



Lenovo Hybrid AI 221 Microfactory with Red Hat Product Guide

Lenovo Hybrid AI 221 provides a compact and low-cost entry point for AI enterprise workloads. In partnership with Red Hat and NVIDIA, Lenovo has developed a full AI platform with minimal setup and configuration requirements. Customers can expect support for all layers of the stack, from hardware up to application layer, so that AI engineers can focus on development, not infrastructure management.

Business Problem

Enterprises want to begin with a right-sized infrastructure footprint but do not want to adopt a dead-end architecture. They need a small deployment that still looks and behaves like the platform they plan to scale in the future, without dealing with the complexities of a large-scale deployment up front. They also want a software stack that can support inference today while leaving room for model customization, lifecycle management, and broader application integration tomorrow.

Solution Overview

Red Hat AI Enterprise with Validated Patterns are central to this Hybrid AI 221 solution with Red Hat, making the AI application layer setup quick and easy with minimal configuration to deploy a tested and maintained GitOps pattern to an existing OpenShift cluster. Automation is the key point of this solution – both the OpenShift cluster bring-up and Validated Pattern installation have automation built in to ensure users can get started as fast as possible. Furthermore, OpenShift GitOps ensures that deployments are consistent and repeatable, because all software artifacts share a single source of truth in a git repository. As the Validated Pattern repository is updated the live pattern deployment will automatically have a rolling upgrade. Red Hat provides a large and growing catalog of [patterns](#) for PoCs or deployments.

The Hybrid AI 221 solution with Red Hat also provides flexibility with OpenShift and OpenShift AI, allowing easy model deployments and workload scheduling/management for any cloud-native workload. Combined with the rich ecosystem and feature set of OpenShift, this solution provides the full stack that an enterprise needs to run modern AI workloads and systems.

For a more information on the Lenovo Hybrid AI 221 platform, please see the [Lenovo Hybrid AI 221 Platform Guide](#).

<u>Application Layer</u> Red Hat Validated Patterns	Red Hat Validated Patterns are a blueprint for automated deployments of ready to run software on OpenShift, utilizing GitOps for deployment management.
<u>AI Platform & Scheduler</u> OpenShift AI	With OpenShift AI, users can perform model serving, fine tuning, and training. The platform includes curated third-party models to get enterprise started with AI quickly.
<u>Container Orchestration</u> Red Hat OpenShift	OpenShift is an enterprise grade Kubernetes platform for deploying applications securely and at scale with minimal configuration or management.
<u>Infrastructure Software</u> NVIDIA AI Enterprise	NVIDIA AI Enterprise includes critical software components such as the GPU Operator, Multi-Instance GPU, and GPU Drivers.
<u>Infrastructure</u> Hybrid AI 221 Platform	Hybrid AI 221 on the ThinkSystem SR650a V4 server is a 2U AI platform for inference workloads, with 2 CPUs, 2 GPUs, and 1 Network Card.

AI Software Stack

Deploying AI to production involves implementing multiple layers of software. The process begins with the system management and operating system of the compute nodes, progresses through a workload or container scheduling and cluster management environment, and culminates in the AI software stack that enables delivering AI agents to users.

Red Hat AI Enterprise

Red Hat AI Enterprise is an integrated AI platform for deploying and managing efficient and cost-effective AI models, agents, and AI-powered applications across hybrid cloud environments. Powered by Red Hat OpenShift, it brings together capabilities for model tuning, high-performance inference, AI application development, agentic workflow management, and lifecycle operations in a tested and supported enterprise stack. For the Lenovo Hybrid AI 221 Microfactory with Red Hat, Red Hat AI Enterprise bundles Red Hat OpenShift Container Platform, Red Hat OpenShift AI, and AI Accelerators to provide the core software foundation for deploying, serving, and managing inference-focused AI workloads with greater efficiency, flexibility, and control.

Red Hat OpenShift

Red Hat OpenShift Container Platform provides the enterprise Kubernetes foundation for this offering. As part of Red Hat AI Enterprise, OpenShift helps organizations deploy, run, and manage containerized AI applications and workloads consistently across hybrid cloud environments. Its orchestration, cluster management, and integrated platform capabilities, including the Red Hat Enterprise Linux CoreOS foundation used by OpenShift clusters, help support secure operations, automated lifecycle management, and consistent infrastructure for inference-focused workloads.

OpenShift AI

Red Hat OpenShift AI provides a comprehensive AI platform for building, tuning, deploying, serving, and monitoring predictive models, generative AI applications, and agentic systems at scale across hybrid cloud environments. As part of Red Hat AI Enterprise, OpenShift AI helps accelerate AI innovation and improve operational consistency through flexible inference, model serving, data connection workflows, observability, and governance capabilities. In this solution, OpenShift AI helps teams move inference-focused AI workloads from isolated experimentation into governed, enterprise-ready operations. For GPU-accelerated use cases, it supports seamless access to NVIDIA GPUs and leverages certified operators to simplify driver and runtime management. This makes it an ideal foundation for organizations looking to operationalize AI across hybrid cloud environments with consistency and control.

