSE Manager for Retail branch server providing network services.

You can prepare a bootable USB device with the image and tools required to deploy a POS terminal using a remote SUSE Manager for Retail branch server. You can create the bootable USB device on the branch server directly, or on the SUSE Manager for Retail Server.



POS devices booted using the USB device are assigned to the SUSE Manager for Retail branch server that created the USB device.

Procedure: Creating a Bootable USB Device

1. On the SUSE Manager for Retail branch server, at the command prompt, as root, create the POS image.

You need to specify the size of the image, in megabytes.

Ensure you allow at least 300 MB:

```
salt-call image_sync_usb.create <usb image name> <size in MB>
```

2. Insert the USB device into the SUSE Manager for Retail branch server machine, and copy the image to the new location:

```
dd bs=1M if=<usb image name> of=<path to usb device>
```

When you have the image on the USB drive, check that the terminals you want to deploy allow local booting. You can check this by editing the Saltboot formula in the SUSE Manager for Retail Web UI. For more information about the Saltboot formula, see [Specialized-guides › Salt](#).

Procedure: Deploying Images to the Terminals using USB

1. Insert the USB device into the terminal.
2. Power on the POS terminal.
3. Boot from the USB device to begin bootstrapping.

5.2.2. Deploy Terminals over a Wireless Network

For terminals that cannot be connected directly to the physical network, you can deploy them over a wireless network. Wireless networks do not support PXE booting, so you must perform the initial booting and initialization of the wireless connection on the terminal using a USB device.

For more information about using USB devices to boot, see [Retail › Retail-deploy-terminals-other](#).



Bootstrapping across a wireless network could expose a security risk if you are using encrypted OS images. The boot `initrd` image and the partition that contains `/etc/salt` must be stored unencrypted on the terminal. This allows SUSE Manager for Retail to set up the wireless network. If this breaches your security requirements, you will need to use a separate production network, with network credentials managed by Salt, so that credentials are not stored on the terminal unencrypted.

Before you begin, you need to have created a bootable USB device. Ensure that the OS image you use to create the USB device has the `dracut-wireless` package included. For more information about using USB devices to boot, see [Retail › Retail-deploy-terminals-other](#).

When you have created the USB device, you need to provide the wireless credentials to the terminal. You can do this in a number of ways:

- Directly during image build.
- Add it to the `initrd` file on the branch server.
- During terminal booting, using the kernel command line.

Procedure: Providing Wireless Credentials During Image Build

1. Ensure that the `dracut-wireless` package is included in the image template.
2. Set the wireless credentials by creating or editing the `etc/sysconfig/network/ifcfg-wlan0` file to the image template, with these details:

```
# ALLOW_UPDATE_FROM_INITRD
WIRELESS_ESSID=<wireless network name>
WIRELESS_WPA_PSK=<wireless network password>
```

If you want to use different credentials for bootstrapping to what is used during normal

operation, you can remove the `ALLOW_UPDATE_FROM_INITRD` line.

3. Build the image.
4. Prepare a USB device using this image, and boot the terminal. For more information about using USB devices to boot, see [Retail › Retail-deploy-terminals-other](#).

Procedure: Providing Wireless Credentials with initrd

1. Set the wireless credentials by creating or editing the `etc/sysconfig/network/ifcfg-wlan0` file, with these details:

```
# ALLOW_UPDATE_FROM_INITRD
WIRELESS_ESSID=<wireless network name>
WIRELESS_WPA_PSK=<wireless network password>
```

2. Copy the file to `initrd` on the branch server:

```
echo ./etc/sysconfig/network/ifcfg-wlan0 | cpio -H newc -o | gzip >>
/srv/saltboot/boot/initrd.gz
```

3. Check that the terminals you want to deploy allow local booting. You can check this by editing the Saltboot formula in the SUSE Manager for Retail Web UI. For more information about the Saltboot formula, see [Specialized-guides › Salt](#).

Procedure: Providing Wireless Credentials During Terminal Boot

1. Mount the USB device on the terminal, and boot from it.
2. Append these commands to the kernel boot parameters:

```
WIRELESS_ESSID=<wireless_network_name>
WIRELESS_WPA_PSK=<wireless_network_password>
```

5.2.2.1. Change Wireless Credentials

After you have set the wireless credentials, you can change them as needed. The way to do this is different if you use one company-wide network, or if you have each branch server on its own separate network.

Procedure: Changing Wireless Credentials for Single Network

1. Rebuild the boot image with updated credentials.
2. Recreate the bootable USB device based on the new boot image.
3. Boot terminal from new USB device.

Procedure: Changing Wireless Credentials for Multiple Networks

1. In the `/srv/salt/` directory, create a salt state called `update-terminal-credentials.sls`, and enter the new wireless network credentials:

```
/etc/sysconfig/network/ifcfg-wlan0
file.managed:
- contents: |
    WIRELESS_ESSID=<wireless_network_name>
    WIRELESS_WPA_PSK=<wireless_network_password>
# regenerate initrd
cmd.run:
- name: 'mkinitrd'
```

2. Apply the Salt state to the terminal:

```
salt <terminal_salt_name> state.apply update-terminal-credentials
```



If you are using a separate network for the boot phase, the managed file might need to be renamed, or extended to `/etc/sysconfig/network/initrd-ifcfg-wlan0`.

5.2.2.2. Use Multiple Wireless Networks

You can configure terminals to use a different set of wireless credentials during the boot process, to what they use during normal operation.

If you provide wireless credentials using `initrd` files, you can create two different files, one for use during boot called `initrd-ifcfg-wlan0`, and the other for use during normal operation, called `ifcfg-wlan0`.

Alternatively, you can use custom Salt states to manage wireless credentials with `saltboot-hook`.

First of all, you need to set the wireless details for normal operation. This will become the default settings. Then you can specify a second Salt state with the wireless details for use during the boot procedure.

Procedure: Using Different Wireless Credentials for Production Network

1. Write a custom Salt state named `/srv/salt/saltboot_hook.sls` containing the wireless details for normal operation. This Salt state is applied by Saltboot after the system image is deployed.

```
{% set root = salt['environ.get']('NEWROOT') %}
{{ root }}/etc/sysconfig/network/ifcfg-wlan0:
file.managed:
- contents: |
    WIRELESS_ESSID=<wireless_network_name>
    WIRELESS_WPA_PSK=<wireless_network_password>
- require:
  - saltboot: saltboot_fstab
- require_in:
  - saltboot: boot_system
```



The boot phase supports only WPA2 PSK wireless configuration. Salt-managed production configuration supports all features supported by all major operating systems.

5.3. Deploy Terminals and Auto-Accept Keys

You can configure SUSE Manager to automatically accept the keys of newly deployed terminals. This is achieved using Salt grains.



Automatically accepting keys is less secure than manually checking and accepting keys. Only use this method on trusted networks.

There are three different ways you can configure auto-signed grains:

- Configure Saltboot to send automatically signed grains once and then delete them. To do this, append the Saltboot configuration to an existing `initrd`. For more information, see [retail:retail-deploy-terminals-auto.pdf](#).
- Choose to keep the automatically signed grains on the Salt client. To do this, include the configuration file in the image source before the client image is built. After booting, the auto-signed grain is stored on the client as a regular Salt grain. For more information, see [retail:retail-deploy-terminals-auto.pdf](#).
- Configure Saltboot during PXE boot using kernel parameters. For more information, see [retail:retail-deploy-terminals-auto.pdf](#).

When you have configured Saltboot using one of these methods, you need to configure the SUSE Manager Server to accept them. For more information, see [retail:retail-deploy-terminals-auto.pdf](#).

5.3.1. Configure Saltboot to Send Auto-Signed Grain Once

Procedure: Configuring Saltboot to Send Auto-Signed Grain Once

1. On the branch server, create a configuration file called `/etc/salt/minion.d/autosign-grains-onetime.conf`.
2. Edit the new configuration file with these details. You can use any value you like as the auto-sign key:

```
# create the grain
grains:
  autosign_key: <AUTOSIGN_KEY>

# send the grain as part of auth request
autosign_grains:
  - autosign_key
```

3. At the command prompt, add the new configuration file to the existing `initrd`:

```
echo ./etc/salt/minion.d/autosign-grains-onetime.conf | /
cpio -H newc -o | gzip >> /srv/saltboot/boot/initrd.gz
```

5.3.2. Configure Saltboot to Keep Auto-Signed Grains

Use different procedure for SLE 15 and SLE 11/12.

Procedure: Configuring Saltboot to Keep Auto-Signed Grains (SLE 15)

1. In the location where the image source is built, such as a build host or source repository, create a configuration file called `etc/salt/minion.d/autosign-grains.conf`.
2. Edit the new configuration file with these details. You can use any value you like as the auto-sign key:

```
# create the grain
grains:
  autosign_key: <AUTOSIGN_KEY>
```

```
# send the grain as part of auth request
autosign_grains:
  - autosign_key
```

Procedure: Configuring Saltboot to Keep Auto-Signed Grains (SLE 11 and SLE 12)

1. In the location where the image source is built, such as a build host or source repository, create a configuration file called `etc/salt/minion.d/autosign-grains.conf`. This must be outside of the `root` directory provided by the template. This way you prevent the inclusion of unwanted files in the `initrd`.
2. Edit the new configuration file with these details. You can use any value you like as the auto-sign key:

```
# create the grain
grains:
  autosign_key: <AUTOSIGN_KEY>

# send the grain as part of auth request
autosign_grains:
  - autosign_key
```

3. Create a tarball of this directory:

```
tar -czf autosign-grains.tgz etc
```

4. Edit the `config.xml` template file. In the `<packages type="image">` element, add:

```
<archive name="autosign.tgz" bootinclude="true"/>
```

5. Save the file and rebuild the image.

5.3.3. Configure Saltboot to Auto-Sign During PXE Boot

Procedure: Configuring Saltboot to Auto-Sign During PXE Boot

1. Configure the PXE formula to specify these kernel parameters during booting:

```
SALT_AUTOSIGN_GRAINS=autosign_key:<AUTOSIGN_KEY>
```

2. PXE boot the Salt client. The formula creates the `./etc/salt/minion.d/autosign-grains-onetime.conf` configuration file and passes it to `initrd`.

