



## Lenovo Open Cloud Automation Product Guide

Lenovo Open Cloud Automation is a software solution that helps to securely accelerate the deployment of distributed Edge infrastructure, at scale. The main users that it serves are the deployment architect, the field technician, and the deployment project manager. With Lenovo Open Cloud Automation, distributed Edge infrastructure is deployed faster, with fewer resources required while also reducing the overall CO2 emissions of the process.

Lenovo Open Cloud Automation achieves this by embodying a simplified deployment process and by automating the steps of the process too complex for humans to execute without mistake. This delivers a standardized experience for both the and the field technicians, while still being customer, deployment, and site specific.

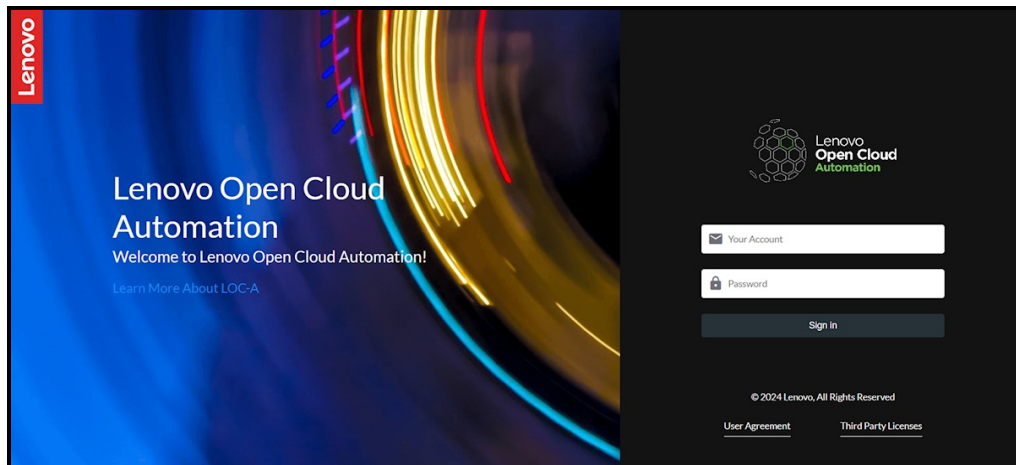


Figure 1. LOC-A for Edge 3.0 login

### Did You Know?

LOC-A enables a single deployment architect to handle many infrastructure installs simultaneously, all from the same deployment template. Running multiple installs saves time, driving more user productivity and delivering value faster for the business.

Lenovo Open Cloud Automation for Edge targets medium (100 sites) to large (>10000 sites) Edge Infrastructure deployments in diverse environments such as Retail, Telco, Oil & Gas, Healthcare, and others.

It uses a no-code, UI driven workflow to orchestrate meta-data handling, edge node onboarding, OS deployment and cluster creation. LOC-A utilizes an encrypted network connection to securely onboard devices and communicate with the BMC, and it implements the concepts of Templates and Instances to enable scalable edge cluster bring-up.

## What's new in LOC-A 3.0

Lenovo recently announced LOC-A version 3.0, improving the overall functionality, including:

- Support deployment of OS and Clusters on SE455 v3
- Server BMC scalable Day 0 setup with Centrally controlled policies
- Scalable Password management through Policies
- Plugin architecture – enablement of partner plugin development with limited assistance.
- OS image side-loading – by-passing network bandwidth limitations.
- ThinkAgile VX Deployer now part of LOC-A

## Part numbers

The following table lists the ordering information for LOC-A.

Table 1. LOC-A Edge version ordering information

Part number	Description
7S0Y0005WW	Lenovo Open Cloud Automation RedHat 1-socket ThinkEdge server with 3 years support. Price per node
7S0Y0007WW	Lenovo Open Cloud Automation RedHat 1-socket ThinkEdge server with 5 years support. Price per node
7S0Y0006WW	Lenovo Open Cloud Automation VMware 1-socket ThinkEdge server with 3 years support. Price per node
7S0Y0008WW	Lenovo Open Cloud Automation VMware 1-socket ThinkEdge server with 5 years support. Price per node

## Features for users

Those designated as LOC-A users have access to portal pages related primarily to Edge infrastructure deployment and maintenance tasks. Users can upload meta-data, on-board devices, create deployment templates and policies and deploy, and monitor their results through the portal pages.

The following features are available to users:

- [Dashboard](#)
- [Deployment planning](#)
- [Image Repository population](#)
- [Device profiles configuration](#)
- [Credential policies configuration](#)
- [Edge Cluster Template creation – with flexible naming convention](#)
- [nZTP Registration image generation](#)
- [Device on-boarding using nZTP](#)
- [Device on-boarding using XCC IP](#)
- [Device on-boarding using XLSX file](#)
- [Device on-boarding using SLP-based discovery](#)
- [Edge Cluster Instance lifecycle management \(creation, expansion, deletion\)](#)

## Dashboard

The dashboard provides an overview of the latest activities that happened in the system. The latest tasks are also listed, indicating the start time and the order of tasks executed. The figure below displays the home menu.

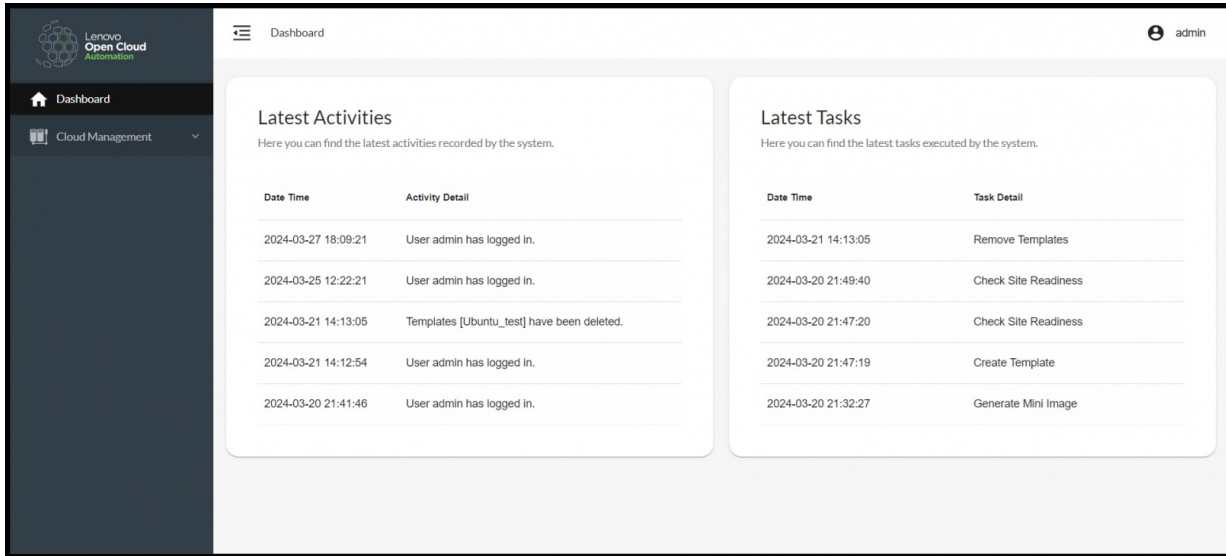


Figure 2. LOC-A Dashboard

## Deployment planning

Deployment planning is a step that the Infrastructure Admin Team is responsible for. LOC-A expects that the output of the planning activity is filled into an excel file with a specific template that the Deployment Engineer can download from the LOC-A support page. Once that file is filled with the correct data about the deployment, the Infrastructure Admin Team can proceed with uploading the file into the LOC-A portal.