Red Hat Validated Patterns

Red Hat Validated Patterns are deployment assets that help accelerate implementation by providing repeatable, GitOps-based solution patterns for complex hybrid cloud environments. In this solution, a validated pattern helps codify the platform, application, and configuration elements needed for a more consistent deployment experience. Managed through Red Hat OpenShift GitOps, validated patterns use Git-based workflows to define and maintain desired-state configuration, helping teams deploy and update the AI software stack more consistently over time.

AI quickstarts

AI quickstarts provide ready-to-run example use cases that teams can deploy, explore, and extend on Red Hat OpenShift. Examples may include retrieval-augmented generation chatbots, product recommendation engines, or assistant-style applications. Because quickstarts are built on OpenShift, they can be deployed within this offering and adapted for customer-specific data, workflows, and application needs. Quickstarts are best positioned as starting points for experimentation and solution development, while Red Hat Validated

Patterns provide a more structured, GitOps-based approach for repeatable deployment.

More on Validated Patterns

Validated Patterns can be used to deploy complex, multi-product solutions where correctness and repeatability are key. They are equally well-suited for proof-of-concept demonstrations due to their ease of deployment. Some example patterns include Multicloud GitOps, Hybrid data services, Industrial Edge, and RAG. These patterns can be used out of the box or customized to fit a use case or PoC.

What separates Validated Patterns with GitOps from traditional architectures is maintenance being baked into the framework as a core feature. There is no need to manually update the deployment, as all software artifacts will constantly sync with the connected git repository. As versions change, APIs degrade, and tools become outdated, the git repository can be updated and the changes will be automatically reflected in the Validated Pattern deployment.

GitOps delivers Infrastructure as code, where the desired state of infrastructure is determined by a codebase instead of manual implementation. Configuration files inside of a git repository outline the resources that should be provisioned, how they behave, and what properties they have. By leveraging git's built in versioning system all changes can be tracked allowing easy updating and rollback. Therefore, by committing to a GitOps deployment, developers and DevOps teams get a standard workflow for application development, increased security, auditability, and consistency. With the OpenShift GitOps operator, the GitOps practice can be directly implemented within an OpenShift cluster. All Validated Patterns implement GitOps.

Validated Patterns are a great fit for organizations who already have a clear use case and are willing to commit fully to GitOps. The catalog of Validated Patterns includes a variety of use cases which many organizations are already implementing, such as RAG, Industrial Edge, and Multicloud applications. In that case, installing the pattern is simple and can become production-ready quickly. Teams already utilizing continuous integration and continuous development pipelines can transition easily to Validated Patterns.

Deploying OpenShift and a Validated Pattern

Deploying the 221 AI Factory can be broken down into 2 main steps: deploying OpenShift and installing the Validated Pattern. These steps assume your server is networked and can reach the internet, which is required for the OpenShift Assisted Installer. Either a static IP assignment or DHCP can be used, but this deployment guide will walk through the steps for a static IP address. For a visual step by step walkthrough of deploying the SNO cluster with the Assisted Installer and deploying the Validated Pattern onto a Lenovo server, click on this [tutorial](#) link.

In the following sections, we take a deeper dive the installations:

- [Single Node OpenShift Cluster Installation](#)
- [Validated Pattern Deployment](#)

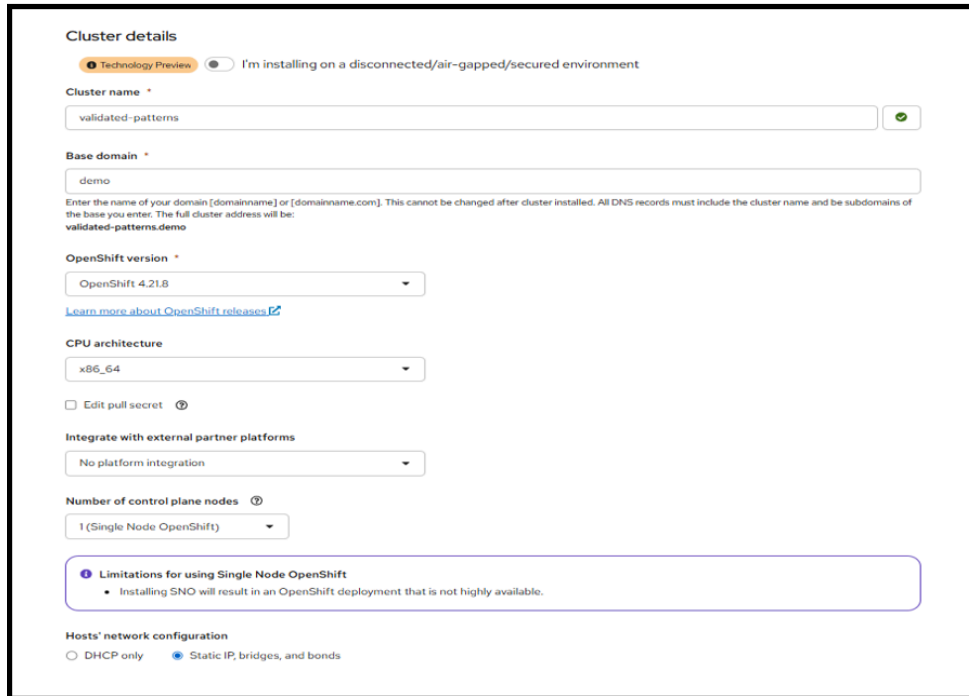
Single Node OpenShift Cluster Installation

Before you begin, if you want to use static IP addressing instead of DHCP you will need the MAC address, subnet, IP address, and DNS address of the node your cluster will be deployed to.