5.3.4. Configure the Server to Auto-Accept

When you have configured Saltboot using one of these methods, you need to configure the server to accept them. The server stores the autosign keys in a file within the `/etc/salt/master.d/` directory. You can enable auto-signing by creating an auto-sign file that contains the key you created when you configured Saltboot.

Procedure: Configuring the Server to Auto-Accept

1. On the SUSE Manager Server, open the master configuration file in the `/etc/salt/master.d/` directory, and add or edit this line:

```
autosign_grains_dir: /etc/salt/autosign_grains
```

2. Create a file at `/etc/salt/autosign_grains/autosign_key`, that contains the auto-sign key you specified with Saltboot:

```
<AUTOSIGN_KEY>
```

For multiple keys, put each one on a new line.

For more information about configuring the server to automatically accept grains, see https://docs.saltstack.com/en/latest/topics/tutorials/autoaccept_grains.html.

5.4. Forced Saltboot image redeployment

Systems provisioned by Saltboot are usually redeployed or repartitioned automatically when a new image is available, or Saltboot partitioning changes.

Occasionally, however, it is needed to force Saltboot to redeploy an image or repartition disk, even when automation would not do so. For these situations, Saltboot offers three ways to force redeployment or repartitioning:

- Force Saltboot redeployment using Salt grains
- Force Saltboot redeployment using custom info values

- Force Saltboot redeployment using Saltboot API call
- Force Saltboot redeployment by custom pillar



Repartitioning of a terminal removes all data stored on the terminal hard disk, including any persistent partitions.

5.4.1. Force Saltboot redeployment using Salt grains

Saltboot redeployment grains have no side effects, and do not require any further configuration. The limitation is that terminal must be accessible by **salt**.

Procedure: Forcing Saltboot to redeploy image

1. On the SUSE Manager Server, as root, apply this Salt state at the command prompt:

```
salt $terminal_minion_id state.apply saltboot.force_redeploy
```

2. Restart the terminal to pick up the changes.

Procedure: Forcing a Saltboot to repartition the hard disk

1. On the SUSE Manager Server, as root, apply this Salt state at the command prompt:

```
salt $terminal_minion_id state.apply saltboot.force_repartition
```

2. Restart the terminal to pick up the changes.

5.4.2. Force Saltboot redeployment using custom info values

Saltboot custom values remove the limitation on terminal being reachable by **salt**, however there are configuration steps.

Custom info keys and values can be also managed using API or **spacecmd** command. For more information, see [Reference › Spacecmd](#).

Procedure: Create custom info key for image redeployment

1. In the SUSE Manager Web UI, navigate to **Systems › Custom System Info**.
2. Click **Create Key** to create new **Custom Info Key**.

3. As **Key Label** fill in `saltboot_force_redeploy`.
4. As **Description** fill in **Force redeploy Saltboot image**.
5. Click **Create Key**.



Creating `saltboot_force_redeploy` custom key is a one time operation.
When created, it is available for repeated use.

Procedure: Assign custom value for image redeployment

1. Navigate to the **Overview** page of the system you want to redeploy.
2. Select tab **Custom Info**.
3. Click on **Create Value**.
4. From the list of available keys click `saltboot_force_redeploy`.
5. As **Value** type **True**.
6. Click **Update Key**.
7. Reboot the terminal to pick up the changes.



After terminal finishes booting, Saltboot redeployment custom info setting is automatically reset to prevent repeated redeployment.

Procedure: Create custom info key for disk repartitioning

1. Navigate to menu::Systems[Custom System Info] page.
2. Click **Create Key** to create new **Custom Info Key**.
3. As **Key Label** fill in `saltboot_force_repartition`.
4. As **Description** fill in **Force terminal disk repartition**.
5. Click **Create Key**.



Creating `saltboot_force_repartition` custom key is one time operation.
Once created, it is available for repeated use.

Procedure: Assign custom value for disk repartitioning

1. Navigate to the **Overview** page of the system you want to redeploy.
2. Select tab **Custom Info**.

3. Click on **Create Value**.
4. From the list of available keys click **saltboot_force_repartition**.
5. As **Value** type **True**.
6. Click **Update Key**.
7. Reboot the terminal to pick up the changes.

5.4.3. Force Saltboot redeployment using Saltboot API call



After terminal finishes booting, Saltboot redeployment setting is automatically reset to prevent repeated redeployment.

An API call **system.setPillar** can be used to set specific Saltboot options through Salt pillar data. SUSE Manager Saltboot integration is using **tuning-saltboot** pillar category to manage Saltboot tuning, including forced redeployment or disk repartition. Using this pillar category allows SUSE Manager to reset Saltboot flag once the terminal is booted up.

Procedure: Forcing Saltboot to redeploy image using API call using **spacecmd** command

1. In the console run following the command, and replace **\$terminal_minion_id** with the actual terminal minion id:

```
spacecmd api -- -A '$terminal_minion_id,tuning-saltboot,{"saltboot":  
{"force_redeployment": "True"}}' system.setPillar
```

Procedure: Forcing Saltboot to repartition disk using API call using **spacecmd** command

1. In the console run the following command, replace **\$terminal_minion_id** with the actual terminal minion id:

```
spacecmd api -- -A '$terminal_minion_id,tuning-saltboot,{"saltboot":  
{"force_repartition": "True"}}' system.setPillar
```

Procedure: Check Saltboot tuning options

1. In the console run the following command, and replace **\$terminal_minion_id** with the actual terminal minion id:

```
spacecmd api -- -A '$terminal_minion_id,tuning-saltboot' system.getPillar
```



Make sure to use **tuning-saltboot** as pillar category in the API call.

5.4.4. Force Saltboot redeployment by custom pillar



Pillars specified outside of SUSE Manager database cannot be reset automatically. Without manual intervention, the terminal will download a new image on each reboot.

Procedure: Force a Saltboot to redeploy image using Saltboot pillar

1. Create new file `/srv/salt/pillar/force_redeploy.sls` with content:

```
saltboot:
  force_redeploy: True
```

2. Create new file or update existing file named `/srv/salt/pillar/top.sls` with content:

```
base:
  '$terminal_minion_id':
    - force_redeploy
```

3. Reboot the terminal to pick up the changes.
4. After the terminal finishes booting, remove modifications made in `/srv/salt/pillar/top.sls` file.

If your terminal encounters a problem with the file system or the partition table, you might need to remove the partition table and reformat the terminal.

Procedure: Force Saltboot to repartition disk using Saltboot pillar

1. Create new file `/srv/salt/pillar/force_repartition.sls` with content:

```
saltboot:
  force_repartition: True
```

2. Create new file or update existing file named `/srv/salt/pillar/top.sls` with content:

```
base:
  '$terminal_minion_id':
```


```
- force_repartition
```

3. Reboot the terminal to pick up the changes.
4. After the terminal finishes booting, remove modifications made in `/srv/salt/pillar/top.sls` file.

5.5. Terminal Boot Process (Saltboot Diagram)

The Saltboot process involves the SUSE Manager Server, a terminal running the Saltboot `initrd`, and the branch server providing the Saltboot services to the terminal.

This sequence diagram explains how the three components interact with each other to boot a terminal.



5.6. Terminal Names

Terminals can be named according to certain parameters, which can make it easier to match the physical device with its record in the SUSE Manager Web UI.

Naming schemes available are **Hostname**, **FQDN**, and **HWType**. Naming scheme can be selected in the **Branch Network** formula. For more information, see [Specialized-guides › Salt](#).

By default, terminals are named according to the **Hostname** naming scheme with the **HWType** scheme as a fallback.

5.6.1. Naming by HWType

Terminal names that are derived from the hardware type use this format:

```
BranchID.Manufacturer-ProductName-SerialNumber-UniqueID
```

For example:

```
B002.TOSHIBA-6140100-41BA03X-c643
```

The **BranchID** is the unique identifier for the branch server that the terminal is connected to. You can configure this value in the [Specialized-guides › Salt](#) settings for the branch server. You can disable this prefix by toggling the **Do not prefix salt client ID with Branch ID** checkbox in the [Specialized-guides › Salt](#).

The **Manufacturer**, **ProductName**, and **SerialNumber** are provided by the terminal hardware BIOS. If the terminal does not provide a serial number, it will be omitted from the terminal name.

The **UniqueID** is the first four characters of a generated machine identification number. Added unique ID is a requirement for successful terminal deployment. Without unique ID, subsequent terminal registration will fail.

5.6.2. Naming by Hostname

Terminal names that are derived from the hostname use this format:

```
BranchID.Hostname-UniqueID
```

For example:

```
B002.terminal-c643
```

The **BranchID** is the unique identifier for the branch server that the terminal is connected to. You can configure this value in the **Specialized-guides › Salt** settings for the branch server. You can disable this prefix by toggling the **Do not prefix salt client ID with Branch ID** checkbox in the **Specialized-guides › Salt**.

The **Hostname** is the plain hostname (without domain part) of the terminal.

The **UniqueID** is the first four characters of a generated machine identification number. You can disable this behavior by toggling the **Do not append unique suffix to the salt client ID** checkbox in the **Specialized-guides › Salt**.

5.6.3. Naming by FQDN

Terminal names that are derived from the Fully Qualified Domain Names (FQDN) use this format:

```
BranchID.FQDN-UniqueID
```

For example:

```
B002.terminal.example.com-c643
```

The **BranchID** is the unique identifier for the branch server that the terminal is connected to. You can configure this value in the **Specialized-guides › Salt** settings for the branch server. You can disable this prefix by toggling the **Do not prefix salt client ID with Branch ID** checkbox in the **Specialized-guides › Salt**.

The **FQDN** is the fully qualified domain name of the terminal.