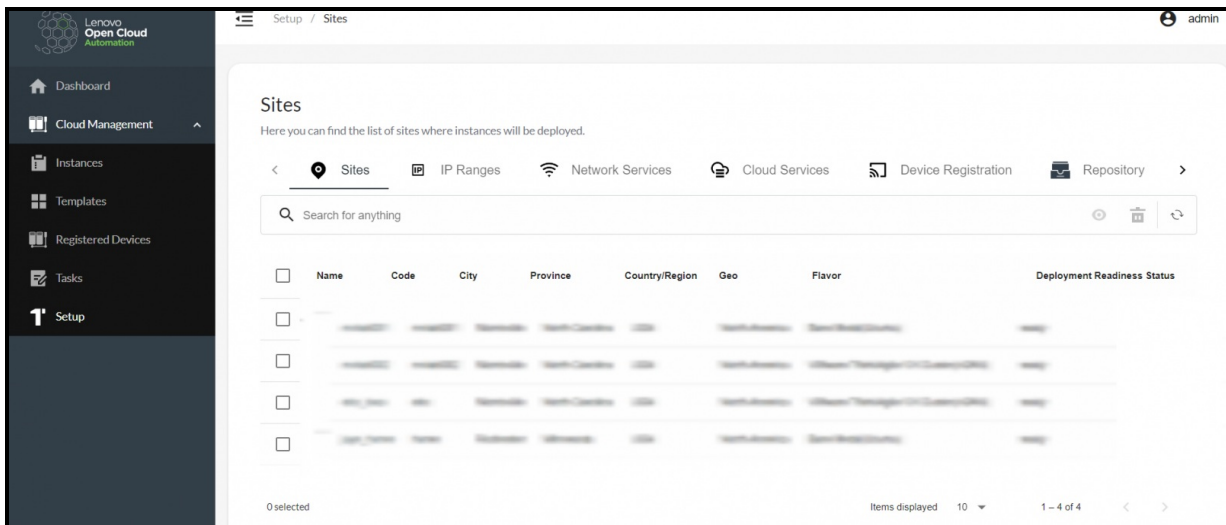


Figure 3. Meta-data onboarded in the Setup page

## Image Repository population

Each Flavor used by a Cluster Template requires at least 1 OS image or artifact (for firmware update). Instead of including the OS image in the Cluster Template, we give users the flexibility to upload their own OS image for a particular flavor. To do that, LOC-A offers a repository in the Setup section. Below is an example of how that repository looks like:

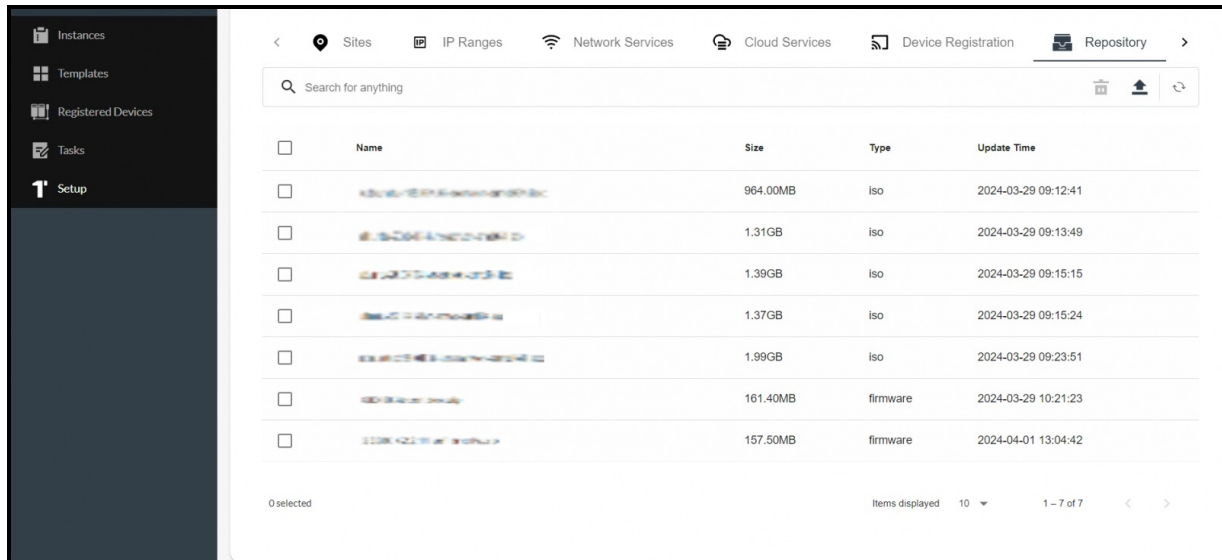


Figure 4. Image repository

### Device profiles configuration

In LOC-A, every Cluster flavor can be mapped to a specific BMC configuration profile to meet the requirements of such flavors. Some of the settings that are needed by such flavors are: XCC: “Power Restore Policy”, UEFI: “Secure Boot”, etc. Device profiles are configured centrally and automatically applied when a device is on-boarded, before the edge cluster is installed.

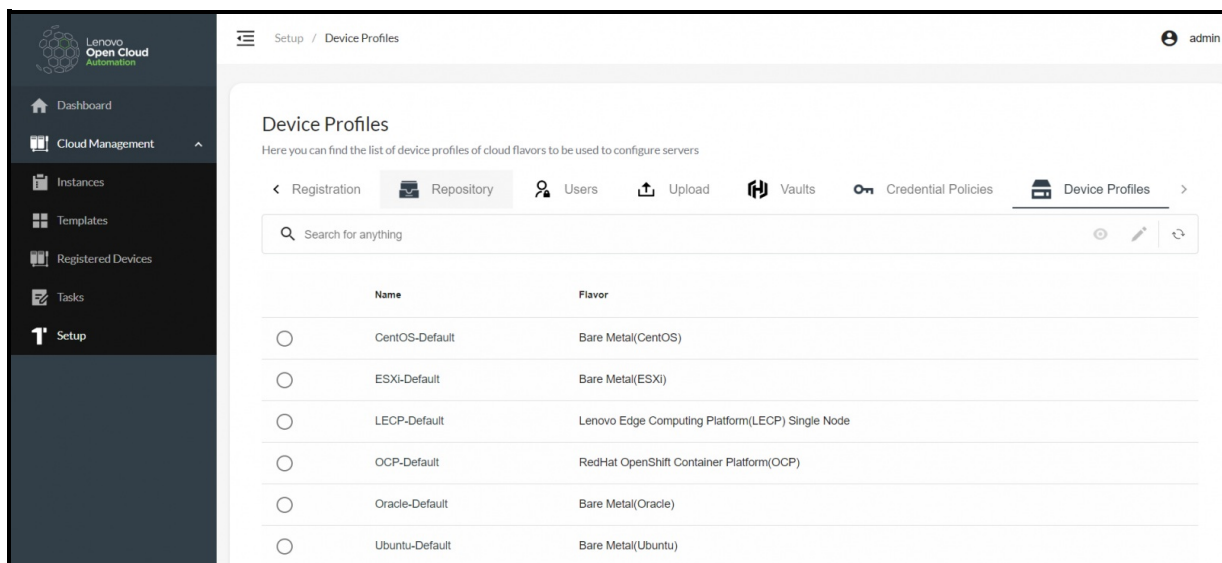


Figure 5. Device Profile selection and configuration




Figure 6. Device Profile editing

## Credential policies configuration

LOC-A uses and sets three types of server credentials: BMC, UEFI, and OS. To improve the security posture and allow only the Infrastructure Admin Team to control how credentials are set, Lenovo Open Cloud Automation implements a few rules applicable to how credentials are generated and used throughout the server pool.

These policies are of 2 types: “static” or “auto generated”:

- **Static:** This type gives the option to set the same Infrastructure Admin Team created credentials for all servers deployed with the same Cluster template, for consistency purposes and easy management.
- **Auto generated:** This type gives the option to generate a new set of credentials for every device and for every type of login (OS, BMC, UEFI). The goal here is to remove the burden of having to generate and memorize unique passwords for every server and login combination, while at the same time not exposing the infrastructure to potential password discovery and re-use attacks.

Other types of credentials are also supported as part of the cluster flavor and cluster template definition (e.g.: pull secrets, certificates, etc.).