Ensure you are logged in and have activated your Red Hat AI Enterprise License. Navigate to <https://console.redhat.com/> and click on Red Hat OpenShift. Then, click on “Create Cluster” under Red Hat.

OpenShift Container Platform. When prompted to select the type of cluster to create, click on the Datacenter tab, and click on “Create cluster” under Assisted Installer.

The Assisted Installer page will open. In Cluster Details, ensure that the option “1 (Single Node OpenShift)” is selected under Number of control plane nodes .



The screenshot shows the 'Cluster details' configuration page. At the top, there are two radio buttons: 'Technology Preview' (selected) and 'I'm installing on a disconnected/air-gapped/secured environment'. Below this, the 'Cluster name' field contains 'validated-patterns' with a green checkmark icon. The 'Base domain' field contains 'demo'. A note below the base domain field states: 'Enter the name of your domain [domainname] or [domainname.com]. This cannot be changed after cluster installed. All DNS records must include the cluster name and be subdomains of the base you enter. The full cluster address will be: validated-patterns.demo'. The 'OpenShift version' dropdown is set to 'OpenShift 4.21.8', with a link to 'Learn more about OpenShift releases'. The 'CPU architecture' dropdown is set to 'x86_64'. There is an unchecked checkbox for 'Edit pull secret'. The 'Integrate with external partner platforms' dropdown is set to 'No platform integration'. The 'Number of control plane nodes' dropdown is set to '1 (Single Node OpenShift)'. A warning box titled 'Limitations for using Single Node OpenShift' contains the text: 'Installing SNO will result in an OpenShift deployment that is not highly available.' At the bottom, the 'Hosts' network configuration' section has two radio buttons: 'DHCP only' (unchecked) and 'Static IP, bridges, and bonds' (checked).

Figure 1. Cluster details

In this example, we assume the IP address for the node is static. If you are using DHCP, the static network configuration can be ignored.

Figure 2. Network configuration

Enter the DNS server address for your node, the subnet your node is in, and the default gateway for that subnet. On the next page, enter the MAC address and IP address.

On the host discovery page, you can generate the discovery ISO that you will boot the server with. For a simple network configuration the minimal image file should suffice. However, for a node with multiple NICs or a complicated network topology, the full image file is recommended and was used to deploy an SR675 V3 with 5 NICs in this demo.

Download the ISO. The ISO can either be flashed to a USB stick or uploaded to media mount in XClarity Controller. If using a USB stick, ensure the software being used to flash it does not alter the installation in any way, as this can cause failure to boot.

Note: If XClarity Controller is used to mount the ISO, the session must stay open for the entire cluster bring-up. This can take up to an hour or even more, so ensure that you will be able to monitor the XCC session and ensure your computer does not go to sleep or the ISO may unmount.

Once the host has booted, you will see a login prompt. No need to login to the node, you can now go back to the Red Hat Console and see the new host pop up in the menu.

Hostname	Role	Status	Discovered on	CPU Cores	Memory	Total storage	
> q0-94-72-72-42-76	Control plane node, Worker (bootstrap)	Ready	1/27/2026, 12:46:11 PM	192	1.00 TiB	719 TiB	⋮

Figure 3. Host discovery

From here, you can progress directly to the Review section of the Assisted Installer and begin cluster

installation.

Once the cluster has been installed, make sure to either add all the OpenShift hostnames to either your machine's hosts file or the DNS server (on Windows you can find the hosts file in C:\Windows\System32\drivers\etc). The installation page contains the exact IPs and hostnames that you can copy and paste into the relevant file.

Validated Pattern Deployment

Before you begin, if you want to use static IP addressing instead of DHCP you will need the MAC address, subnet, IP address, and DNS address of the node your cluster will be deployed to.

Before proceeding, please note that once a Validated Pattern is deployed, it can be difficult to uninstall due to the nature of GitOps. Depending on the exact Validated Pattern chosen, tearing down the cluster and redeploying it can be the easiest option.

You can explore the catalogue of Validated Patterns here: <https://validatedpatterns.io/patterns/>

On the OpenShift dashboard, select Ecosystem > Software Catalog.