The **UniqueID** is the first four characters of a generated machine identification number. You can disable this behavior by toggling the **Do not append unique suffix to the salt client ID** checkbox in the **Specialized-guides › Salt**.

5.6.4. Assign Hostnames to Terminals

If you want terminal names to be derived from the hostname, you will need to ensure your terminals have a static hostname. This requires a static IP address to be able to resolve the static hostname.

There are a number of different ways to assign hostnames to terminals. This section describes how to do this when DNS and DHCP services are managed by the branch server.

Procedure: Assigning IP Address and Hostname with Formulas

1. In the DHCP formula settings, navigate to **Hosts with Static IP Address** and click [**Add Item**].
For more information on the DHCP formula, see [Specialized-guides › Salt](#).
2. In the **Hostname** field, type the hostname of the branch server.
3. In the **IP Address** field, type the static IP address for the terminal. Ensure the IP address is within the range used by the branch server.
4. In the **Hardware Type and Address** field, type the hardware type and address in this format:

```
ethernet <terminal_MAC_address>
```

5. OPTIONAL: For multiple terminals, click [**Add Item**] and fill in the details for each terminal.
6. Click [**Save Formula**] to save the changes.
7. In the Bind formula settings, navigate to the A records of the appropriate non-reverse zone, and click [**Add Item**]. For more information on the bind formula, see [Specialized-guides › Salt](#).
8. In the **Hostname** field, type the hostname of the branch server.
9. In the **IP Address** field, type the static IP address you assigned to the terminal in the DHCP formula settings.
10. OPTIONAL: For multiple terminals, click [**Add Item**] and fill in the details for each terminal.
11. Click [**Save Formula**] to save the changes.
12. Apply the highstate on the branch server to apply the changes.



If the terminal was previously registered using a name based on the hardware type instead of the hostname, you will need to delete the previous registration. When you re-register the terminal, use the new terminal name.

Procedure: Assigning IP Address and Hostname with YAML

1. At the command prompt on the branch server, export a YAML configuration file:

```
retail_yaml --to-yaml retail.yaml
```

2. Open the YAML file and navigate to the end of the branch server section. Add a new **terminals** section if it does not already exist.
3. Add the IP address, MAC address, and hardware type for the terminal, using this format:

```
$hostname:
  IP: <IP_Address>
  hwAddress: <MAC_Address>
  hwtype: <HWTYPE_Group_name_without_HWTYPE:_prefix>
```

4. Import the modified YAML file:

```
retail_yaml --from-yaml retail.yaml
```

5. Apply the highstate on the branch server to apply the changes.



If the terminal was previously registered using a name based on the hardware type instead of the hostname, you will need to delete the previous registration. When you re-register the terminal, use the new terminal name.

For more information about using YAML configuration files, see [Retail › Retail-mass-config](#).

5.7. Offline Use

If the SUSE Manager Server is offline, you can still perform some operations on the terminals. This is useful if the connection between the branch server and the SUSE Manager Server is unstable or has low bandwidth. This feature uses caching to perform updates.

5.7.1. Offline Terminal Reboot

If the SUSE Manager Server is offline, and a terminal is rebooted, it will fall back to a previously installed image.

This will occur in these situations:

- If the Saltboot formula has not started within a specified time (default value is 60 seconds).
- If the terminal does not acknowledge that the Saltboot formula has started.
- If the root partition is specified on the kernel command line (handled by the PXE formula), is mountable (and is not encrypted), and contains `/etc/ImageVersion` (which is created by Kiwi).

You can adjust the timeout value by changing the `SALT_TIMEOUT` kernel parameter. The parameter is measured in seconds, and defaults to 60.

```
SALT_TIMEOUT = 60
```

For more about kernel parameters, see [Specialized-guides › Salt](#).

5.7.2. Cached Terminal Updates

If the bandwidth between the branch server and the terminal is low, or for optimization of the terminal update process, POS images can be cached in advance on the terminal. The upgrade can then be performed on the terminals when suitable.

This functionality requires the terminal to have a dedicated service partition. A service partition is a partition mounted as `/srv/saltboot`. This partition must be created before the system partition and large enough to store a POS image. To ensure that terminals will always have such a partition, use the Saltboot formula for terminal hardware type to specify the partition details. For more information, see [Specialized-guides › Salt](#).

When the service partition is set up on the terminal, a POS image can be downloaded in advance by applying the `saltboot.cache_image` state:

```
salt $TERMINALID state.apply saltboot.cache_image
```

This can be done regularly to ensure that terminals always have an up-to-date POS image downloaded.

When the terminal is rebooted and an up-to-date POS image is found in the service partition, the terminal will automatically use this cached image for system redeployment.

5.8. Rate Limiting Terminals

Salt is able to run commands in parallel on a large number of terminals. This can potentially create heavy load on your infrastructure. You can use rate-limiting parameters to control the load in your environment.

For more information about rate limiting on terminals, see [Specialized-guides › Salt](#).

5.8.1. Troubleshooting

Sometimes when attempting to reboot a terminal after attempting to apply the Saltboot formula, the terminal will hang at the boot screen. This can be caused by a presence ping timeout value being set at a value that is too low. You can adjust the presence ping timeout value to fix this problem.

For more information about rate limiting on terminals, see [Specialized-guides › Salt](#).

Chapter 6. Introduction to Retail Formulas

Formulas are pre-written Salt states, that are used to configure your SUSE Manager for Retail installation.

You can use the SUSE Manager Web UI to apply common SUSE Manager formulas. For the most commonly used formulas, see [Specialized-guides › Salt](#).

All formulas must be accurately configured for your SUSE Manager for Retail installation to function correctly. If you are unsure of the correct formula configuration details, run the `retail_branch_init` command before you begin to create the recommended formula configuration. You can then manually edit the formulas as required.

6.1. Branch Server Formulas

Branch servers are configured using formulas. Formulas can be configured using SUSE Manager Web UI, or the SUSE Manager XMLRPC API. To fully configure SUSE Manager for Retail, these formulas need to be enabled and configured on the branch server:

- Branch network formula, see [Specialized-guides › Salt](#)
- Bind formula, see [Specialized-guides › Salt](#)
- DHCPD formula, see [Specialized-guides › Salt](#)
- PXE formula, see [Specialized-guides › Salt](#)
- TFTP formula, see [Specialized-guides › Salt](#)
- VSFTP formula, see [Specialized-guides › Salt](#)

Optionally, you can also enable the image synchronization formula. For more information, see [Specialized-guides › Salt](#).



Badly configured formulas can result in the branch server failing to work as expected. Due to the generic nature of formulas it is difficult to do overall validation. We recommend that you configure the branch server using the SUSE Manager for Retail command line utilities, and use individual formula settings for further tuning if required. For more information, see [Retail › Retail-install-setup](#).



If a formula uses the same name as an existing Salt state, the two names will collide. This could result in the formula being used instead of the state. Always check the names of states and formulas to avoid name collisions.

When you have made changes to your formula, ensure you apply the highstate. The highstate propagates your changes to the appropriate services.

6.2. Partitioning and Image Deployment Formula

Use the Saltboot formula to specify disk partitioning, and to select which image should be deployed. For more information about the Saltboot formula, see [Specialized-guides › Salt](#).

Chapter 7. Image Pillars

If the built image type is **PXE**, a Salt pillar is also generated.

Image pillars are stored in the database and Salt subsystem can access details about the generated image. Details include where the image files are located and provided, image checksums, information needed for network boot, and more.

The generated pillar is available to all connected clients.

This is an example of image pillar:

```
{
  "images": {
    "POS_Image_JeOS7": {
      "7.1.0-1": {
        "url": "https://ftp/saltboot/image/POS_Image_JeOS7.x86_64-7.1.0-1/POS_Image_JeOS7.x86_64-7.1.0",
        "arch": "x86_64",
        "hash": "7368c101e96826053c6efde0588cf365",
        "size": 1548746752,
        "sync": {
          "url": "https://manager.example.com/os-images/1/POS_Image_JeOS7-7.1.0-1/POS_Image_JeOS7.x86_64-7.1.0",
          "hash": "7368c101e96826053c6efde0588cf365",
          "local_path": "image/POS_Image_JeOS7.x86_64-7.1.0-1"
        },
        "type": "pxe",
        "fstype": "ext3",
        "filename": "POS_Image_JeOS7.x86_64-7.1.0",
        "inactive": false,
        "boot_image": "POS_Image_JeOS7-7.1.0-1"
      }
    }
  },
  "boot_images": {
    "POS_Image_JeOS7-7.1.0-1": {
      "arch": "x86_64",
      "name": "POS_Image_JeOS7",
      "sync": {
        "initrd_url": "https://manager.example.com/os-images/1/POS_Image_JeOS7-7.1.0-1/POS_Image_JeOS7.x86_64-7.1.0.initrd",
        "kernel_url": "https://manager.example.com/os-images/1/POS_Image_JeOS7-7.1.0-1/POS_Image_JeOS7.x86_64-7.1.0-5.14.21-150400.24.55-default.kernel",
      }
    }
  }
}
```

```

    "local_path": "POS_Image_JeOS7.x86_64-7.1.0-1"
  },
  "initrd": {
    "url": "https://ftp/saltboot/boot/POS_Image_JeOS7.x86_64-7.1.0-1/POS_Image_JeOS7.x86_64-7.1.0.initrd",
    "hash": "d38b74a373bc6c9def1f069a8533d99f",
    "size": 118252253,
    "version": "7.1.0",
    "filename": "POS_Image_JeOS7.x86_64-7.1.0.initrd"
  },
  "kernel": {
    "url": "https://ftp/saltboot/boot/POS_Image_JeOS7.x86_64-7.1.0-1/POS_Image_JeOS7.x86_64-7.1.0-5.14.21-150400.24.55-default.kernel",
    "hash": "946dac0a19125d78e282afe0e3ebf0b6",
    "size": 11444416,
    "version": "5.14.21-150400.24.55-default",
    "filename": "POS_Image_JeOS7.x86_64-7.1.0-5.14.21-150400.24.55-default.kernel"
  },
  "basename": "POS_Image_JeOS7.x86_64-7.1.0"
}
}
}

```

Under the **Images** section, there is a specific image called **POS_Image_JeOS7** with a version **7.1.0-1**. It provides various details such as the image URL for download, architecture, hash, size, synchronization information, image type (in this case, **PXE**), file system type, file name, and an inactive flag to exclude it from auto-selection. Additionally, it references the corresponding boot image named **POS_Image_JeOS7-7.1.0-1**.