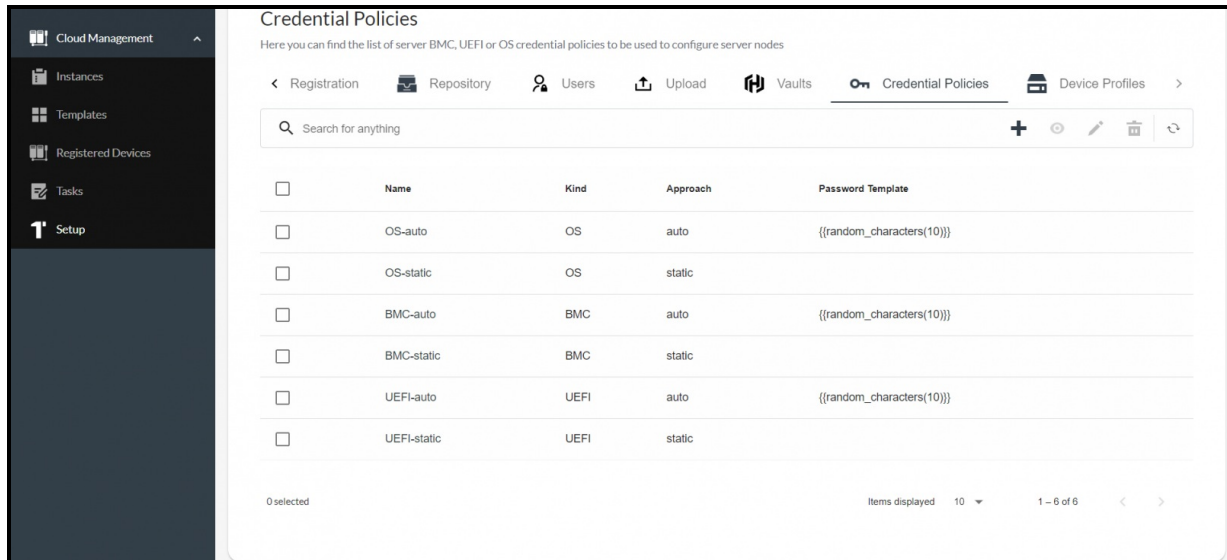


Figure 7. Credential Policies list

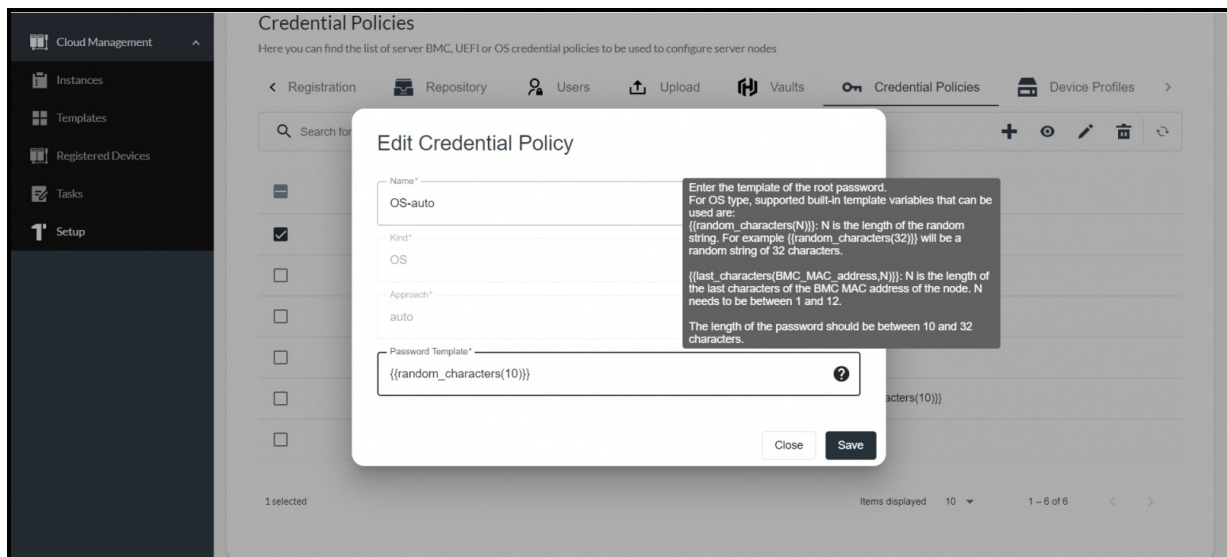


Figure 8. Credential Policy Editing – with rules

## Edge Cluster Template creation – with flexible naming convention

One of the most used concepts from the Lenovo Open Cloud Automation portal is the infrastructure definition Template. Once meta-data is loaded and servers are on-boarded, the Infrastructure Admin Team uses the infrastructure Template to create instances of clusters or bare-metal nodes, in parallel and at scale. To do that, the Infrastructure Admin Team must first create one or more infrastructure Templates.

A sample Template for an Ubuntu bare-metal deployment is displayed in the figure below.

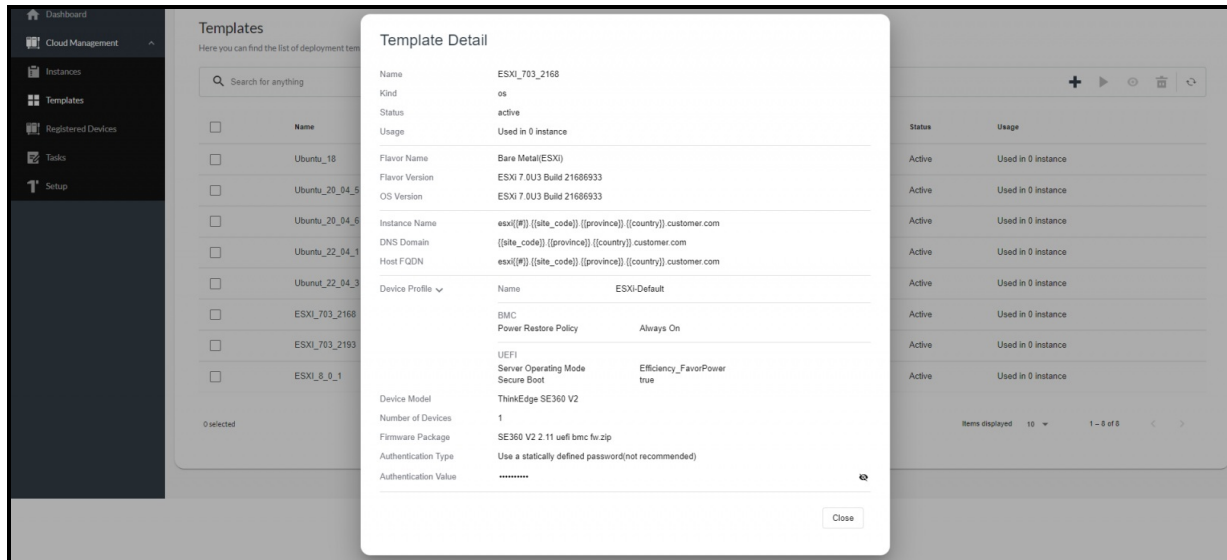


Figure 9. Bare-metal Ubuntu definition Template

## nZTP Registration image generation

To enable near-Zero Touch Provisioning of ThinkEdge servers, after uploading the meta-data describing the deployment, the Infrastructure Admin Team generates one or more Registration Images that allows the Field Technicians to perform the registration process using their Laptop and the Lenovo Open Cloud Automation Registration Utility. In the figure below you can see a generated registration image.

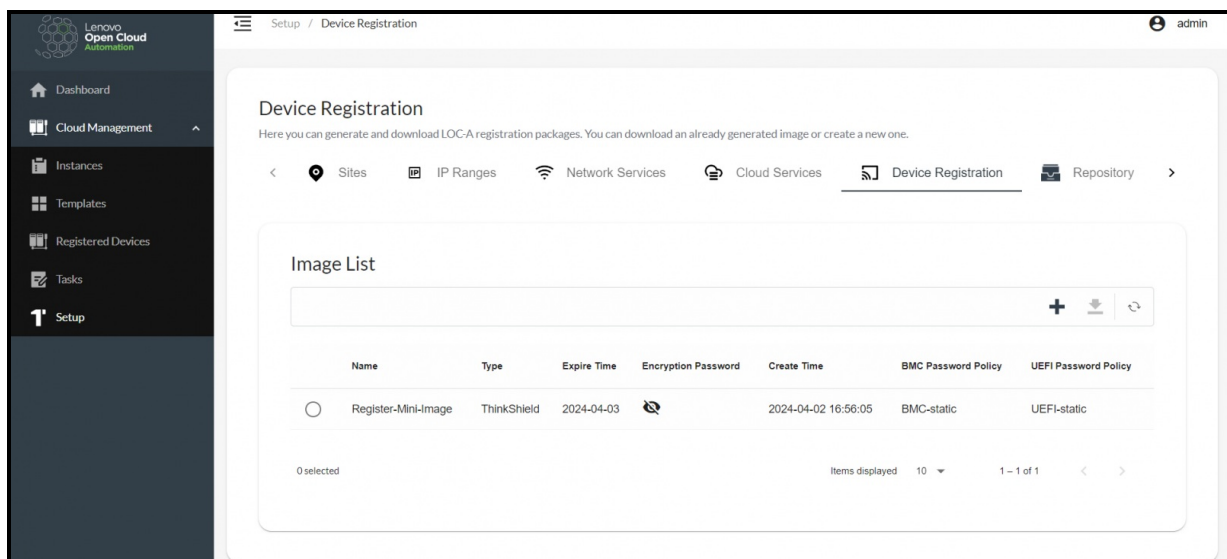


Figure 10. Generated Registration Image in the Lenovo Open Cloud Automation portal

