

AI on Azure Red Hat OpenShift Workshop


Martin Seman

Sr. EMEA Black Belt, Red Hat

Yury Titov

Sr. EMEA Black Belt, Red Hat





Agenda

- **14:00 Workshop Introduction:** Digital Transformation at Parasol Insurance
 - Introduction to Parasol Insurance and the Hands-on Lab
 - Current claims handling process and proposed improvements
 - Connection and Setup
- **14:15 - Hands-On Lab:** Insurance claims lab
- **17:00 – Wrap-up: feedback and conclusion**

Mini presentations:

- Collaboration between the App Developer and Data Scientist
 - Operationalize AI with Red Hat OpenShift AI & Azure services
 - RH AI: quick web ui walkthrough
 - MCP protocol
- **Complete quiz to get digital credential**

Enabling your AI apps:

Collaboration between the app
developer and data scientist



Reality of enterprise IT environments

Mixed infrastructure environments, diverse app portfolios, & limited automation

Infrastructure



Bare metal



Virtualization



Edge



Private cloud



Public cloud

Applications



AI/ML



Analytics



Serverless



Cloud-native and
microservices



Java™



.Net



ISV

Development, IT processes, and skills



Developer
tools



Pipeline
and
processes

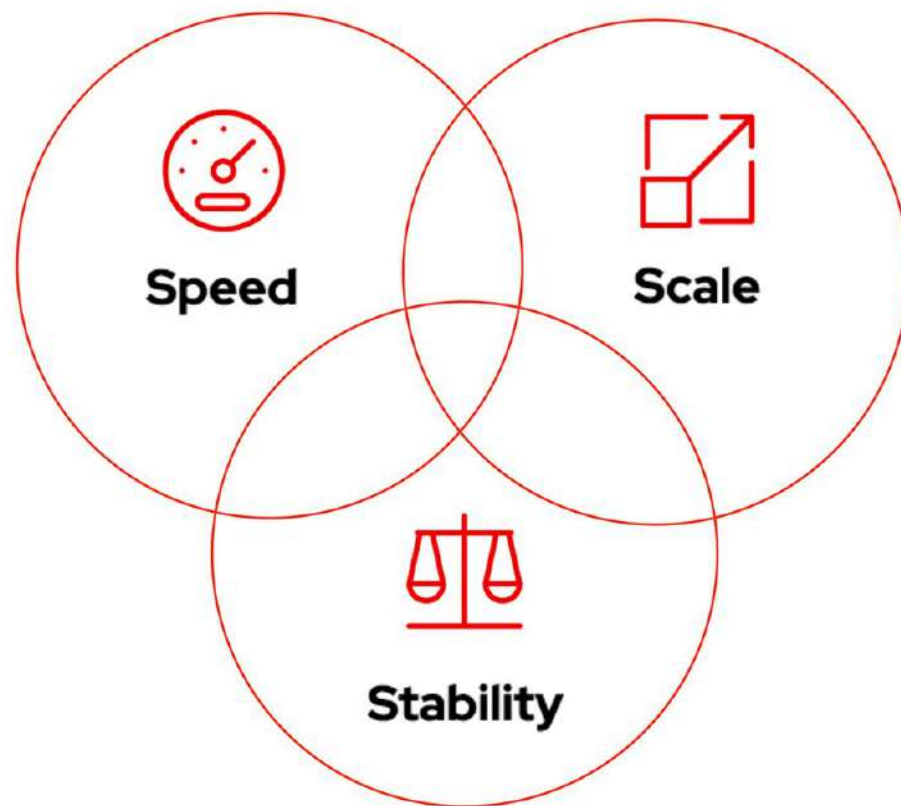


People and
policies



The right
skills

While the new world demands





100-x%

Skills

Manage workloads

Build cloud-native apps

Data-driven insights

Developer productivity

Security

Platform services

Application services

Data services

Developer services

Compliance

Kubernetes cluster services

Standardization

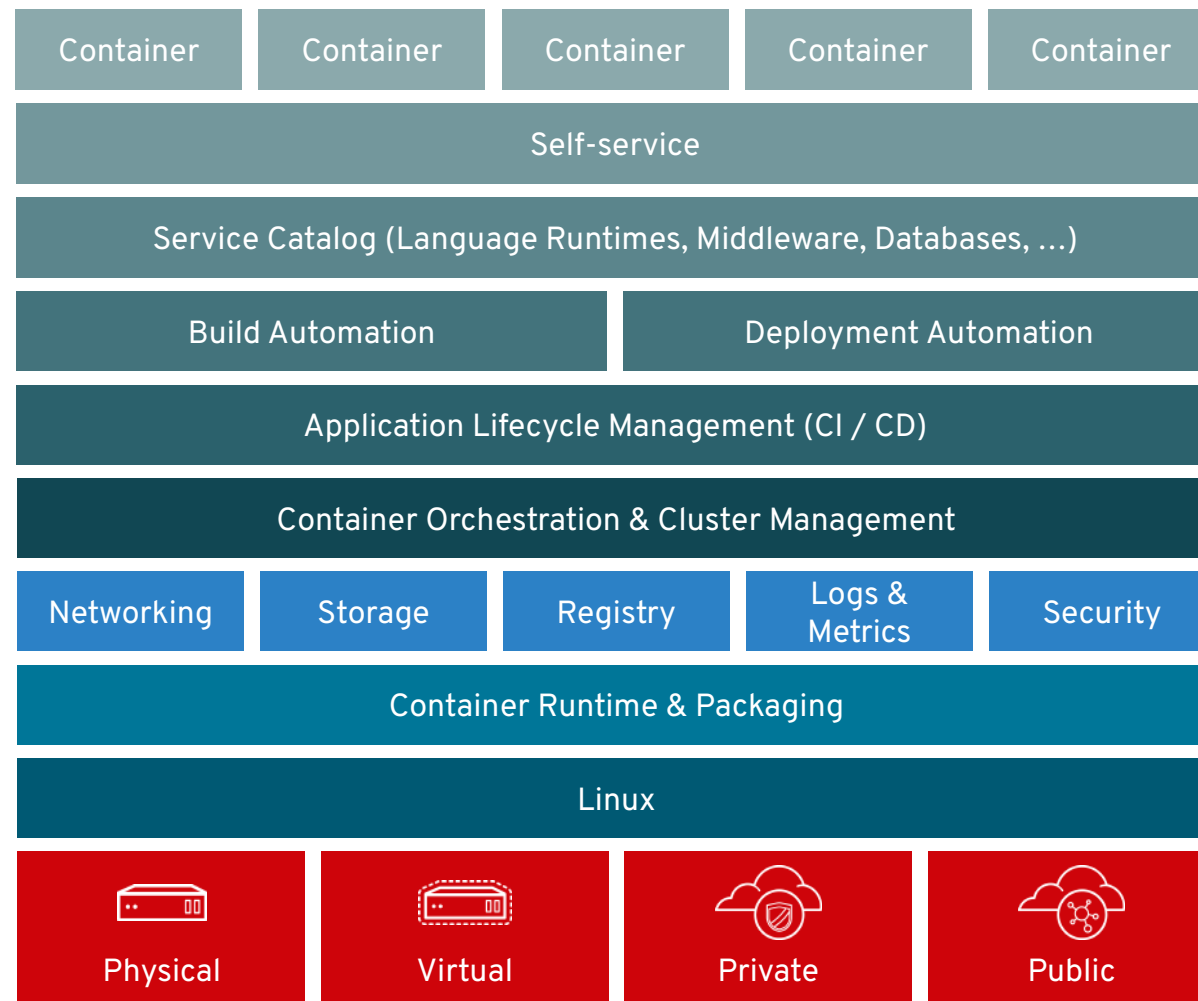
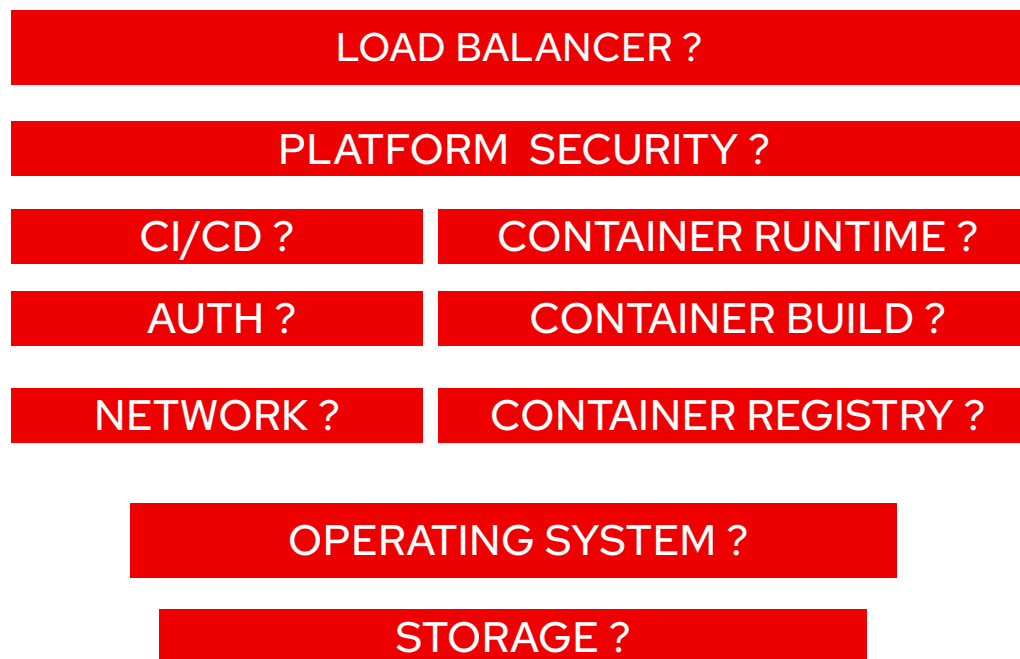
Kubernetes