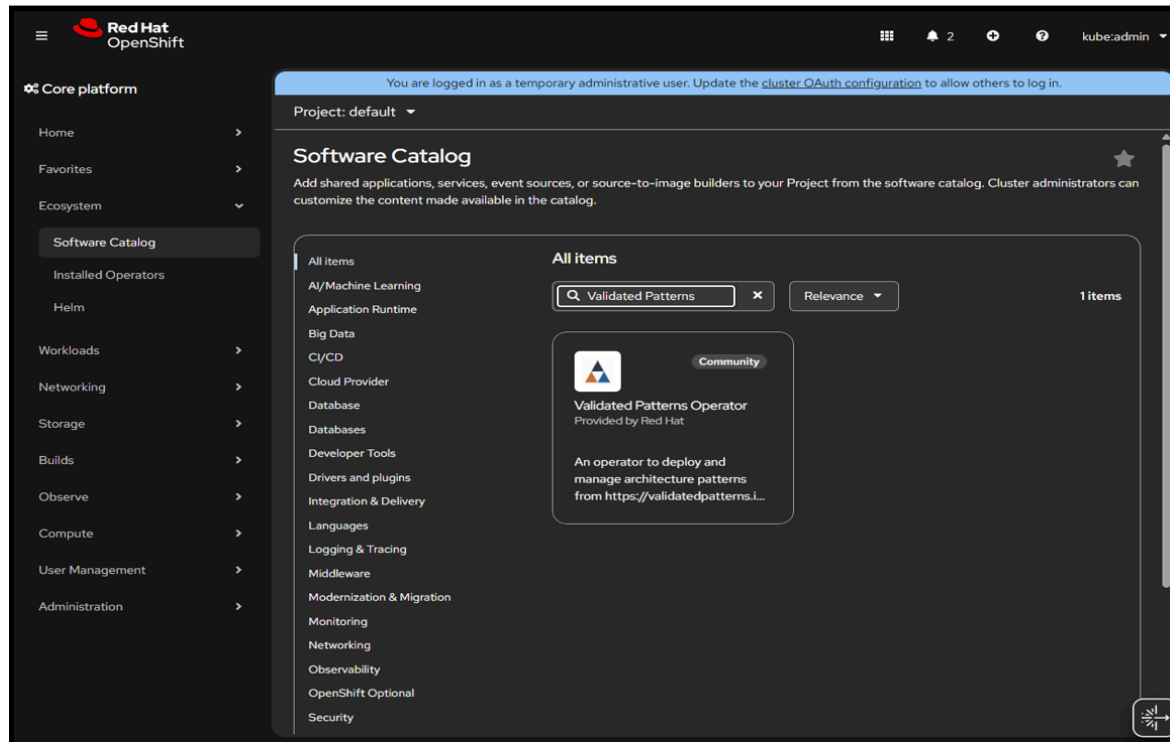


Figure 4. OpenShift dashboard

Search for the Validated Patterns Operator and install it (default settings are fine). After the operator is installed, click “View Operator”. Then, click “Create Instance” to install a Validated Pattern directly from OpenShift.