## Device on-boarding using nZTP

For greenfield deployment scenarios, Lenovo Open Cloud Automation implements a near Zero Touch Provisioning mechanism. The Field Technician, after setting up the devices on the location (racking, cabling, etc.), uses the Lenovo provided Registration Utility application which they have installed on their laptop and go through the process of connecting to each ThinkEdge server's

In the Registration utility, the Field Technician loads the Registration image and selects the Site location for the server that they are onboarding. This activity needs to be performed for all servers to be registered. In the figure below you can see the result of the server registration process:

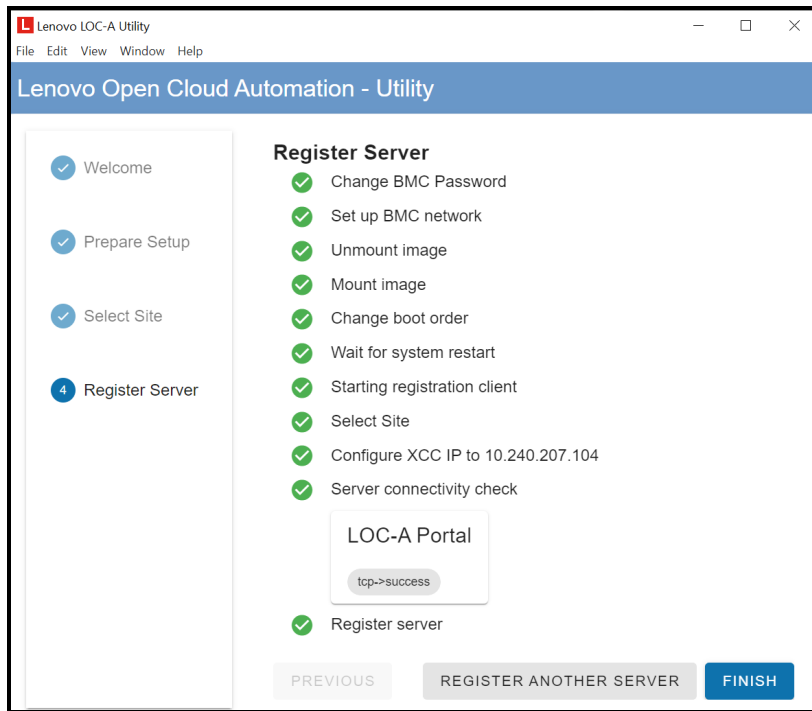


Figure 11. Registration Utility – successful server registration

### Device on-boarding using XCC IP

For brownfield scenarios where the BMC networking and credentials have already been configured, the field technician can use the same Lenovo provided Registration Utility to onboard the respective devices, following a similar process while using the prior knowledge of the IP and credentials of the pre-configured BMC.



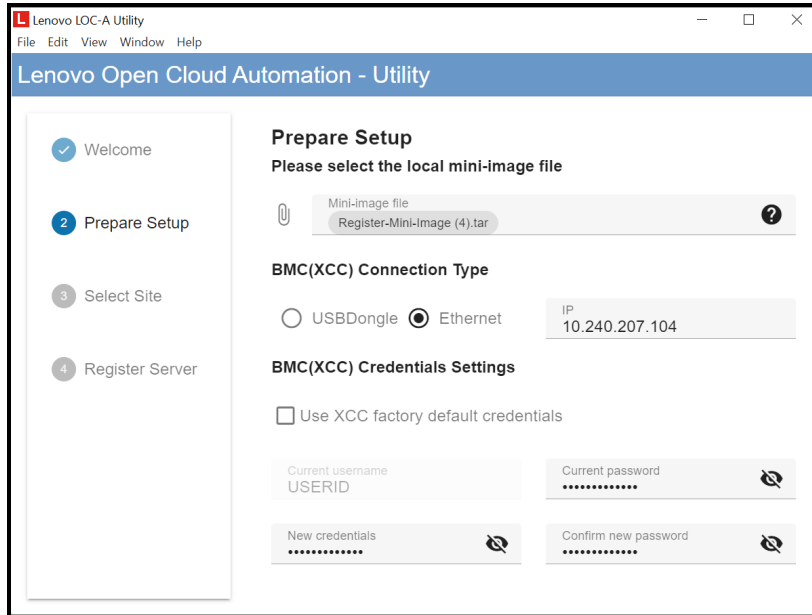


Figure 12. Registration Utility configuration for IP-based device on-boarding

### Device on-boarding using XLSX file

When brownfield deployments involve large number of nodes, Lenovo Open Cloud Automation also offers the possibility of on-boarding servers in bulk using an excel file that the cloud admin uploads by navigating to the Registered Devices page. Below you can see an example of that modal.

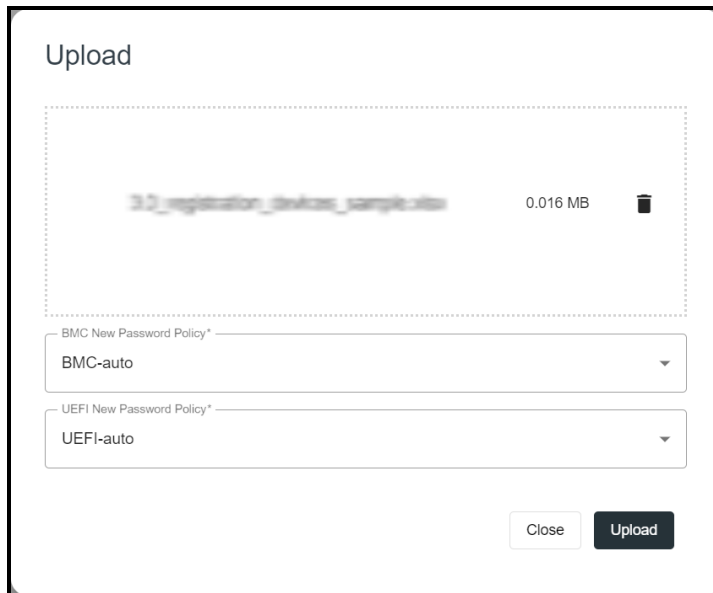


Figure 13. Bulk upload for IP-based server on-boarding

### Device on-boarding using SLP-based discovery

For scenarios where LOC-A has access to the Layer 2 BMC network, to ease the bulk server registration, Lenovo Open Cloud Automation offers the possibility that the Infrastructure Admin Team can discover servers using SLP (Service Location Protocol) and do the on-boarding with a few clicks, for further provisioning.