The **boot_images** section contains information about the boot image **POS_Image_JeOS7-7.1.0-1**. It specifies the details about kernel and **initrd**.

Image pillar can be modified via API. This example retrieves the image pillar for a given image ID, and changes the **Inactive** flag to **True** and writes it back.

```

#!/usr/bin/env python3
from xmlrpc.client import ServerProxy

MANAGER_URL = "http://manager.example.com/rpc/api"

MANAGER_LOGIN = "admin"
MANAGER_PASSWORD = "adminpass"

IMAGE_ID=15

```

```
client = ServerProxy(MANAGER_URL)
key = client.auth.login(MANAGER_LOGIN, MANAGER_PASSWORD)

pillar = client.image.getPillar(key, IMAGE_ID)
[(name, version_dict)] = pillar['images'].items()
[(version, image_data)] = version_dict.items()
image_data['inactive'] = True
client.image.setPillar(key, IMAGE_ID, pillar)

client.auth.logout(key)
```


Chapter 8. Administration

This section contains notes on administering your SUSE Manager for Retail installation. For general administration tasks, see the SUSE Manager documentation at <https://documentation.suse.com/suma/>.

8.1. Mass Configuration

Mass configuration is possible with branch servers and terminals.

8.1.1. Branch Server Mass Configuration

Branch servers are configured individually using formulas. If you are working in an environment with many branch servers, it often helps to configure branch servers automatically with a pre-defined configuration file, rather than configuring each one individually.



Before working with the mass configuration tool, back up the existing branch servers configuration.

The mass configuration tool overwrites the existing configuration with data specified in the provided YAML file.

The mass configuration tool does not support all possible formula configurations. Always make sure on a small sample that the mass configuration tool can configure systems as expected.

8.1.1.1. Configure Multiple Branch Servers

Configuring multiple branch servers requires the `python-susemanager-retail` package, which is provided with SUSE Manager for Retail. Install the `python-susemanager-retail` package on the SUSE Manager server.

Procedure: Configuring Multiple Branch Servers

1. Create a YAML file describing the infrastructure you intend to install. For an example YAML file, see **Retail › Retail-mass-config-yaml**.
2. On a clean SUSE Manager for Retail installation, import the YAML file you have created:

```
retail_yaml --from-yaml filename.yaml
```

See the `retail_yaml --help` output for additional options.

3. In the SUSE Manager Web UI, check that your systems are listed correctly. Also check that the formulas you require are available.
4. Create activation keys for each of your branch servers, register them using bootstrap, and configure them as proxy servers. For more information, see [Retail › Retail-install-unified](#).
5. In the **States** tab, click **[Apply Highstate]** to deploy your configuration changes for each branch server.

If you need to change your configuration, you can edit the YAML file at any time, and use the `retail_yaml --from-yaml` command to upload the new configuration.

Use empty profiles together with activation keys to onboard all the systems of your infrastructure. Use an activation key to assign the channels listed in [Retail › Retail-install-setup](#).

8.1.2. Terminal Mass Configuration

If you are working in an environment with many terminals, it often helps to configure terminals automatically with a pre-defined configuration file, rather than configuring each one individually.

You will need to have your infrastructure planned out ahead of time, including the IP addresses, hostnames, and domain names of branch servers and terminals. You will also require the hardware (MAC) addresses of each terminal.



The settings specified in the configuration file cannot always be successfully applied. In cases where there is a conflict, SUSE Manager will ignore the request in the file. This is especially important when designating hostnames. You should always check that the details have been applied as expected after using this configuration method.

8.1.2.1. Configure Multiple Terminals

Procedure: Configuring Multiple Terminals

1. Create a YAML file describing the infrastructure you intend to install, specifying the hardware address for each terminal. For an example YAML file, see [Retail › Retail-mass-config-yaml](#).

2. On a clean SUSE Manager installation, import the YAML file you have created:

```
retail_yaml --from-yaml filename.yaml
```

See the `retail_yaml --help` output for additional options.

3. In the SUSE Manager Web UI, check that your systems are listed and displaying correctly, and the formulas you require are available.
4. Create activation keys for each of your branch servers, connect them using bootstrap, and configure them as proxy servers. For more information, see [Retail › Retail-install-unified](#).
5. In the **States** tab, click [**Apply Highstate**] to deploy your configuration changes for each branch server.
6. Connect your terminals according to your infrastructure plan.

If you need to change your configuration, you can edit the YAML file at any time, and use the `retail_yaml --from-yaml` command to upload the new configuration.

8.1.3. Export Configuration to Mass Configuration File

If you already have a configuration that you would like to export to a YAML file, use:

```
retail_yaml --to-yaml filename.yaml
```

This can be a good way to create a basic mass configuration file. However, it is important to check the file before using it, because some mandatory configuration entries may be missing.



When you are designing your configuration and creating the YAML file, ensure the branch server ID matches the fully qualified hostname, and the Salt ID. If these do not match, the bootstrap script could fail.

8.2. Example YAML File for Mass Configuration

You can use the `retail_yaml` command to import configuration parameters from a manually prepared YAML file. This section contains a YAML example file with comments.

Listing 1. Example: YAML Mass Terminal Configuration File

```
branches:
```

```

# there are 2 possible setups: with / without dedicated NIC
#
# with dedicated NIC
branchserver1.branch1.cz: # salt ID of branch server
  branch_prefix: branch1 # optional, default guessed from salt id
  server_name: branchserver1 # optional, default guessed from salt id
  server_domain: branch1.cz # optional, default guessed from salt id
  nic: eth1 # nic used for connecting terminals, default taken from hw info in
SUMA
  dedicated_nic: true # set to true if the terminals are on separate network
  salt_cname: branchserver1.branch1.cz # hostname of salt master / broker for
terminals, mandatory
  configure_firewall: true # modify firewall configuration
  branch_ip: 192.168.2.1 # default for dedicated NIC: 192.168.1.1
  netmask: 255.255.255.0 # default for dedicated NIC: 255.255.255.0
  dyn_range: # default for dedicated NIC: 192.168.1.10 - 192.168.1.250
    - 192.168.2.10
    - 192.168.2.250
# without dedicated NIC
# the DHCP formula is not used, DHCP is typically provided by a router
# the network parameters can be autodetected if the machine is already connected to
SUSE Manager
branchserver2.branch2.cz: # salt ID of branch server
  branch_prefix: branch2 # optional, default guessed from salt id
  server_name: branchserver2 # optional, default guessed from salt id
  server_domain: branch2.cz # optional, default guessed from salt id
  salt_cname: branchserver2.branch1.cz # FQDN of salt master / broker for terminals,
mandatory
  branch_ip: 192.168.2.1 # optional, default taken from SUMA data if the machine is
registered
  netmask: 255.255.255.0 # optional, default taken from SUMA data if the machine is
registered
  exclude_formulas: # optional, do not configure listed formulas
    - dhcp # without dedicated NIC the dhcp service is typically provided by a
router
  hwAddress: 11:22:33:44:55:66 # optional, required to connect pre-configured entry with
particular machine
# during onboarding
terminals: # configuration of static terminal entries
  hostname1: # hostname
  hwAddress: aa:aa:aa:bb:bb:bb # required as unique id of a terminal
  IP: 192.168.2.50 # required for static dhcp and dns entry
  saltboot: # optional, alternative 1: configure terminal-specific pillar data
  partitioning: # partitioning pillar as described in saltboot documentation
  disk1:
    device: /dev/sda
    disklabel: msdos

```

```

partitions:
  p1:
    flags: swap
    format: swap
    size_MiB: 2000.0
  p2:
    image: POS_Image_JeOS6
    mountpoint: /
    type: DISK
hostname2:          # hostname
hwAddress: aa:aa:aa:bb:bb:cc # required as unique id of a terminal
IP: 192.168.2.51      # required for static dhcp and dns entry
hwtype: IBMCORPORATION-4838910 # optional, alternative 2: assign the terminal to
hwtype group
# if neither of hwtype nor saltboot is specified, a group is assigned according to
hwtype
# reported by bios on the first boot
hwtypes:
  IBMCORPORATION-4838910:  # HWTYPE string (manufacturer-model) as returned by
  bios
  description: 4838-910    # freetext description
saltboot:
  partitioning:          # partitioning pillar as described in saltboot documentation
  disk1:
    device: /dev/sda
    disklabel: msdos
    partitions:
      p1:
        flags: swap
        format: swap
        size_MiB: 1000.0
      p2:
        image: POS_Image_JeOS6
        mountpoint: /
        type: DISK
TOSHIBA-6140100:
  description: HWTYPE:TOSHIBA-6140100
  saltboot:
    partitioning:
    disk1:
      device: /dev/sda
      disklabel: msdos
      partitions:
        p1:
          flags: swap
          format: swap
          size_MiB: 1000.0

```

```
p2:
  image: POS_Image_JeOS6
  mountpoint: /
  type: DISK
```

8.3. Delta Images

If you have very large images that you need to synchronize to the branch server, you can use delta images to save network bandwidth.

A delta image contains only the differences between two regular images. If there are only a few changes between two images, the delta image can be very small. Synchronizing a delta image to the branch consumes less network bandwidth but it requires some extra hardware resources on the branch server to rebuild the installable image.

8.3.1. Building Delta Images

The **retail_create_delta** tool creates a delta image on the SUSE Manager server. The tool uses **xdelta3** internally.