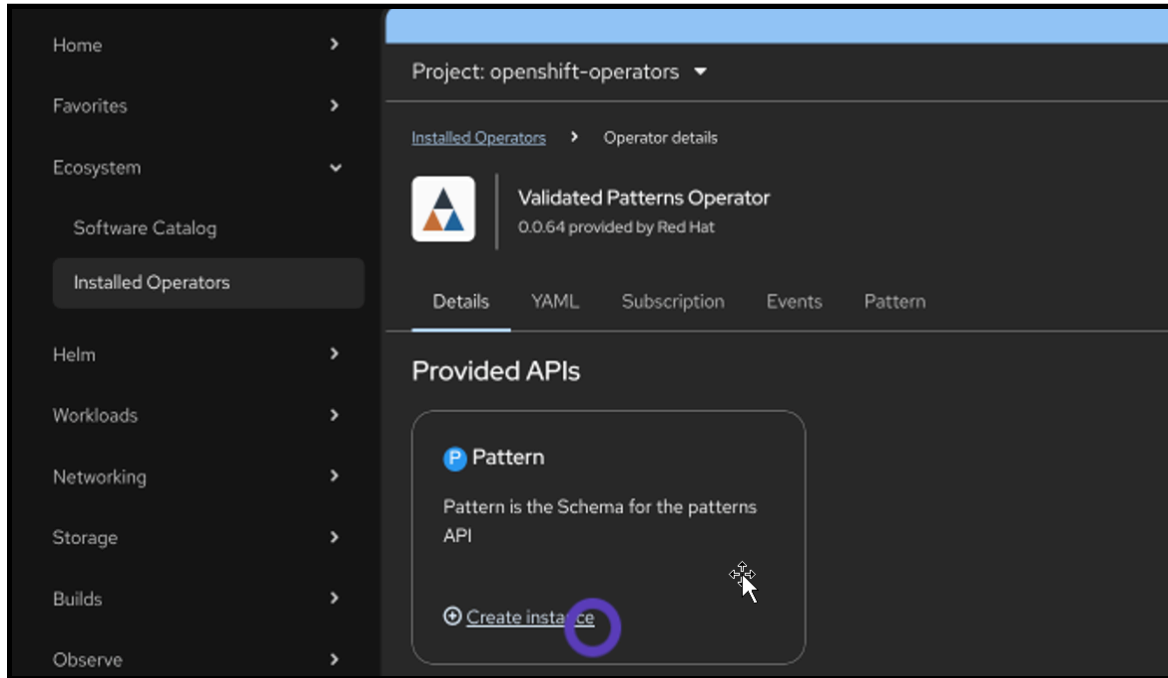


Figure 5. Validated Patterns Operator

As an example, let's take the "AI Generation with LLM and RAG" pattern. The GitHub source code is found here <https://github.com/validatedpatterns/rag-llm-gitops>.

You can name your pattern anything, however the Cluster Group Name must be specific to the pattern. From the GitHub link above, in the file values-global.yaml, you can see that the clusterGroupName is hub, this is what must be entered into the Cluster Group Name Field. Next, expand the Git Config section and change the "Target Repo" to the Git Hub url of the pattern you would like to deploy.

Finally, click Create to begin deploying the pattern. To check on progress, click the bento menu and select Prod ArgoCD.

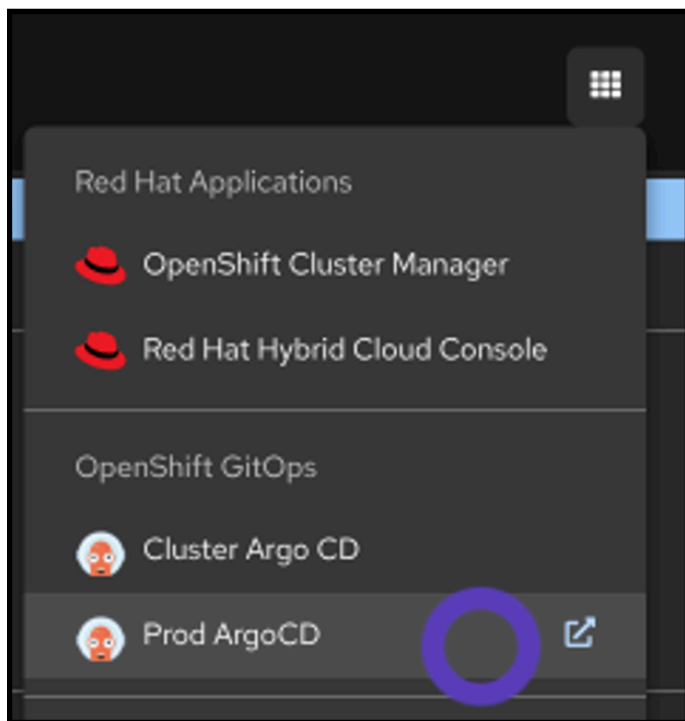


Figure 6. Bento Menu

This will show the progress of each application bring up, and once each application is completely in sync the Validated Pattern is ready to be used.

Application Example

In the following sections, we take a deeper dive into the Validated Patterns:

- [Validated Pattern: RAG LLM on Single-Node OpenShift Cluster](#)
- [Validated Pattern: OpenShift AI Install and Management](#)

Validated Pattern: RAG LLM on Single-Node OpenShift Cluster

This [Validated Pattern](#) is for an LLM Chatbot with RAG on custom documentation and an OpenWebUI-based frontend. You can interact with the LLM through the frontend for a traditional chatbot experience, or this pattern can be built upon to create custom agents drawing from internal enterprise data.