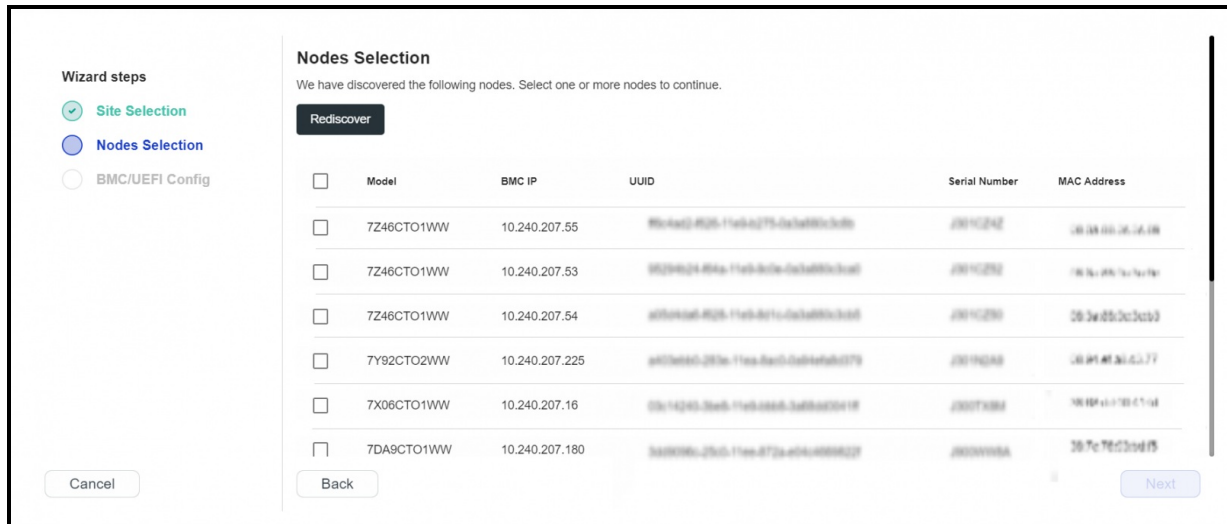


Figure 14. Edge server discovery from Lenovo Open Cloud Automation portal

### Edge Cluster Instance lifecycle management (creation, expansion, deletion)

The next step after Edge servers have been on-boarded is to create a cluster instance. This step requires that the correct meta-data has been populated for that cluster, that an Edge Cluster deployment template has been created and that the minimum number of Edge Servers needed for 1 instance have been on-boarded into Lenovo Open Cloud Automation. The Edge Cluster instance is created by navigating to the Instances page and by clicking the “+” button in the top right corner of the page.

The Infrastructure Admin Team selects the template that they want to instantiate, in what sites they want to instantiate it (they can select all the sites that are considered “Ready for deployment” and meet the conditions mentioned above) and then hits deploy.

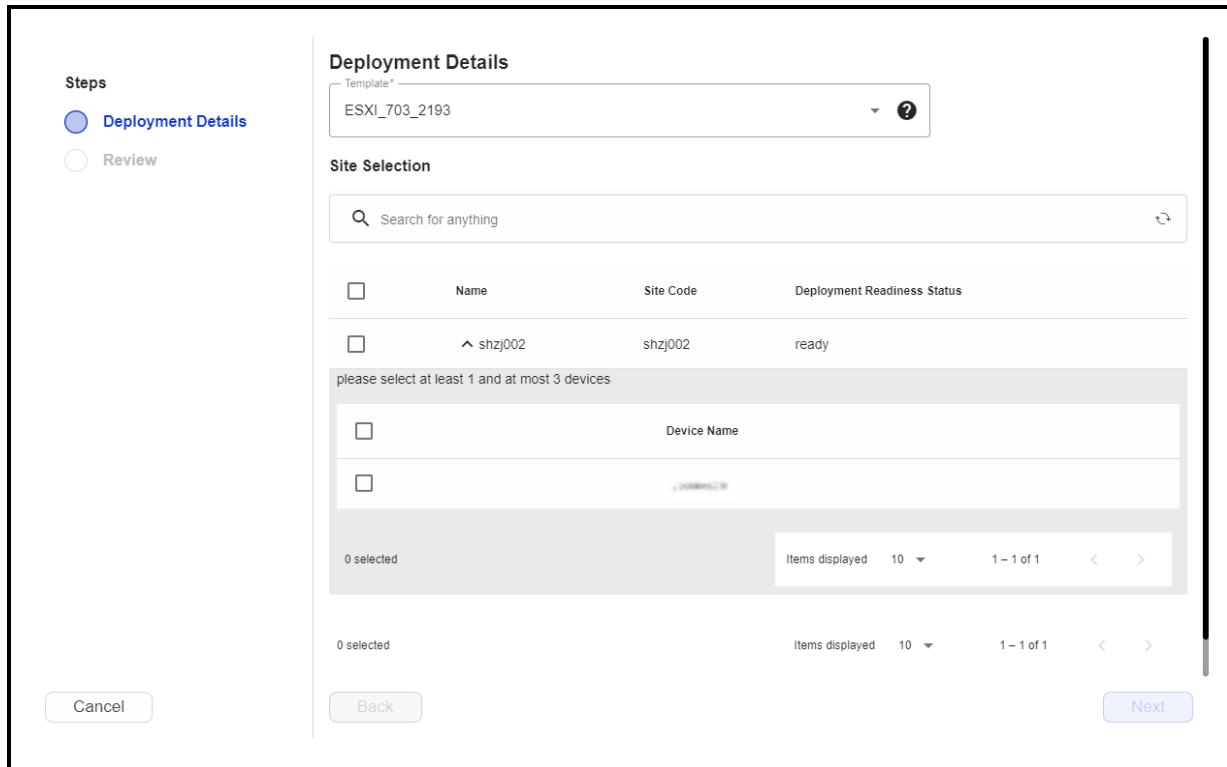


Figure 15. Instance creation, first step – template, site, and device selection

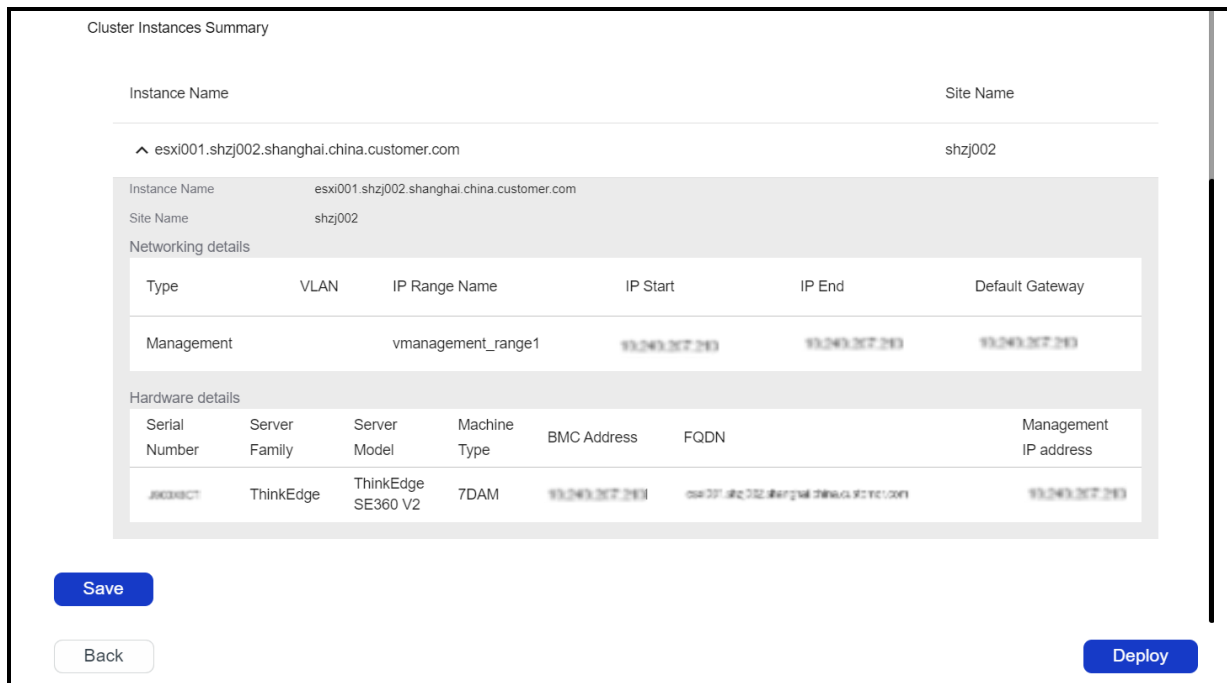


Figure 16. Instance creation, 2nd step – Meta-data population, IP address assignment and device selection details

## Benefits to users

Benefits to users include the following:

- Designed around a simplified logistic process where ThinkEdge servers are deployed straight to the edge location, skipping staging environment.
- Offers a web-based portal to on-board and deploy distributed Edge infrastructure.
- Standard and plugin-based deployment flavors for many common Edge use cases and scenarios.
- Lenovo pre-defined deployment templates and device profiles to provide an intuitive starting point for the Infrastructure Admin Team.
- Centralized credential policies that give the Infrastructure Admin Team control over how the infrastructure is secured.
- Onboard and deploy edge infrastructure across multiple sites simultaneously, from a single, simple workflow.
- Operating System image side-loading during device on-boarding to avoid transferring images during deployment.
- Simplified server onboarding process for the field-technician
- Task progress monitoring and log download