Use the name and the version strings of the target and the source image as parameters to the command. The format of the parameters must be **<NAME>-<VERSION>** and they must correspond to the image names and versions available in the pillar. For example, if the pillar contains:

```
images:
  POS_Image_JeOS6:
    6.0.0:
      ...
    6.0.1:
      ...
  POS_Image_Graphical6:
    6.0.0:
      ...
```

Then the **retail_create_delta** command is:

```
retail_create_delta POS_Image_JeOS6-6.0.1 POS_Image_JeOS6-6.0.0
```

This command will generate the delta image between version 6.0.0 and version 6.0.1. The resulting delta file is saved in **/srv/www/os-images** and the corresponding pillar file in **/srv/susemanager/pillar_data/images/**.

8.3.2. Tuning Delta Generation

Performance tuning is possible with the `-B <VALUE>` option, which is passed to `xdelta3`. With higher values you achieve smaller deltas at the cost of higher memory requirements. For more information, see the `xdelta3` documentation (`man xdelta3`).

8.3.3. Image Synchronizing to the Branch Server

When an image is synchronized to the branch server, the `image-sync-formula` first checks whether the source image is available on the branch server. Only if the source image is available, the delta will be downloaded to save network bandwidth.

8.4. Network Administration

If you are intending to set up either an external or a shared network architecture, you need to ensure that the server providing DHCP services has PXE support enabled.

For example:

```
next-server: <branch server IP>
if option arch = 00:07 {
    filename "boot/shim.efi";
} else {
    filename "boot/pxelinux.0";
}
```

Configure your DNS servers to resolve `salt` and `ftp` as CNAMEs to the correct branch server FQDN.

If using a CNAME is not possible, there are several workarounds:

- For `salt`, set this kernel parameter when the terminal boots:

```
MASTER=<branch_server_fqdn>
```

You can configure this using the PXE formula. For more information, see [Specialized-guides › Salt](#).

- For `ftp`, you can use an A record that resolves to an IP address instead of a CNAME. Alternatively, you can change the terminal boot process using Salt pillars. For more information, see [Retail › Retail-deploy-terminals](#).

For a description of the different networking architectures, see [Retail › Retail-network-arch.](#)

Chapter 9. Retail Migration

This section provides instructions for migrating SUSE Linux Enterprise Point of Service II, SUSE Manager for Retail 3.1, or SUSE Manager for Retail 3.2 to the newest version of SUSE Manager for Retail.

9.1. Before You Migrate

This document is intended to guide you through migration your SUSE Linux Enterprise Point of Service or older SUSE Manager for Retail installation (3.1 or 3.2) to the newest version of SUSE Manager for Retail.

This document is divided into scenarios. Pick the scenario that best suits your environment, and follow the instructions in that section to migrate your installation.



Ensure your existing installation is fully updated, and that you have performed a backup, before you begin your migration.

9.1.1. Prepare to Migrate from SUSE Linux Enterprise Point of Service

SUSE Linux Enterprise Point of Service cannot be upgraded directly to SUSE Manager for Retail. The migration requires you to perform some manual configuration. To assist you in the migration, as much information as possible about the existing hardware configuration and network infrastructure is recorded. Then this information is used for rebuilding the new SUSE Manager for Retail installation.

In some cases, this will require a lengthy downtime to perform the migration. If you are not able to manage downtime, you can install new servers and run them in parallel to the existing ones while you perform the migration. This is especially relevant for large installations.

It is possible to run a SLEPOS Admin server and SUSE Manager Server in parallel. In such a scenario, branches that have been migrated will run on the SUSE Manager server, while those that have not yet been migrated can continue to run on the SLEPOS Admin server. This includes all operations, such as adding new terminals, or building and deploying new images.

However, if you run network services (especially DHCP) on the branch servers, you will not be able to run both old and new branch servers in parallel on the same branch, because they can conflict with each other. This can result in multiple terminals having the same IP address, or terminals randomly assigned to different branch servers. If you need to migrate in this

environment, and you want to configure a new branch server while the branch is still running on old infrastructure, make sure that the new branch server is not connected to the network with the terminals.

If your branch server does not provide DHCP services, you can configure the new one in parallel and, when you are ready, change the configuration of your DHCP server from the old to the new branch server.

9.2. Migrate SUSE Linux Enterprise Point of Service II to SUSE Manager for Retail

This section describes migrating from an existing SUSE Linux Enterprise Point of Service II installation to a new SUSE Manager installation. You can perform this migration all at once by creating a data dump in a single file, and then moving it to the new server.

Alternatively, you can perform the migration in stages by creating a data dump for each branch, and moving them to the new server one by one. Importing and deploying the converted data can also be done in one or multiple steps, depending on your environment.

9.2.1. Migration with Complete Data Dump

In this procedure, you create a single data dump in an XML file, convert it to YAML, and migrate it to the new infrastructure all at once.

1. Install the SUSE Manager for Retail server 4.3. For more information, see [Retail › Retail-install-unified](#).
2. On the SLEPOS Admin server export all the data stored in LDAP to an XML file. Run this command as an administrator:

```
posAdmin --export --type xml --file dumpfile.xml
```

The resulting **dumpfile.xml** file will contain global information, with parts about images, hardware and its partitioning, and the description of the branch servers with networking data, services, and attached terminals.

3. Move the XML file to the newly created SUSE Manager server, and convert it to YAML:

```
retail_migration dumpfile.xml retail.yml
```

4. Review the generated YAML file (**retail.yml**) and adjust it as necessary. Consider **HWTYPE** group naming and image name and version changes in the partitioning data. Group names must not exceed the 56 character limit. You can shorten the names as needed, and the image names must match the images in SUSE Manager. The **--save-mapping** option can help you with this task.

Also check whether there are duplicate MAC addresses of the terminals in the generated YAML file. Choose which entry you want to keep. If there are duplicate MAC addresses, importing the YAML file will fail.



SUSE Linux Enterprise Point of Service images will not be migrated. You must rebuild the images using the OS image building functionality. For more information about building images, see **Administration › Image-management**.

5. Import the complete data (YAML) with:

```
retail_yaml --from-yaml retail.yml
```

You can see statistical data while the import is in progress. Check the results in the Web UI by navigating to **Systems › System List › All** and find empty profiles. To see the groups for the hardware configuration, branches, servers, and terminals, navigate to **Systems › Groups**.

To finalize the branch server migration, you must install the branch server machines as Salt clients and bootstrap them as proxies. For more information about proxy installation, see **Retail › Retail-install**. For more information about using an activation key to assign the required channels, see **Retail › Retail-install-setup**. After onboarded to SUSE Manager, the branch servers machines are connected with the empty profiles (by FQDN), and so they will get the Retail configuration.

After all the branches are migrated, shutdown and remove the old SLEPOS Admin Server.

9.2.2. Migration with Branch by Branch Data Dump

In this procedure, you migrate the SLEPOS infrastructure and the branches one by one, first exporting and then importing.

1. Install the SUSE Manager for Retail server 4.3 For more information, see **Retail › Retail-install-unified**.

2. On every branch server:

```
posAdmin --export --type xml --file dumpfile.xml
```

These dumps will contain only the LDAP data of the branch, and any global data.

3. Similarly, you can export the LDAP data of every branch if you run the command on the Admin server with the branch credentials explicitly specified:

```
posAdmin --export --type xml --file dumpfile.xml --user $branch_dn \
--password $password
```

For background information about SLEPOS branch server configuration, see https://documentation.suse.com/sle-pos/11-SP3/html/SLEPOS-guide/cha.slepos_branchserv.html.

4. Review the generated YAML file (**retail.yml**) and adjust it as necessary. Consider **HWTYPE** group naming and image name and version changes in the partitioning data. You can shorten the names as needed, and the image names must match the images in SUSE Manager. The **--save-mapping** option can help you with this task.
5. Check whether there are duplicate MAC addresses of the terminals in the generated YAML file. Choose which entry you want to keep. As long as there are duplicate MAC addresses, SUSE Manager will refuse importing the YAML file.



SUSE Linux Enterprise Point of Service images will not be migrated. You must rebuild the images using the OS image building functionality. For more information about building images, see **Administration › Image-management**.

The data can be imported branch by branch. For each branch perform the following steps:

1. Run the import command for one branch after the other:

```
retail_yaml --from-yaml retail.yml --branch <branch_name>
```

Repeat the command for every branch.

2. To finalize each branch server migration, you must install the branch server machine as a Salt-based client and bootstrap it as a proxy. For more information about proxy installation,

see [Installing and Registering](#). For more information about using an activation key to assign the required channels, see [Configuring Server](#). After onboarded to SUSE Manager for Retail, the branch server machine is connected with the empty profile (by FQDN), and so it will get the Retail configuration.

3. Apply Highstate on the branch server; this will happen automatically if **Configuration File Deployment** is enabled.
4. Boot the terminals of the branch.

After all the branches are migrated, shut down and remove the old SLEPOS Admin Server.

9.2.3. Converting XML to YAML

When you perform a migration using one of the methods in this chapter, one of the steps takes the XML data dump file from SUSE Linux Enterprise Point of Service, and converts it to a YAML file for SUSE Manager for Retail. The tool that performs this conversion has additional features, which are outlined in this section.

To validate the XML file before conversion, and print any errors:

```
retail_migration dumpfile.xml
```

To write a mapping file called `map.yml`:

```
retail_migration dumpfile.xml --save-mapping map.yml
```

The mapping file contains two dictionaries:

1. **images**, which maps old SUSE Linux Enterprise Point of Service images to new images built in SUSE Manager.
2. **groups**, which maps legacy SUSE Linux Enterprise Point of Service **scCashRegister** objects to SUSE Manager **HWTYP**E groups. Group names must not exceed the 56 character limit.

The mapping file should be edited as required for your environment.