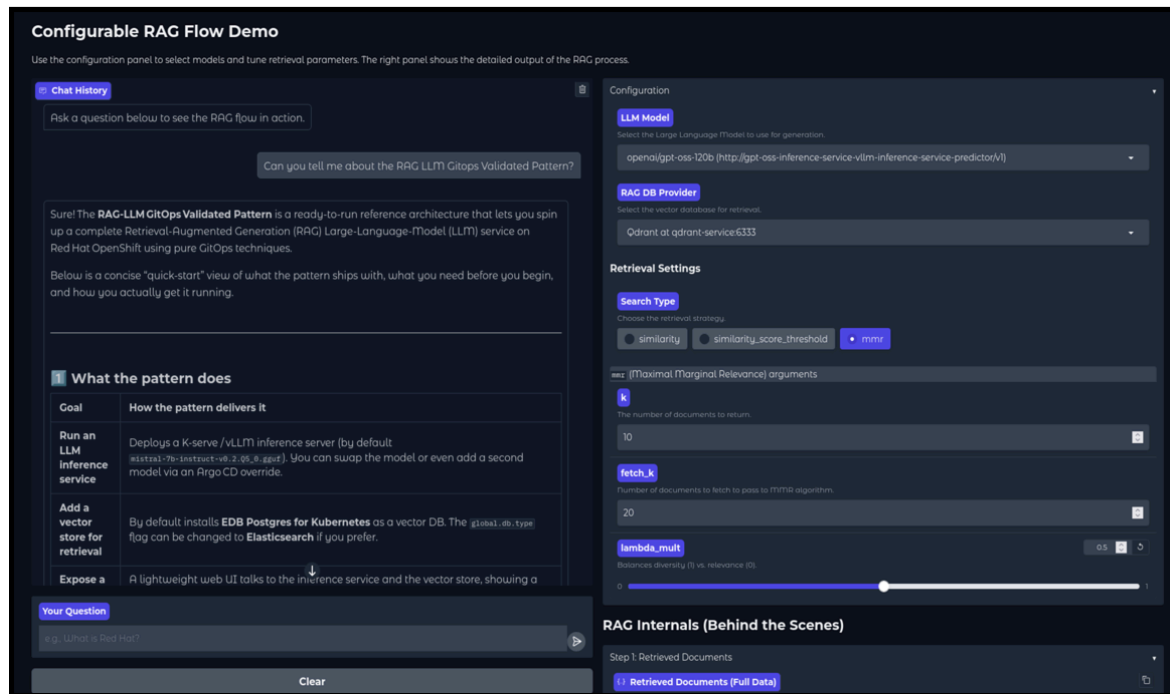


Figure 7. OpenWebUI-based frontend

The RAG LLM Pattern uses these components out of the box:

- AI Model: IBM Granite 4 Small or GPT-OSS 120B
- Inference Service: vLLM
- Orchestration: OpenShift AI
- Vector Database: Qdrant

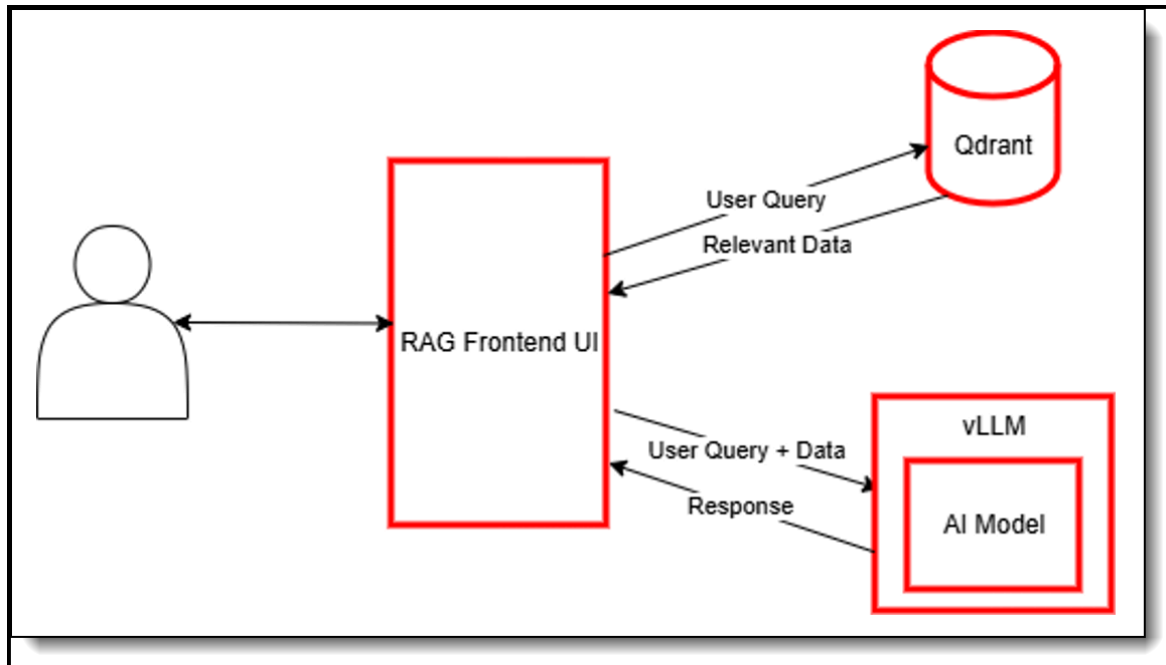


Figure 8. RAG model flow

In order to customize the model and RAG database data used, you must fork the [original repository](#) and edit the necessary yaml files. To edit the default RAG data sources for example, you must edit the `vectorEmbedJob.repoSources` and `vectorEmbedJob.webSources` fields in `overrides/vector-db-values.yaml`.

```

5   vectorEmbedJob:
6     repoSources:
7       - repo: https://github.com/validatedpatterns/docs.git
8         globs:
9           - "**/*.md"
10          - "**/*.adoc"
11    webSources: []
12    chunking:
13      size: 4096
  
```

Figure 9. GitHub repository

Either GitHub repositories or web sources can be ingested into the vector database as data sources. Adding or removing URLs from this file will cause the relevant files to be loaded into the vector database once the deployment syncs with the forked GitHub repository.

Validated Pattern: OpenShift AI Install and Management

Automating and managing OpenShift AI deployment can be made easy with the [OpenShift AI Validated Pattern](#). The quickstart involves only three simple steps: Fork the GitHub repository, clone it to the local machine, and run “./pattern.sh make install”. The installation can also be done with the Validated Patterns Operator for those who want to use the web console as demonstrated earlier in this product guide.