## Usage workflows

The following figure is the workflow illustrating the main actions that can be performed and by what persona when nZTP is used:

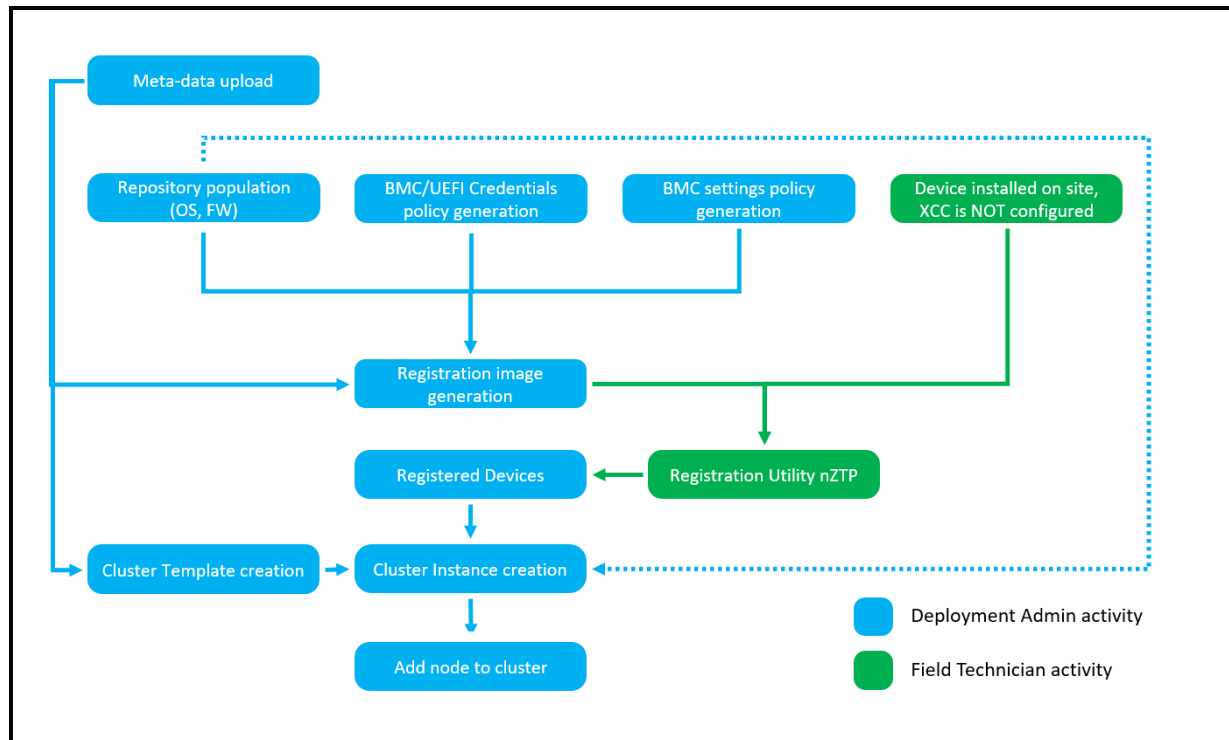


Figure 17. LOC-A user flow - greenfield deployment scenario – nZTP

For greenfield scenarios where server discovery with SLP is used, the user workflow is illustrated below:

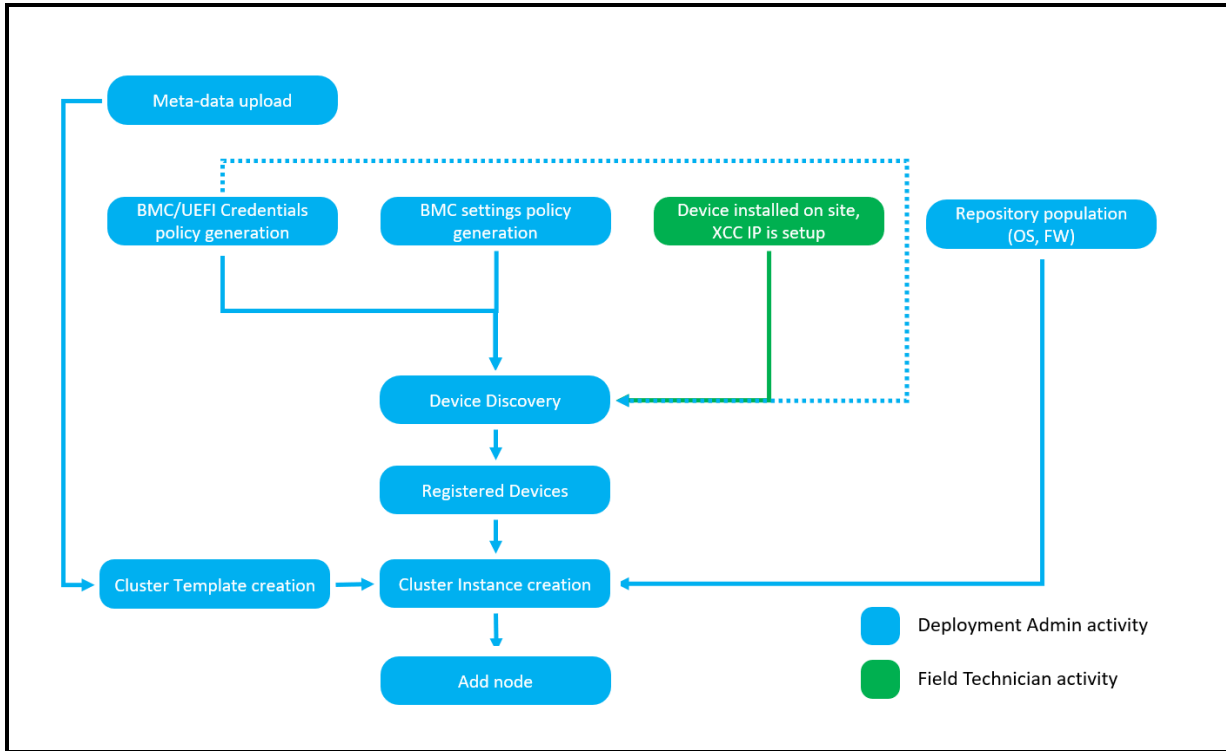


Figure 18. LOC-A user flow - greenfield deployment scenario – SLP discovery

In case of brownfield scenarios, but with field technician doing site surveys and on-boarding of devices with the Registration Utility, the user workflow is shown below.