To perform a conversion using a mapping file:

```
retail_migration dumpfile.xml retail.yml --mapping map.yml
```

If you are performing a branch-by-branch migration, the resulting **retail.yml** file will contain a new version of SUSE Linux Enterprise Point of Service LDAP data. If you want to preserve any global changes in your SUSE Manager for Retail settings, remove the **global** hardware types from the resulting **retail.yml** file before importing it. Alternatively, you can import **retail.yml** using this command to import only the new systems and groups defined in the file, and leave any existing configuration settings untouched:

```
retail_yaml --only-new
```

9.3. Upgrade SUSE Manager for Retail Branch Server

This section describes upgrading the SUSE Manager for Retail Branch Server to the next SP (service pack).

The SUSE Manager for Retail Branch Server is a client system similar to the SUSE Manager Proxy, with additional SUSE Manager for Retail features.



Upgrade the SUSE Manager Server before starting the SUSE Manager for Retail upgrade.

Procedure: Upgrading the SUSE Manager for Retail Branch Server

1. For general information about upgrading a proxy client, see [Installation-and-upgrade › Proxy-intro](#).
2. After the proxy upgrade is complete, apply the highstate on the SUSE Manager for Retail Branch Server. When applying the highstate, the retail functionality will also be updated.

Chapter 10. Example configurations

This section contains some example configurations using various configuration techniques.

- **Retail › Dedicated-with-formulas**

Manual configuration of all branch services using formulas with forms to configure Saltboot with dedicated network for POS terminals.

- **Retail › Dedicated-with-scripts**

Automatic configuration of all branch services using SUSE Manager for Retail scripts to configure Saltboot with dedicated network for POS terminals.

- **Retail › Shared-central-dns**

Configuration of basic branch services using formulas with forms to configure Saltboot in shared network environment.

- **Retail › Containerized-saltboot**

Configuration of Saltboot using containerized proxy.

10.1. Set Up the SUSE Manager for Retail Environment with dedicated network for terminal

To set up the SUSE Manager for Retail environment, you will need to have already installed and configured:

- SUSE Manager for Retail Server
- one or more SUSE Manager for Retail branch server proxies
- one or more POS images built

10.1.1. Assumptions

In this example we consider architecture with dedicated network for POS terminals where branch server provides all required services including DHCP and DNS.

For dedicated network we assume 192.168.86.0/24 network.

In this example we provide **salt** and **ftp** CNAMEs for Saltboot service discovery.

We assume to have one regular branch server with salt minion id **branch1.mystore.com**, which is equal to the branch server fully qualified domain name, and branch id **B0001**.

As a terminal we assume to have one terminal with hardware manufacturer **TerminalOEM** and model **T1000**.

For POS image we assume to have one with name **POS_Image_JeOS7**.

This example covers two configuration methods with both of them achieving same result.

10.1.2. Create required system groups

Unlike when using provided tools, manually setting up formulas requires manual creation of needed system groups. Follow guide [Reference › Systems](#) to create the system groups:

- **TERMINALS**
- **HWType:TerminalOEM-T1000**
- **B0001**

First group is generic optional group for collecting all POS terminals. Second group is hardware type group for our POS terminal. Third group is mandatory branch group.

For more information about Saltboot groups, see [Retail › Retail-install-setup](#).

10.1.3. Assign and configure branch server formulas

We assign following formulas to the branch server:

- Branch Network
- Dhcpd
- Bind
- Pxe
- Tftpd

Procedure: Assign formulas to the branch server

1. Select the **Formulas** tab for the system **branch1.mystore.com**
2. You should see list of formulas in **Configuration** sub-tab
3. In displayed list select **Branch Network**
4. In displayed list select **Pxe**
5. In displayed list select **Dhcpd**
6. In displayed list select **Tftpd**
7. In displayed list select **Bind**
8. Save assigned formulas by clicking [**Save**] in top right corner

Procedure: Configure Branch Network formula

1. Select Branch Network formula from branch server formula list
2. Keep **Dedicated Nic** check-box selected, because we are using dedicated network for terminals
3. Keep the rest with default values
4. In the **Terminal Naming and Identification** section edit **Branch Identification** and set it to our branch id **B0001**
5. Save the formula

Procedure: Configure Pxe formula

1. Select Pxe formula from branch server formula list
2. At the end of **Kernel Command Line Parameters** append **MASTER=branch1.mystore.com**
3. Keep the rest with default values
4. Save the formula

Procedure: Configure Dhcpd formula

1. Select Dhcpd formula from branch server formula list
 - Set the **Domain Name** to the **branch1.mystore.com**
 - Set the **Domain Name Server** to the IP address of branch server **192.168.86.1**)
 - Set the **Listen Interfaces** to the **eth1**, which is network interface of branch server connected to dedicated terminal network

2. Navigate to the **Network Configuration (subnet)** section, and use these parameters for Network1:
 - In the **Network IP** field, enter the IP address of the terminal network **192.168.86.0**.
 - In the **Netmask** field, enter the network mask of the terminal network **255.255.255.0**.
 - In the **Domain Name** field, enter the domain name for the terminal network **branch1.mystore.com**).
3. In the **Dynamic IP Range** section, use these parameters to configure the IP range to be served by the DHCP service:
 - In the first input box, set the lower bound of the IP range **192.168.86.10**
 - In the second input box, set the upper bound of the IP range **192.168.86.250**
4. In the **Broadcast Address** field, enter the broadcast IP address for the terminal network **192.168.86.255**
5. In the **Routers** field, enter the IP address to be used as default route in the terminal network **192.168.86.1**
6. In the **Next Server** field, enter the IP address of the branch server for PXE booting **192.168.86.1**
7. Set the **Filename** to the **/boot/pxelinux.0**
8. Set the **Filename Efi** to the **/boot/shim.efi**
9. Set the **Filename Http** to the **<http://192.168.86.1/saltboot/boot/shim.efi>**
10. Save the formula

Procedure: Configure Tftpd formula

1. Select Tftpd formula from branch server formula list
2. Set the **Internal Network Address** to **192.168.86.1**, this way tftp will be listening on any IPv4
3. Set the **TFTP base directory** to **/srv/saltboot**
4. Save the formula

Procedure: Configuring Bind with reverse name resolution

1. Select Bind formula from branch server formula list
2. In the **Config** section, select **Include Forwarders** so DNS server can forward queries to next DNS server

3. In the **Options** section click **+** to add an option
4. Set the **Option** to **empty-zones-enable**
 - Set the **Value** to **No**
5. In the **Configured Zones** section, use these parameters for Zone 1, which is our primary zone:
 - Set the **Name** to **branch1.mystore.com**
 - In the **Type** field, select **master**
6. Click [**Add item**] to add a second zone, and set these parameters for Zone 2, which is used for reverse name resolution:
 - Set the **Name** to **com.mystore.branch1**
 - In the **Type** field, select **master**
7. In the **Available Zones** section, use these parameters for Zone 1:
 - In the **Name** field, enter the domain name **branch1.mystore.com**
 - In the **File** field, type the name of your configuration file **branch1.mystore.com.txt**
8. In the **Start of Authority (SOA)** section, use these parameters for Zone 1:
 - In the **Nameserver (NS)** field, use the FQDN of the branch server **branch1.mystore.com**
 - In the **Contact** field, use the email address for the domain administrator
 - Keep all other fields as their default values
9. In the **Records** section, in subsection **A**, use these parameters to set up an A record for Zone 1:
 - In the **Hostname** field, use the hostname of the branch server **branch1.mystore.com.**, notice trailing **.** which are required here
 - In the **IP** field, use the IP address of the branch server **192.168.86.1**
10. In the **Records** section, subsection **NS**, use these parameters to set up an NS record for Zone 1:
 - In the input box, use the branch server **branch1.mystore.com.**
11. In the **Records** section, subsection **CNAME**, use these parameters to set up CNAME records for Zone 1:
 - In the **Key** field, enter **salt**, and in the **Value** field, type the branch server **branch1.mystore.com.**
 - Click [**Add item**] to add another entry.
 - In the **Key** field, enter **ftp**, and in the **Value** field, type the branch server

branch1.mystore.com.

12. Set up Zone 2 using the same parameters as for Zone 1, but ensure you use the reverse details:
 - The same SOA section as Zone 1.
 - Empty A and CNAME records.
 - Additionally, configure in Zone 2:
 - Set **Generate Reverse** to the network IP address **192.168.86.0/24**
 - Set **For Zones** to the domain name of your branch network **branch1.mystore.com**
13. Click [**Save Formula**] to save your configuration.
14. Apply the highstate.

10.1.4. Setup partitioning

Partitioning is specific to the hardware type and configured using **Saltboot** formula.

Procedure: Assign Saltboot formula to hardware type group

1. Navigate to **Systems › System Groups**
2. Select group **HWType:TerminalOEM-T1000**, which is our hardware type group
3. Select the **Formulas** tab once in group details
4. You should see list of formulas in **Configuration** sub-tab
5. In displayed list select **Saltboot**
6. Save assigned formulas by clicking [**Save**] in top right corner

Procedure: Configure Saltboot formula

1. Select Saltboot formula from **HWType:TerminalOEM-T1000** group formula list
2. Set **Disk Symbolic ID** to **Disk1**
3. Set **Device Type** to **DISK**
4. Set **Disk Device** to *****
5. Set **Partition table type** to **gpt**
6. Click [**+**] to add a partition

- Set Partition Symbolic ID to p1
 - Set Partition Size (MiB) to 512
 - Set Device Mount Point to /boot/efi
 - Set Filesystem Format to vfat
 - Set Partition Flags to boot
7. Click [+] to add a partition
- Set Partition Symbolic ID to p2
 - Set Device Mount Point to /
 - Set OS Image to Deploy to POS_Image_JeOS7
8. Save the formula

After all procedures are done, apply highstate on the branch server.

10.1.5. Synchronize images

After highstate is applied, we proceed with synchronizing images as usual with apply `image-sync` state.

Terminal can now be started and will be automatically provisioned, pending salt key acceptance.

10.2. Set Up the SUSE Manager for Retail Environment with dedicated network for terminal using bundled scripts

To set up the SUSE Manager for Retail environment, you will need to have already installed and configured:

- SUSE Manager for Retail Server
- one or more SUSE Manager for Retail branch server proxies
- one or more POS images built

10.2.1. Assumptions

In this example we consider architecture with dedicated network for POS terminals where branch server provides all required services including DHCP and DNS.