Furthermore, utilizing Validated Patterns with OpenShift GitOps provides a platform for unified management, which includes secrets management. This framework provides security and ensures consistency, making it much easier to dynamically manage the cluster and its security.

Ordering Through DCSC

Pre-configured Hybrid AI 221 Single Nodes can be found in [Lenovo Data Center Solution Configurator](#) (DCSC) under Deployment Ready Solutions / AI / Datacenter / Hybrid AI 221. Look for “Hybrid AI 221 Single Node Red Hat AI Enterprise” for a single SR650i V4 server with RTX Pro 6000 and the Red Hat AI Enterprise License along with NVIDIA Enterprise. The node can be ordered as-is, or customized to fit specific needs.

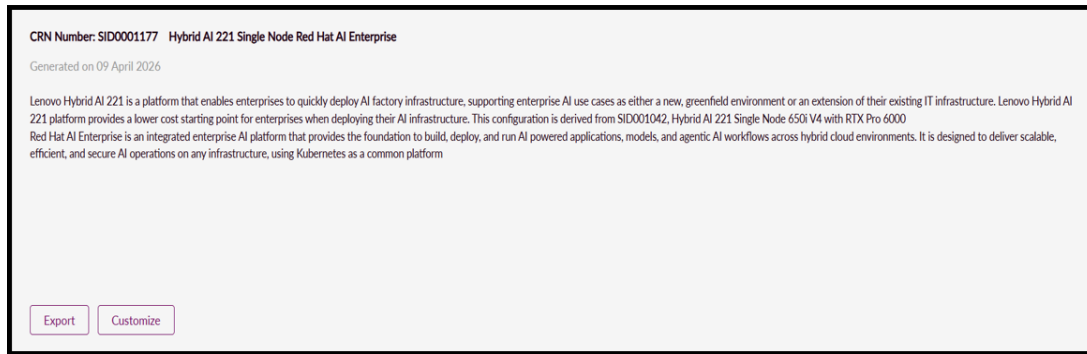


Figure 10. DCSC CRN config

Licensing

The Red Hat AI Enterprise license bundles licenses for Red Hat OpenShift Container Platform, OpenShift AI, and AI Accelerators. It is licensed on a per-node basis, and includes all components needed for the Lenovo Hybrid AI 221 platform with Red Hat.

7S0FCTOMWW	Red Hat AI Enterprise	1
SEVY	Red Hat AI Enterprise, Standard 3Yr w/Red Hat support	1

The Lenovo Hybrid AI 221 platform can also be used with NVIDIA Enterprise, which will include support for the GPU Operator used in most Validated Patterns, as well as a suite of full stack AI tools that can either integrate with OpenShift or be used standalone. NVIDIA Enterprise is priced per GPU.

7S02CTO1WW	NVIDIA Software	1
S6Z3	NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, 3 Year	2

Bill of Materials

The following table lists Bill of Materials.

Table 1. Licenses Bill of Materials

Description	Qty
ThinkSystem SR650a V4 2.5"/EDSFF 3.S Chassis	1
Intel Xeon 6530P 32C 225W 2.3GHz Processor	2
ThinkSystem V4 2U Standard Heatsink	2
ThinkSystem 64GB TruDDR5 6400MHz (2Rx4) RDIMM	8
ThinkSystem 2.5" U.2 VA 3.84TB Read Intensive NVMe PCIe 5.0 x4 HS SSD	2
ThinkSystem 2U V4 8x2.5" NVMe Backplane	1
ThinkSystem M.2 RAID B545i-2i SATA/NVMe Adapter	1
ThinkSystem M.2 ER3 960GB Read Intensive SATA 6Gb NHS SSD	2
ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port PCIe Ethernet Adapter	1
ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition 96GB PCIe Gen5 Passive GPU	2
ThinkSystem SR650/a V4 x16 Rear Direct Riser Slot 5	1
ThinkSystem 2700W 230V Platinum CRPS Hot-Swap Power Supply v2.3	2
2.5m, 16A/100-250V, C19 to C20 Jumper Cord	2
ThinkSystem 2U 6056 24K Ultra Fan Module for 600W PCIe Adapter	6
ThinkSystem Long Travel Toolless Slide Rail Kit V4 with 2U CMA	1
ThinkSystem SR650a V4 Standard Left Rack Latch	1
3m Green Cat6 Cable	2
Lenovo 3m Passive 25G SFP28 DAC Cable	2
ThinkSystem SR650 V4 Processor Board	1
ThinkSystem 1x1 2.5" HDD Filler	6
OCP3.0 Filler with screw	2
Lenovo ThinkSystem Memory Dummy	24
ThinkSystem SR650 V4 System I/O Board	1
ThinkSystem V4 CPU HS Clip	2
ThinkSystem 2U 2LP Riser Cage Filler	2
ThinkSystem SR650a V4 2U F GPU 8x25 HDD Cage	1
ThinkSystem SR650 V4 Main air duct mylar for Front AC GPU	1
ThinkSystem 2U Main Air Duct	1
ThinkSystem 2U FGPU Riser Cage for 600w	2
ThinkSystem 2U Front Double Width Air Duct for 600w	1
ThinkSystem 2U V4 3FH Riser Cage	1
ThinkSystem 2U 3FH Riser Cage Filler	1
ThinkSystem M.2 Signal & Power Cable ,ULP 82P-SLX4/2X10 SB, 400/400mm	1
ThinkSystem GPU Power Cable, PCIe16-PCIe16, 400mm	2
Power Cable, 2x3+6 P-MTK PWR, 900mm	1
Think System,PCIe Gen5 Cable, MTK 74-MCIOx8, 800mm	2
ThinkSystem, PCIe 5.0 Cable, MCIO 8X STR TO MCIO 8X STR, 550mm	2
NVIDIA Software	1