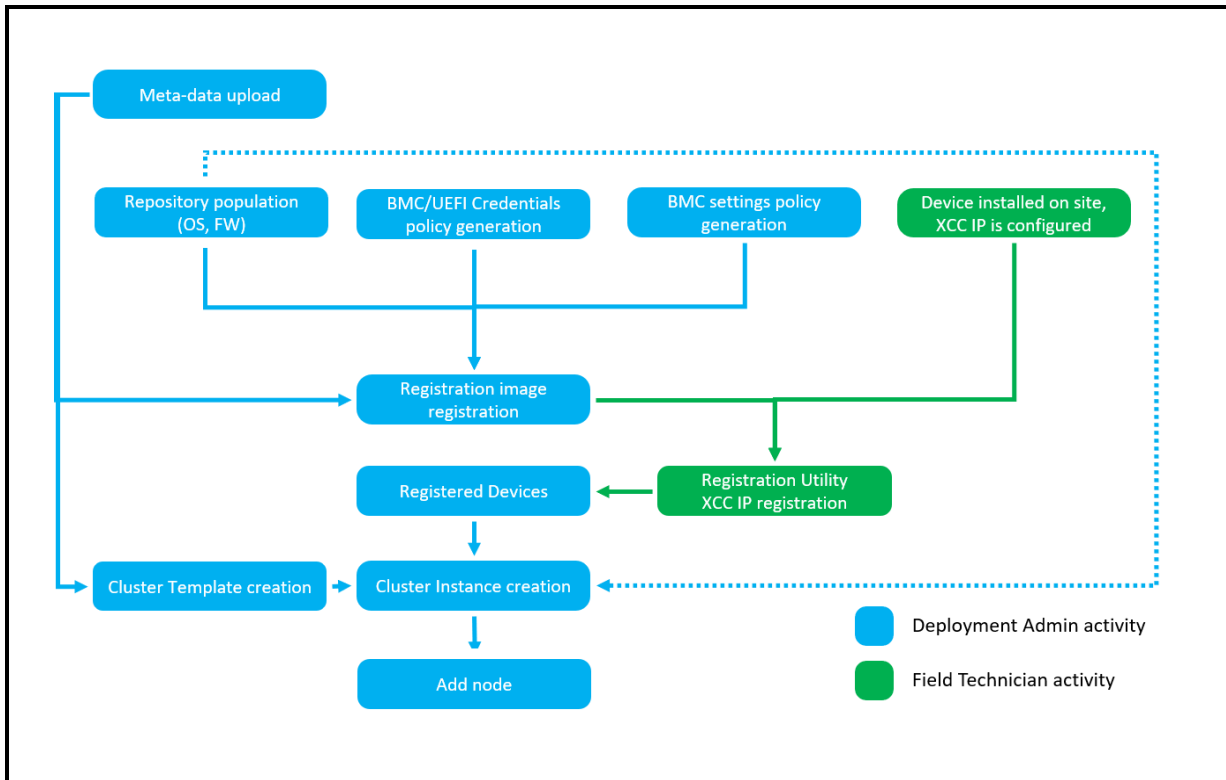


Figure 19. LOC-A user flow - brownfield deployment scenario – IP based on-boarding

For brownfield scenarios where the devices are already in the field and the Infrastructure Admin Team has an accurate inventory of the devices that contains information such as Serial Number, BMC IP address, Location, etc for each server, the user flow looks like in the below diagram.

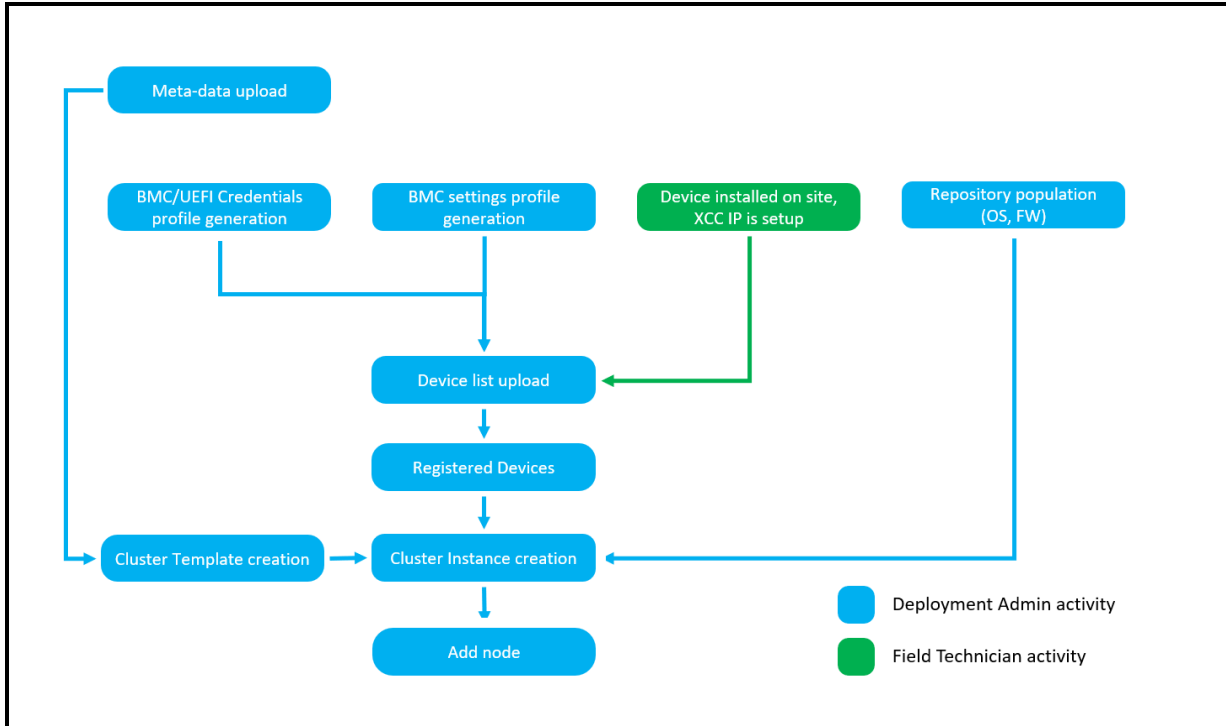


Figure 20. LOC-A user flow – brownfield deployment scenario – Inventory file based on-boarding

## Features for LOC-A administrators

For LOC-A administrators the Setup page contains a few tabs that enable the administrators to perform tasks relevant for the day-to-day maintenance of the platform.

- **Vault Configuration** – In the Vaults tab, from the Setup page, Administrators can add, edit or delete their organization’s vaults. A secrets’ vault is used by LOC-A to both read already defined credentials for existing cloud services, and also to store the credentials for OS, BMC or UEFI which LOC-A automatically generates when the credentials policy is set to \*-auto.
- **User Administration – Local User Management** – provides the option to create, edit, search for, and delete local users.
- **User Administration - LDAP configuration** – Administrators can enable and configure their LDAP server parameters as well as the priority policy for Local vs. LDAP defined users.
- **Tasks Management** - The Tasks page displays a list of tasks executed by the users and their status. When clicking on an individual task the user can see details about the task such as the start- and end-time, the detailed error cause if the task is failed, the sub-tasks performed and their execution time. From the same modal that displays the details for a particular task, the user can download the logs specific for that task which helps narrow the investigation scope in case troubleshooting is needed. If logs for all tasks are needed, the user can download them from the main Tasks page where all the tasks are listed.

## Deployment

Lenovo Open Cloud Automation comes as an OVA appliance that can be easily on-boarded using the customer's vCenter. Other Delivery options include a QCOW2 appliance. There is no feature-function distinction between the 2 delivery methods and the customer can choose based on their existing infrastructure. In terms of scalability, underneath the covers, the two appliances make use of Kubernetes, and the appliances can be scaled both vertically (increase number of allocated resources) but also horizontally (by creating more appliances that are used as worker nodes by the embedded Kubernetes cluster).

## Subscription and support

Lenovo Open Cloud Automation is enabled through a per-node subscription and support entitlement model, which once entitled for the all the nodes contained within the cluster, gives the customer access to LOC-A software updates and Lenovo support for the duration of the acquired term.

Lenovo will provide interoperability support for all software tools defined as validated with LOC-A, and development support (Level 3) for specific Lenovo-supported tools only. Open source and supported-vendor bugs/issues will be logged and tracked with their respective communities or companies if desired, with no guarantee from Lenovo for bug fixes. Full support details are provided at the support links below for each respective version of LOC-A. Additional support options may be available; please contact your Lenovo sales representative for more information.

- Lenovo provides support in English globally and in Chinese for China (24x7)
- Support response times are as follows:
  - Severity 1 issues response is 1 business day.
  - Other issues: 3 business days

LOC-A has 1-year lifecycle for each release, customer should upgrade to the latest version if out of support. New versions of LOC-A are released quarterly, usually in the last working day of the quarter.

The following table lists end of support for LOC-A versions.

Table 2. End of support list

Version	Date (mm/dd/yyyy)
LOC-A 3.0	3/29/2025
LOC-A 2.9	1/22/2025
LOC-A 2.8	9/28/2024
LOC-A 2.7	6/28/2024
LOC-A 2.6	4/7/2024
LOC-A 2.5	9/30/2023
LOC-A 2.4	4/29/2023
LOC-A 2.3	10/18/2022
LOC-A 2.2	6/15/2022
LOC-A 2.1	1/15/2022
LOC-A 2.0	6/30/2021
LOC-A 1.0 for VCF	3/24/2021

## Validated software components

The LOC-A Core Framework appliance provides a self-contained image, for quick installation, that contains all the services required to do the automated cloud deployment and management for edge sites. The services within the image run as services on top of a built-in K3S cluster.