For dedicated network we assume 192.168.86.0/24 network.

In this example we provide **salt** and **ftp** CNAMEs for Saltboot service discovery.

We assume to have one regular branch server with salt minion id **branch1.mystore.com**, which is equal to the branch server fully qualified domain name, and branch id **B0001**.

As a terminal we assume to have one terminal with hardware manufacturer **TerminalOEM** and model **T1000**.

For POS image we assume to have one with name **POS_Image_JeOS7**.

This example covers two configuration methods with both of them achieving same result.

10.2.2. Quick set up of SUSE Manager for Retail using for Retail tools

- SUSE Manager for Retail provides command line tool **retail_branch_init**
- SUSE Manager for Retail provides mass import command line tool **retail_yaml**

The branch server can be configured automatically using the **retail_branch_init** command, as shown in this section.

Procedure: Configuring Branch Server Formulas With a Helper Script

1. Run following command on SUSE Manager Server:

```
retail_branch_init branch1.mystore.com --dedicated-nic eth1 \
    --branch-prefix B0001 \
    --server-domain branch1.mystore.com \
    --branch-ip 192.168.86.1 \
    --dyn-range 192.168.86.10 192.168.86.250 \
    --netmask 255.255.255.0
```

You can use the **retail_branch_init --help** command for additional options.

2. Verify that your changes have been configured correctly by checking the SUSE Manager Web UI branch server system formulas.
3. Apply highstate on the branch server.

Similar results can be achieved by using mass import command line tool.

Procedure: Configuring Branch Server Formulas With a Mass Import Tool

1. Prepare branch specific YAML file:

For example, create **branch.yaml** file with content:

```
branches:
  branch1.mystore.com:
    branch_prefix: B0001
    server_domain: branch1.mystore.com
    nic: eth1
    dedicated_nic: true
    configure_firewall: true
    branch_ip: 192.168.86.1
    netmask: 255.255.255.0
    dyn_range:
      - 192.168.86.10
      - 192.168.86.250
```

For more information about mass import tool, see **Retail › Retail-mass-config**.

2. Import branch information from YAML file to SUSE Manager

```
retail_yaml --from-yaml branch.yaml
```

3. Verify that your changes have been configured correctly by checking the SUSE Manager Web UI branch server system formulas.
4. Apply highstate on the branch server.



Both **retail_branch_init** and **retail_yaml** commands override existing configuration settings of the specified branch server.

After the initial configuration done by command line tools, branch server configuration can be further adjusted in SUSE Manager Web UI through branch server formulas.

10.2.3. Setup partitioning

Follow **Reference › Systems** to create **HWType:TerminalOEM-T1000** system group.

Partitioning is specific to the hardware type and configured using **Saltboot** formula.

Procedure: Assign Saltboot formula to hardware type group

1. Navigate to **Systems › System Groups**
2. Select group **HWType:TerminalOEM-T1000**, which is our hardware type group
3. Select the **Formulas** tab once in group details
4. You should see list of formulas in **Configuration** sub-tab
5. In displayed list select **Saltboot**
6. Save assigned formulas by clicking **[Save]** in top right corner

Procedure: Configure Saltboot formula

1. Select Saltboot formula from **HWType:TerminalOEM-T1000** group formula list
2. Set **Disk Symbolic ID** to **Disk1**
3. Set **Device Type** to **DISK**
4. Set **Disk Device** to *****
5. Set **Partition table type** to **gpt**
6. Click **[+]** to add a partition
 - Set **Partition Symbolic ID** to **p1**
 - Set **Partition Size (MiB)** to **512**
 - Set **Device Mount Point** to **/boot/efi**
 - Set **Filesystem Format** to **vfat**
 - Set **Partition Flags** to **boot**
7. Click **[+]** to add a partition
 - Set **Partition Symbolic ID** to **p2**
 - Set **Device Mount Point** to **/**
 - Set **OS Image to Deploy** to **POS_Image_JeOS7**
8. Save the formula

After all procedures are done, apply highstate on the branch server.

10.2.4. Synchronize images

After highstate is applied, we proceed with synchronizing images as usual with apply `image-sync` state.

Terminal can now be started and will be automatically provisioned, pending salt key acceptance.

10.3. Set Up the SUSE Manager for Retail Environment with shared network between terminals, branch servers and SUSE Manager

To set up the SUSE Manager for Retail environment, you will need to have already installed and configured:

- SUSE Manager for Retail Server
- one or more SUSE Manager for Retail branch server proxies
- one or more POS images built

10.3.1. Assumptions

In this example we consider architecture with shared network between POS terminals, branch server and SUSE Manager where branch server provides only PXE configuration and images.

In this example we provide **salt** on kernel command line option and **ftp** in pillar for Saltboot service discovery.

We assume to have one regular branch server with salt minion id **branch1.mystore.com**, which is equal to the branch server fully qualified domain name, and branch id **B0001**.

As a terminal we assume to have one terminal with hardware manufacturer **TerminalOEM** and model **T1000**.

For POS image we assume to have one with name **POS_Image_JeOS7**.

There is already existing external DHCP and DNS service. DHCP service is correctly configured to PXE boot from branch server.

10.3.2. Create required system groups

Follow guide [Reference › Systems](#) to create the system groups:

- **TERMINALS**
- **HWType:TerminalOEM-T1000**
- **B0001**

First group is generic optional group for collecting all POS terminals. Second group is hardware type group for our POS terminal. Third group is mandatory branch group.

For more information about Saltboot groups, see [Retail › Retail-install-setup](#).

10.3.3. Assign and configure formulas

We assign following formulas to the branch server:

- Branch Network
- Pxe
- Tftpd

Procedure: Assign formulas to the branch server

1. Select the **Formulas** tab for the system **branch1.mystore.com**
2. You should see list of formulas in **Configuration** sub-tab
3. In displayed list select **Branch Network**
4. In displayed list select **Pxe**
5. In displayed list select **Tftpd**
6. Save assigned formulas by clicking [**Save**] in top right corner

Procedure: Configure Branch Network formula

1. Select Branch Network formula from branch server formula list
2. Unselect **Dedicated Nic** check-box because we are not using dedicated network
3. Keep the rest with default values
4. In **Terminal Naming and Identification** section edit **Branch Identification** and set it to our

branch id **B0001**

5. Save the formula

Procedure: Configure Pxe formula

1. Select Pxe formula from branch server formula list
2. At the end of **Kernel Command Line Parameters** append **MASTER=branch1.mystore.com**
3. Keep the rest with default values
4. Save the formula

Procedure: Configure Tftpd formula

1. Select Tftpd formula from branch server formula list
2. Set **Internal Network Address** to **0.0.0.0**, this way tftp will be listening on any IPv4
3. Set **TFTP base directory** to **/srv/saltboot**
4. Save the formula



Next two formulas are group formulas. Make sure you are editing those formulas in their respective groups.

Group formulas are visible as assigned also on system level, but editing them will create system specific setting instead of group setting.

Assign **Saltboot Group** formula to the **B0001** group and **Saltboot** formula to the **HWType:TerminalOEM-T1000** group.

Procedure: Assign Saltboot Group formula to Branch group

1. Navigate to **Systems › System Groups**
2. Select group **`B0001**, which is our branch group
3. Select the **Formulas** tab once in group details
4. You should see list of formulas in **Configuration** sub-tab
5. In displayed list select **Saltboot Group**
6. Save assigned formulas by clicking [**Save**] in the top right corner

Procedure: Configure Saltboot Group formula

1. Select Saltboot Group formula from **B0001** group formula list
2. Set **Image download server** to **branch1.mystore.com**
3. Unselect **Branch server is containerized proxy** because this is not containerized proxy
4. Save the formula

Procedure: Assign Saltboot formula to hardware type group

1. Navigate to **Systems › System Groups**
2. Select group **HWType:TerminalOEM-T1000**, which is our hardware type group
3. Select the **Formulas** tab once in group details
4. You should see list of formulas in **Configuration** sub-tab
5. In displayed list select **Saltboot**
6. Save assigned formulas by clicking [**Save**] in top right corner

Procedure: Configure Saltboot formula

1. Select Saltboot formula from **HWType:TerminalOEM-T1000** group formula list
2. Set **Disk Symbolic ID** to **Disk1**
3. Set **Device Type** to **DISK**
4. Set **Disk Device** to *****
5. Set **Partition table type** to **gpt**
6. Click [**+**] to add a partition
 - Set **Partition Symbolic ID** to **p1**
 - Set **Partition Size (MiB)** to **512**
 - Set **Device Mount Point** to **/boot/efi**
 - Set **Filesystem Format** to **vfat**
 - Set **Partition Flags** to **boot**
7. Click [**+**] to add a partition
 - Set **Partition Symbolic ID** to **p2**
 - Set **Device Mount Point** to **/**
 - Set **OS Image to Deploy** to **POS_Image_JeOS7**

8. Save the formula

After all procedures are done, apply highstate on the branch server.

10.3.4. Synchronize images

After highstate is applied, we proceed with synchronizing images as usual with apply `image-sync` state.

Terminal can now be started and will be automatically provisioned, pending salt key acceptance.

10.4. Set Up the SUSE Manager for Retail Environment using containerized proxy

To set up the SUSE Manager for Retail environment, you will need to have already installed and configured:

- SUSE Manager for Retail Server 4.3 or newer
- one or more SUSE Manager for Retail containerized proxies
- one or more SUSE Manager build hosts

This section covers how to configure your environment using containerized SUSE Manager Proxy for **Saltboot** deployment.



Containerized SUSE Manager for Retail environment does not use SUSE Manager for Retail Branch Servers.



Containerized workflow requires POS images build using SUSE Manager Server 4.3 or newer. Older images will not work.



Containerized workflow not longer implicitly configure DHCP for PXE booting. See [Specialized-guides › Salt](#) how to use DHCP formula to configure DHCP server.