NVIDIA Enterprise (NVIDIA AI Enterprise and NVIDIA Omniverse Enterprise) Subscription per GPU, 3 Year	2
XClarity Controller Prem-FOD	1
Lenovo XClarity XCC3 premier - FOD	1
Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1
Red Hat AI Enterprise	1
Red Hat AI Enterprise, Standard 3Yr w/Red Hat support	1

Seller training courses

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

1. **Excerpts from VTT: Systems and Data Center Cooling for Energy Efficient HPC and AI Computing**

2026-05-12 | 29 minutes | Employees and Partners

Watch this video to see Dr Vinod Kamath speak on systems power and cooling for HPC and AI.

Topics include:

Systems of the Future 2026-2029 - Platform Layouts and Cooling Challenges
Technologies that Deliver Energy Efficient Cooling for the System

Tags: Advanced DataCenter, Artificial Intelligence (AI), High-Performance Computing (HPC), Technical Sales

Published: 2026-05-12

Length: 29 minutes

Start the training:

Employee link: Grow@Lenovo

Partner link: [Lenovo 360 Learning Center](#)

Course code: DVHPC233

2. **Excerpts from VTT: Understanding CDUs and their role in the Data Center**

2026-05-12 | 50 minutes | Employees and Partners

Join this session to gain a clear understanding of Coolant Distribution Units (CDUs) and their role in modern liquid cooling environments. Our speaker - Matthew Ziegler will walk you through what CDUs are, how they function within solutions like Lenovo's Neptune direct water cooling architecture, and why they are essential for high density compute environments.

Published: 2026-05-12

Length: 50 minutes

Start the training:

Employee link: Grow@Lenovo

Partner link: [Lenovo 360 Learning Center](#)

Course code: DVSYS215

3. **Lenovo Infrastructure Services Overview**

2024-10-28 | 25 minutes | Employees and Partners

This e-learning course provides an overview of the full portfolio of Lenovo ISG Services for our Lenovo sellers and business partners.

At the end of this course, you will be able to:

- Provide an overview of the full portfolio of service offerings
- Describe each service of the IT lifecycle from beginning to end
- Understand the benefits for your customers

Tags: Services, Services Lifecycle

Published: 2024-10-28

Length: 25 minutes

Start the training:

Employee link: Grow@Lenovo

Partner link: [Lenovo 360 Learning Center](#)

Course code: DSVC102r5

Related publications and links

For more information, see these resources:

- **Lenovo Hybrid AI 221 Platform Guide:**
<https://lenovopress.lenovo.com/lp2313-lenovo-hybrid-ai-221-platform-guide>
- **NVIDIA Software Product Guide**
<https://lenovopress.lenovo.com/lp2311-lenovo-hybrid-ai-software-platform>
- **Implementing AI Workloads using NVIDIA GPUs on ThinkSystem Servers:**
<https://lenovopress.lenovo.com/lp1928-implementing-ai-workloads-using-nvidia-gpus-on-thinksystem-servers>
- **Making LLMs Work for Enterprise Part 3: GPT Fine-Tuning for RAG:**
<https://lenovopress.lenovo.com/lp1955-making-llms-work-for-enterprise-part-3-gpt-fine-tuning-for-rag>
- **Lenovo to Deliver Enterprise AI Compute for NetApp AIPod Through Collaboration with NetApp and NVIDIA**
<https://lenovopress.lenovo.com/lp1962-lenovo-to-deliver-enterprise-ai-compute-for-netapp-aipod-nvidia>

Related product families

Product families related to this document are the following:

- [AI Servers](#)
- [Artificial Intelligence](#)
- [Hybrid AI Factory](#)
- [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 2026. All rights reserved.

This document, LP2433, was created or updated on May 11, 2026.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2433>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2433>.

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®

ThinkSystem®

XClarity®

The following terms are trademarks of other companies:

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

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Windows® is a trademark of Microsoft Corporation in the United States, other countries, or both.

IBM® and IBM Granite® are trademarks of IBM in the United States, other countries, or both.

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