Each LOC-A software release is validated against a defined configuration of software tools and Lenovo systems, to make deployment more straightforward and enable support. Other management tools, hardware systems and configurations outside the defined stack may be compatible with LOC-A, though not formally supported; to determine compatibility with other solutions, please check with your Lenovo sales representative.

The following software components are validated by Lenovo as part of the overall LOC-A software solution entitlement:

- **Inventory Service (LIS)** The Inventory service is the source-of-truth for the infrastructure that handles planning data and edge site resources, including sites, IP addresses and VLANs, cloud services, network services, and the cloud objects, such as tenants and clusters. The metadata for resources can be imported or created by users in the planning phase.
- **Configuration Service (LCS)** The Configuration service is an execution orchestrator built on AWX. LOC-A LCS is configured with predefined automation workflows and job templates that make managing the infrastructure easy and efficient.
- **Hardware Management Service (LMS)** The Hardware Management service helps to provision hardware and performs hardware management operations during the lifecycle of Lenovo servers. LOC-A includes Confluent and Lenovo OneCli as components of its Hardware Management Service. LMS is responsible for:
  - Server inventory
  - Server power operations (as required by the deployment process, not on-demand)
  - Server operating system deployment
  - Server firmware updates
  - Server configuration

Technical Support for software means the provision of telephone or web-based technical assistance by Lenovo to the Customer's technical contact(s) with respect to any software defects, errors, and product problems exhibited on supported Lenovo configurations.

Technical support does not cover help with the initial installation of the product, software how-to, training, and or configuring the production environment. Please contact your local Lenovo Sales Representative or Lenovo Business Partner for the best service offerings if you need assistance in these areas.



## Supported ThinkEdge servers

LOC-A seamlessly integrates with Lenovo ThinkEdge servers, offering robust support for Lenovo hardware within the cluster.

The following Lenovo systems are supported to be deployed with LOC-A:

- **ThinkSystem SE350** – The ThinkSystem SE350 is a purpose-built server that is half the width and significantly shorter than a traditional server, making it ideal for deployment in tight spaces. It can be mounted on a wall, stacked on a shelf or mounted in a rack. The ThinkSystem SE350 puts increased processing power, storage and network closer to where data is generated, allowing actions resulting from the analysis of that data to take place more quickly. For more information, see the [SE350 product guide](#).
- **ThinkEdge SE350 V2** – The SE350 V2 is a purpose-built server that is half the width and significantly shorter than a traditional server, making it ideal for deployment in tight spaces. It can be mounted on a wall, ceiling mount, placed on a shelf or mounted in a rack. The ThinkEdge SE350 V2 server puts increased processing power, storage, and network closer to where data is generated. For customers that want to install server outside data center looking for reduced latency by processing at the edge. For more information, see the [SE350 V2 product guide](#).
- **ThinkEdge SE360 V2** – The SE360 V2 is a purpose-built server that is a 2U high and half width making it significantly smaller than a traditional server, ideal for deployment in tight spaces. It can be mounted on a wall, ceiling, placed on a desk or mounted in a rack. The ThinkEdge SE360 V2 server puts increased processing power, storage and network closer to where data is generated. For customers that want on-premise deployments and outside data center looking for reduced latency by processing at the edge. For more information, see the [SE360 V2 product guide](#).
- **ThinkEdge SE450** – The ThinkEdge SE450 is a single-socket server with a 2U height and short depth case, making it suitable for deployment in shallow cabinets. It is based on the 3rd Gen Intel Xeon Scalable processor. The SE450 can be mounted on a wall, stacked on a shelf or mounted in a rack. The server puts increased processing power, storage and network closer to where data is generated, allowing actions resulting from the analysis of that data to take place more quickly. For more information, see the [SE450 product guide](#).
- **ThinkEdge SE455 V3** - The ThinkEdge SE455 V3 is a single-socket server with a 2U height and short depth case, making it suitable for deployment in shallow cabinets. It can be mounted in a 2-post or 4-post rack. The SE455 V3 uses the new AMD EPYC 8004 Series "Siena" processors for an ideal mix of performance and power efficiency. The SE455 V3 puts processing power, storage and network closer to where data is generated, allowing actions resulting from the analysis of that data to take place more quickly. For more information, see the [SE455 V3 product guide](#)

Additional Lenovo ThinkSystem servers may be compatible with LOC-A. Contact your Lenovo sales representative for more information.

## Client PC requirements

A web browser is used to access the LOC-A portal. To fully utilize LOC-A's capabilities, the client PC should meet the following specifications:

- Hardware: CPU of 2.0 GHz or above and 4 GB or more of RAM
- Display resolution: 1280 x 800 or higher
- Browser: Chrome (v62.0 or higher) or Firefox (v56.0 or higher) is recommended

A Microsoft Windows® based laptop or tablet is required to be used by the field technician when running the Lenovo Open Cloud Automation Registration Utility. Since the utility is delivered as a fully stand-alone, self-contained application, it does not have any dependencies on the Operating System and the requirements on the laptop specifications are as follows:

- Hardware: CPU of 1.8 GHz or above and 2 GB or more of RAM; RJ45 port or USB to Ethernet adapter for connecting with an Ethernet cable to the server XCC RJ45 port.
- Display resolution: 640 x 480 or higher
- Browser: N/A – the utility does not run in a browser, it's a native application

## Seller Training Courses

The following sales training courses are offered for employees and partners (login required). Courses are listed in date order.

1. **Lenovo Open Cloud Automation Education**  
2023-03-24 | 25 minutes | Employees and Partners

In this course we will cover an overview of what is Lenovo Open Cloud Automation. By the end of this course, you will be able to describe the benefits of the Lenovo Open Cloud Automation, identify the user extensible technologies that LOC-A integrates, and explain why a customer should choose the Lenovo Open Cloud Automation. The course was last updated in March 2023.

Published: 2023-03-24  
Length: 25 minutes  
Employee link: [Grow@Lenovo](#)  
Partner link: [Lenovo Partner Learning](#)  
Course code: DLOCO101r2

2. **Partner Technical Webinar - Smart Cities with ThinkEdge**  
2022-12-13 | 60 minutes | Partners Only

In this 60-minute replay, Alejandro Perez, Lenovo Americas Edge Solutions Leader, discussed the Edge applications of Smart Cities, specifically the City of Barcelona. Alejandro reviewed the key uses of the solution and the underlying architecture including the Lenovo SE350, SE450 and Lenovo Open Cloud Automation software.

Published: 2022-12-13  
Length: 60 minutes  
Partner link: [Lenovo Partner Learning](#)  
Course code: 120922

3. **Technical Champions Webinar: Accelerate Edge Deployments with Lenovo Open Cloud Automation**  
2022-10-26 | 59 minutes | Employees Only

When speaking with customers about Edge Computing it is easy to be dragged down the path of speeds, feeds, and costs. View this webinar and pivot the conversation to focus on Lenovo Cloud Automation and faster deployment at the Edge.

Webinar topics include:

- Lenovo Edge Portfolio
- Ways to deploy Edge
- Proposed streamlining deployment with LoC-A
- Demo

Published: 2022-10-26  
Length: 59 minutes  
Employee link: [Grow@Lenovo](#)  
Course code: DLOC200

## Related links

For more information, see the following resources:

- LOC-A web page:  
<https://www.lenovo.com/us/en/servers-storage/software/open-cloud-automation/>
- LOC-A Support website:  
<https://support.lenovo.com/us/en/solutions/HT509884>
- Lenovo DCSC configurator:  
<https://dcsc.lenovo.com>

## Related product families

Product families related to this document are the following:

- [Edge Computing and Internet of Things \(IoT\)](#)
- [Edge Servers](#)
- [Lenovo XClarity](#)

## Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.  
8001 Development Drive  
Morrisville, NC 27560  
U.S.A.  
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2024. All rights reserved.

This document, LP1974, was created or updated on June 21, 2024.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:  
<https://lenovopress.lenovo.com/LP1974>
- Send your comments in an e-mail to:  
[comments@lenovopress.com](mailto:comments@lenovopress.com)

This document is available online at <https://lenovopress.lenovo.com/LP1974>.

## Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkAgile®

ThinkEdge®

ThinkSystem®

The following terms are trademarks of other companies:

AMD and AMD EPYC™ are trademarks of Advanced Micro Devices, Inc.

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Microsoft® and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

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