Without using DHCP formula, make sure your DHCP server has correct PXE booting setting pointing to containerized proxy. See below for example.

10.4.1. Assumptions

In this example we are going to use branch id **B0001**.

As a terminal we assume to have one terminal with hardware manufacturer **TerminalOEM** and model **T1000**.

For POS image we assume to have one with name **POS_Image_JeOS7**.

10.4.2. Create required system groups

Follow guide [Reference › Systems](#) to create the system groups:

- **TERMINALS**
- **HWType:TerminalOEM-T1000**
- **B0001**

First group is generic optional group for collecting all POS terminals. Second group is hardware type group for our POS terminal. Third group is mandatory branch group.

For more information about Saltboot groups, see [Retail › Retail-install-setup](#).

We assign **Saltboot Group** formula to group **B0001** we just created. With this, our branch group is converted to **Saltboot Group**.

10.4.3. Saltboot group

Containerized SUSE Manager for Retail is configured in **Saltboot Group**.

Saltboot Groups are branch groups, system group with branch id as its name, with **Saltboot Group** formula enabled.

Saltboot Group
Expand All Sections

Search by formula's group name

Settings for Saltboot deployment group

Image download server:

?

Branch server is containerized proxy:

☒

?

Default boot image for new registrations:

?

Kernel parameters for the group:

?

Saltboot client naming scheme:

Hostname

?

Do not prefix salt client ID with saltboot group ID:

☐

?

Do not append unique suffix to the salt client ID:

☐

?

Saltboot Group formula is a successor of **Branch Network** formula, **PXE** formula and **TFTP** formula used in regular SUSE Manager for Retail setups.

Name of the **Saltboot group** is automatically used as a **branch id**, an identifier for group of machines booted through particular containerized SUSE Manager Proxy.

All **saltboot** deployed machines though containerized proxy will automatically become members of its **Saltboot group**.

To connect **Saltboot group** with containerized proxy fill **Image Download Server** entry with Fully Qualified Domain Name (**FQDN**) of the containerized proxy.

With this, mandatory configuration is finished. The rest of configuration is optional.

10.4.3.1. Default boot image

Configure **Default boot image for new registrations** to specify what boot image should be booted by not yet registered POS terminal. This is useful when stable boot image is wanted for initial deployments. Without this setting, newest built boot image is used as default boot image.

If **Default boot image for new registrations** is set, option to set its version appears under name **Default boot image version** where specific image version can be set.

10.4.3.2. Kernel option for the Saltboot group

Option **Kernel parameters for the group** can be used to pass extra kernel options to all POS

terminals registered with this Saltboot group.

10.4.3.3. Naming scheme for new registrations

Last three options are related to how will be newly registered machine visible in SUSE Manager Server.

See [Retail › Retail-terminal-names](#) for explanation of possible configurations.

10.4.4. Comparing containerized and non-containerized workflows

External DHCP service must be used with containerized Saltboot.

For more information about how to enable PXE booting in DHCP service, see [Retail › Retail-install-setup](#).

Containerized workflow relies on updated image building in SUSE Manager Server 4.3 where PXE images are no longer collected as bundle, but kernel, initrd and filesystem image are collected individually.

Containerized workflow uses new TFTP container which instead of providing files present on the proxy, routes TFTP requests as HTTP requests through local proxy to SUSE Manager Server.

Containerized proxy is not a Salt client, it is not possible to call **image-sync** state.



Once POS image is build and made available on SUSE Manager Server, it is immediately available to the Saltboot clients as well. Image synchronization is not needed, nor available. This may have implications on how images are deployed to production.

The following sections differentiate between containerized and regular workflow. Both are assuming proxy (containerized or in form of SUSE Manager for Retail Branch Server) are available.

Containerized workflow:

1. Build POS image
2. Configure DHCP server for PXE booting for given network
3. Create **Saltboot group** and configure it for existing containerized proxy
4. Boot system to be deployed

Non-Containerized workflow:

1. Build POS image
2. Configure and apply Retail formulas on SUSE Manager for Retail Branch server
3. Apply **highstate** state on the Branch server
4. Create **branch group** with the name of Branch ID as set in retail formulas
5. Apply **image-sync** state on configured SUSE Manager for Retail Branch server
6. Boot system to be deployed

10.4.5. Validating Saltboot group configuration

Containerized **Saltboot** utilizes **Cobbler** system underneath for managing PXE and UEFI configuration.

When new PXE image is built (such as SUSE Manager for Retail POS_Image_JeOS images) **cobbler distro** and **cobbler profile** are automatically generated for this image.

For example when first image **POS_Image_JeOS** version **7.0.0** is build under organization with number 1 Cobbler list will show:

```
# cobbler list

distros:
  1-POS_Image_JeOS7-7.0.0-1

profiles:
  1-POS_Image_JeOS7-7.0.0-1
```

These entries contain information about kernel and initrd. These entries are however not yet available for PXE booting.

Only when **Saltboot group** is created, new Cobbler profile is created for this **Saltboot group** which points to **cobbler distro** based on default boot image configuration.

For example, when system group **B0001** is created and **Saltboot group formula** is assigned and configured for this group, new Cobbler profile is created.

```
# cobbler list

distros:
```

```
1-POS_Image_JeOS7-7.0.0-1
```

```
profiles:
```

```
1-POS_Image_JeOS7-7.0.0-1
```

```
1-B0001
```

When inspecting this new group using command **cobbler profile report --name 1-B0001** details of this profile reveal configuration of this Saltboot group.

```
# cobbler profile report --name 1-B0001
```

```
Name : 1-B0001
```

```
Comment : Saltboot group B0001 of organization SUSE default profile
```

```
Distribution : 1-POS_Image_JeOS7-7.0.0-1
```

```
Kernel Options : {'MASTER': ['downloadserver.example.org'], 'MINION_ID_PREFIX':  
['B0001']}
```

Kernel options in example are always present and are internal for Saltboot functionality.

With this information **Cobbler** is able to generate required PXE and UEFI Grub configurations which can be checked in `/srv/tftpboot/pxelinux.cfg/default` and `/srv/tftpboot/grub/x86_64/menu_items.cfg`.

These files contain the end result which will be used by PXE client when determining what to boot and with what parameters.

Chapter 11. Best practices

11.1. Deploying new images

This file describes how to deploy a new image on selected terminals first, test it, and then deploy it to all terminals.

The hardware type group (i.e. **HWTYPE** group) is used for booting new terminals, while already deployed terminals can be moved to any group that provides the Saltboot formula. This flexibility allows for testing of new images or configurations.

Procedure: Dual group image deployment

1. Build new image. Use image synchronization formula to synchronize the image to all Branch Servers.
2. Create a test group.
 - Name the group, for example, **TEST:POSVendor-Terminal**. The name can be chosen freely.
 - Add the Saltboot formula to this group. Use the same configuration as in the original hardware type group (**HWTYPE**), except for the image version. Use the name and version of the new image.
3. Move selected terminals.
 - Identify the terminals you want to test the new image on.
 - Transfer these terminals from the group **HWTYPE:POSVendor-Terminal** group to the newly created group **TEST:POSVendor-Terminal**. It is recommended to use System Set Manager for this task.
4. Reboot terminals.
 - Power on the selected POS terminals. They should boot the new image.
5. Verify and evaluate.
 - Monitor the terminals in the **TEST:POSVendor-Terminal** group.
 - Check if the new image performs as expected and if there are any issues or errors encountered.
6. Update **HWTYPE** group.
 - Once you have confirmed that the new image works correctly, the **HWTYPE** group can be

updated.

- Modify the image version in the **HWTYP**:**POSVendor-Terminal** group to match the tested and approved version.
- This ensures that all other terminals receive the new image.

11.1.1. Controlling image versions

By default, Saltboot formula selects the specified image name and version or selects the highest one if the version is not specified.

There are multiple methods to control the image version within the **HWTYP** and **TEST** groups:

- Specify exact version.
 - In both the **HWTYP** and **TEST** groups, specify explicitly the exact version of the image to be deployed.
 - Make sure that the specified image version is already synchronized.
- Mark new image as **Inactive**.
 - In the **HWTYP** group, do not specify the image version.
 - Mark the new image as **Inactive** after building and before synchronizing it to Branch Server.
 - The **Inactive** flag ensures, that the image is not auto-selected when the image version is not specified.
 - After testing, remove the **Inactive** flag so the image can be auto-selected in the **HWTYP** group.
 - The **Inactive** flag can be accessed through the API. For more information, see [Retail › Retail-image-pillars](#).
- Utilize **Freeze Image** flag.
 - In the **HWTYP** group, omit the image version.
 - Before synchronizing the new image, set the **Freeze Image** flag in Saltboot formula in **HWTYP** group. This ensures that the already deployed terminals in this group will keep the old image.
 - Synchronize and test the new image in the **TEST** group.
 - Reset the **Freeze Image** flag in **HWTYP** group, so the terminals in this group will be

updated to the new image too.



Do not use the whitelist in Image Synchronization formula to control which terminal boots which image. It does not provide the required granularity and does not work with containerized Branch Server.

Chapter 12. What Next?

This document covers only a sub-section of information about your SUSE Manager for Retail installation. If you need further information or support, try one of these options.

12.1. More Documentation

For SUSE Manager documentation, visit <https://documentation.suse.com/suma/4.3/>.

For legacy SUSE Linux Enterprise Point of Service documentation, see <https://documentation.suse.com/sle-pos/11-SP3/>. For legacy SUSE Manager for Retail documentation, see <https://documentation.suse.com/suma-retail/4.1/>. Note, however, that SUSE Manager for Retail documentation supersedes this information.

12.2. Support

For personalized support, log in to your SUSE Customer Center account at <https://scc.suse.com/login>.

For assistance with planning and installing your SUSE Manager for Retail environment, contact the SUSE Consulting team.

Chapter 13. GNU Free Documentation License

Copyright © 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that

overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely

XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in

the Addendum below.

- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under

this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document

is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".