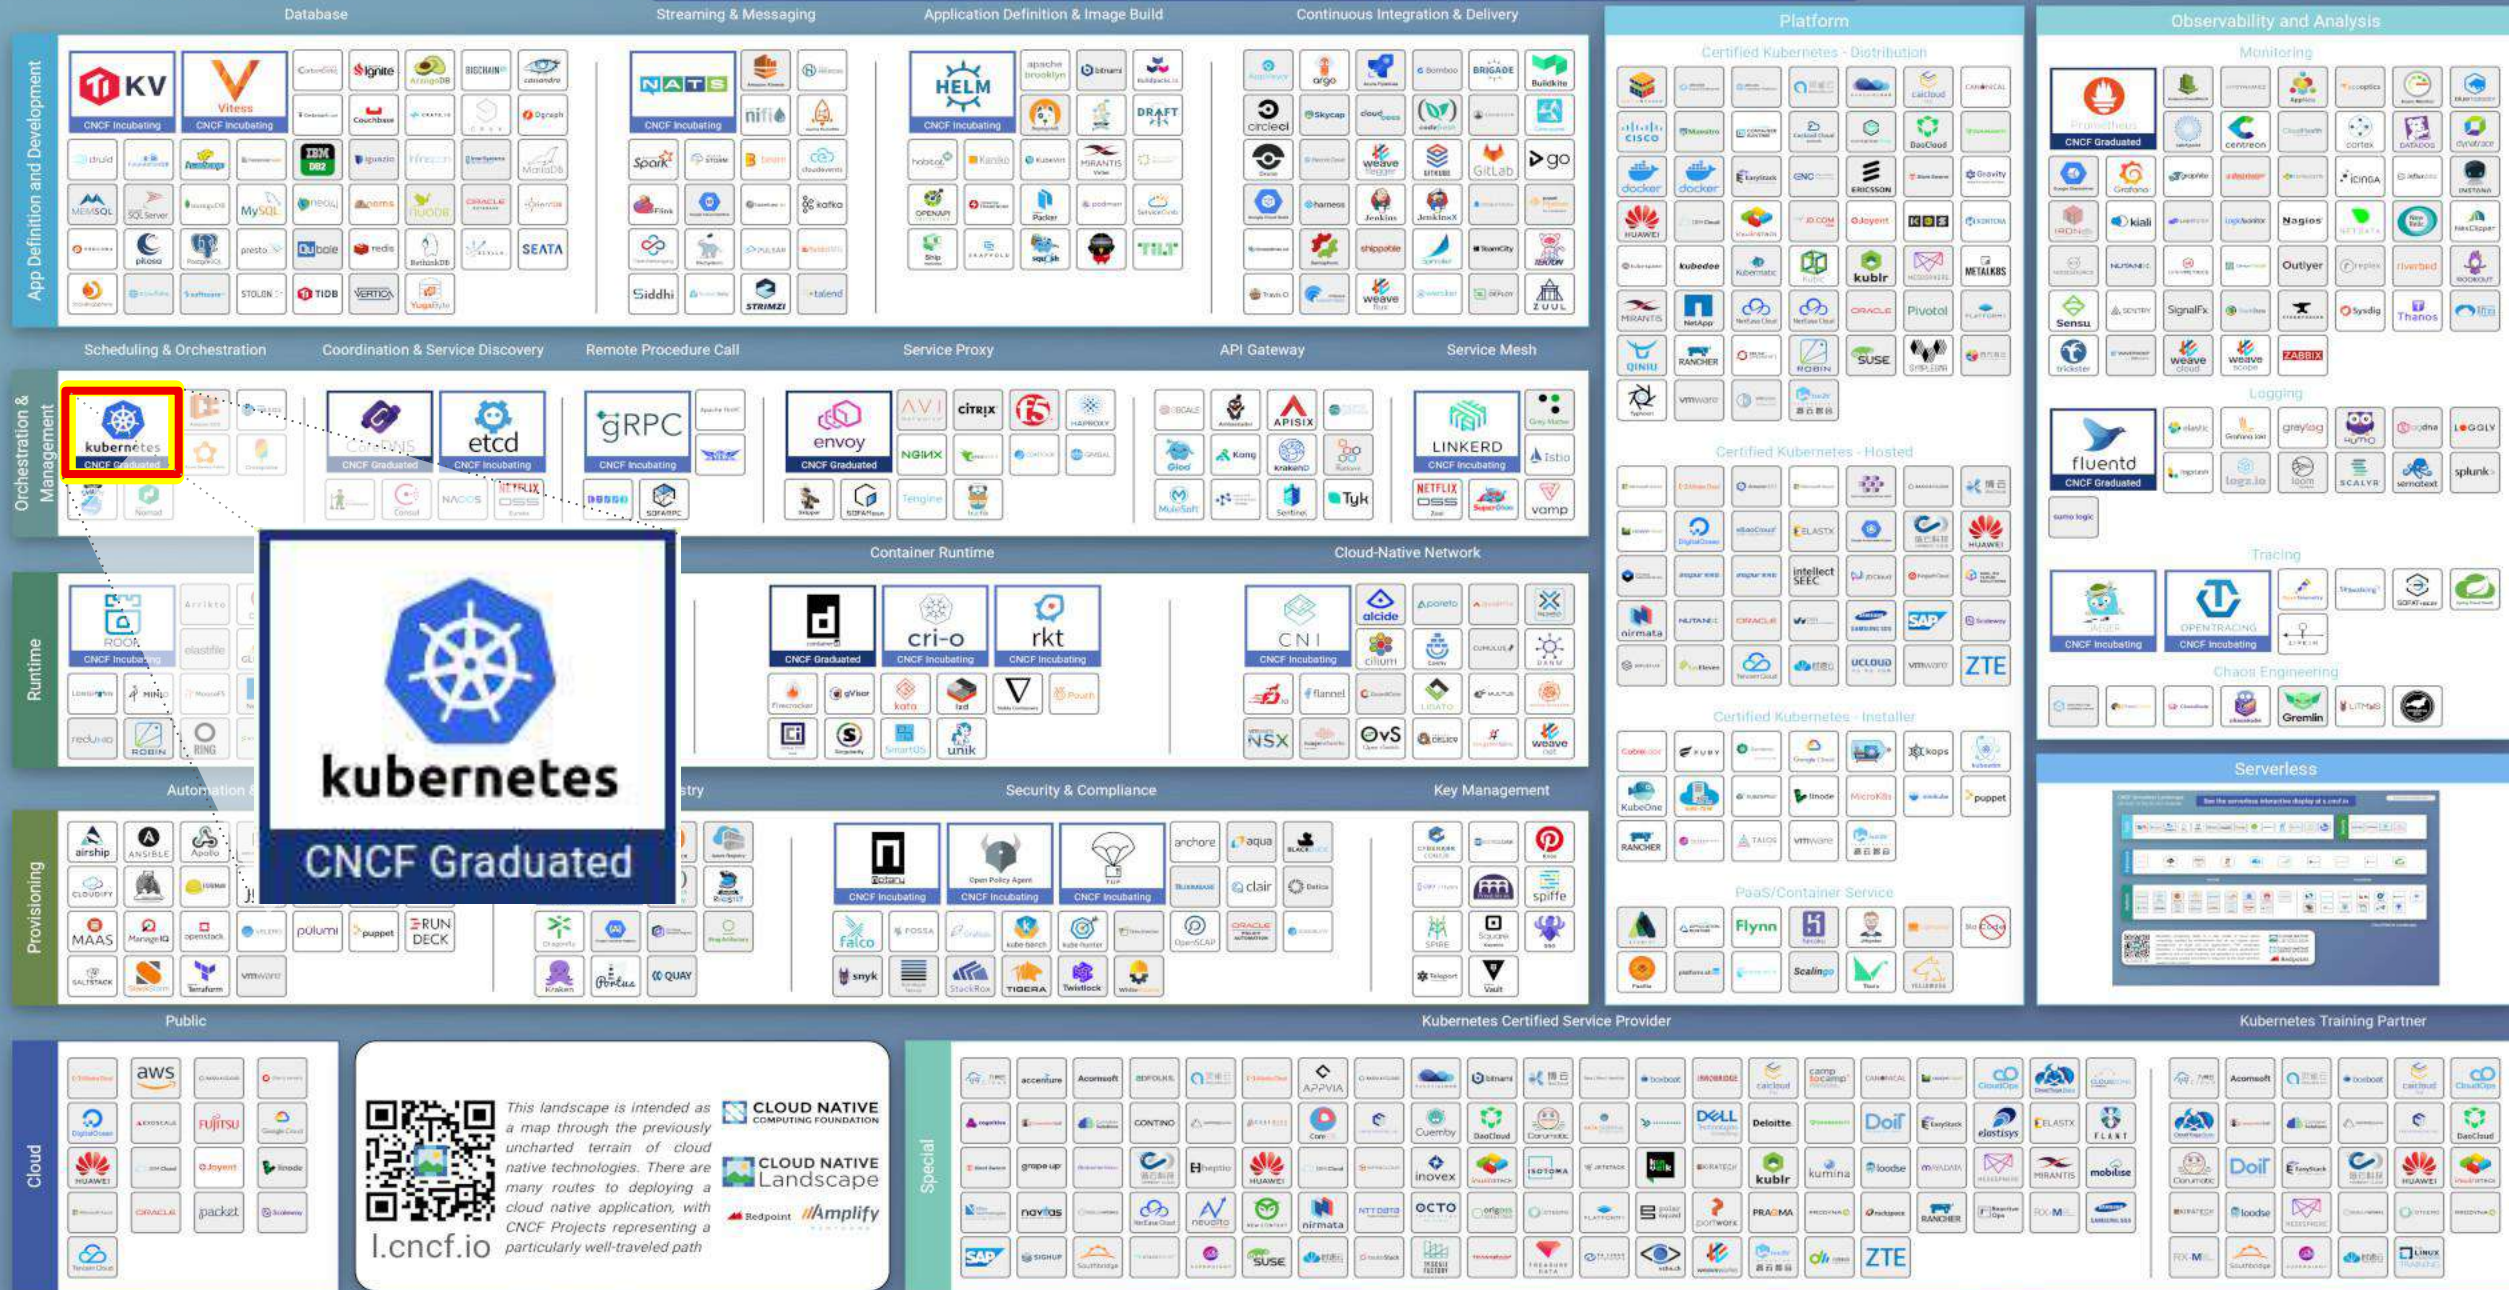
Automation

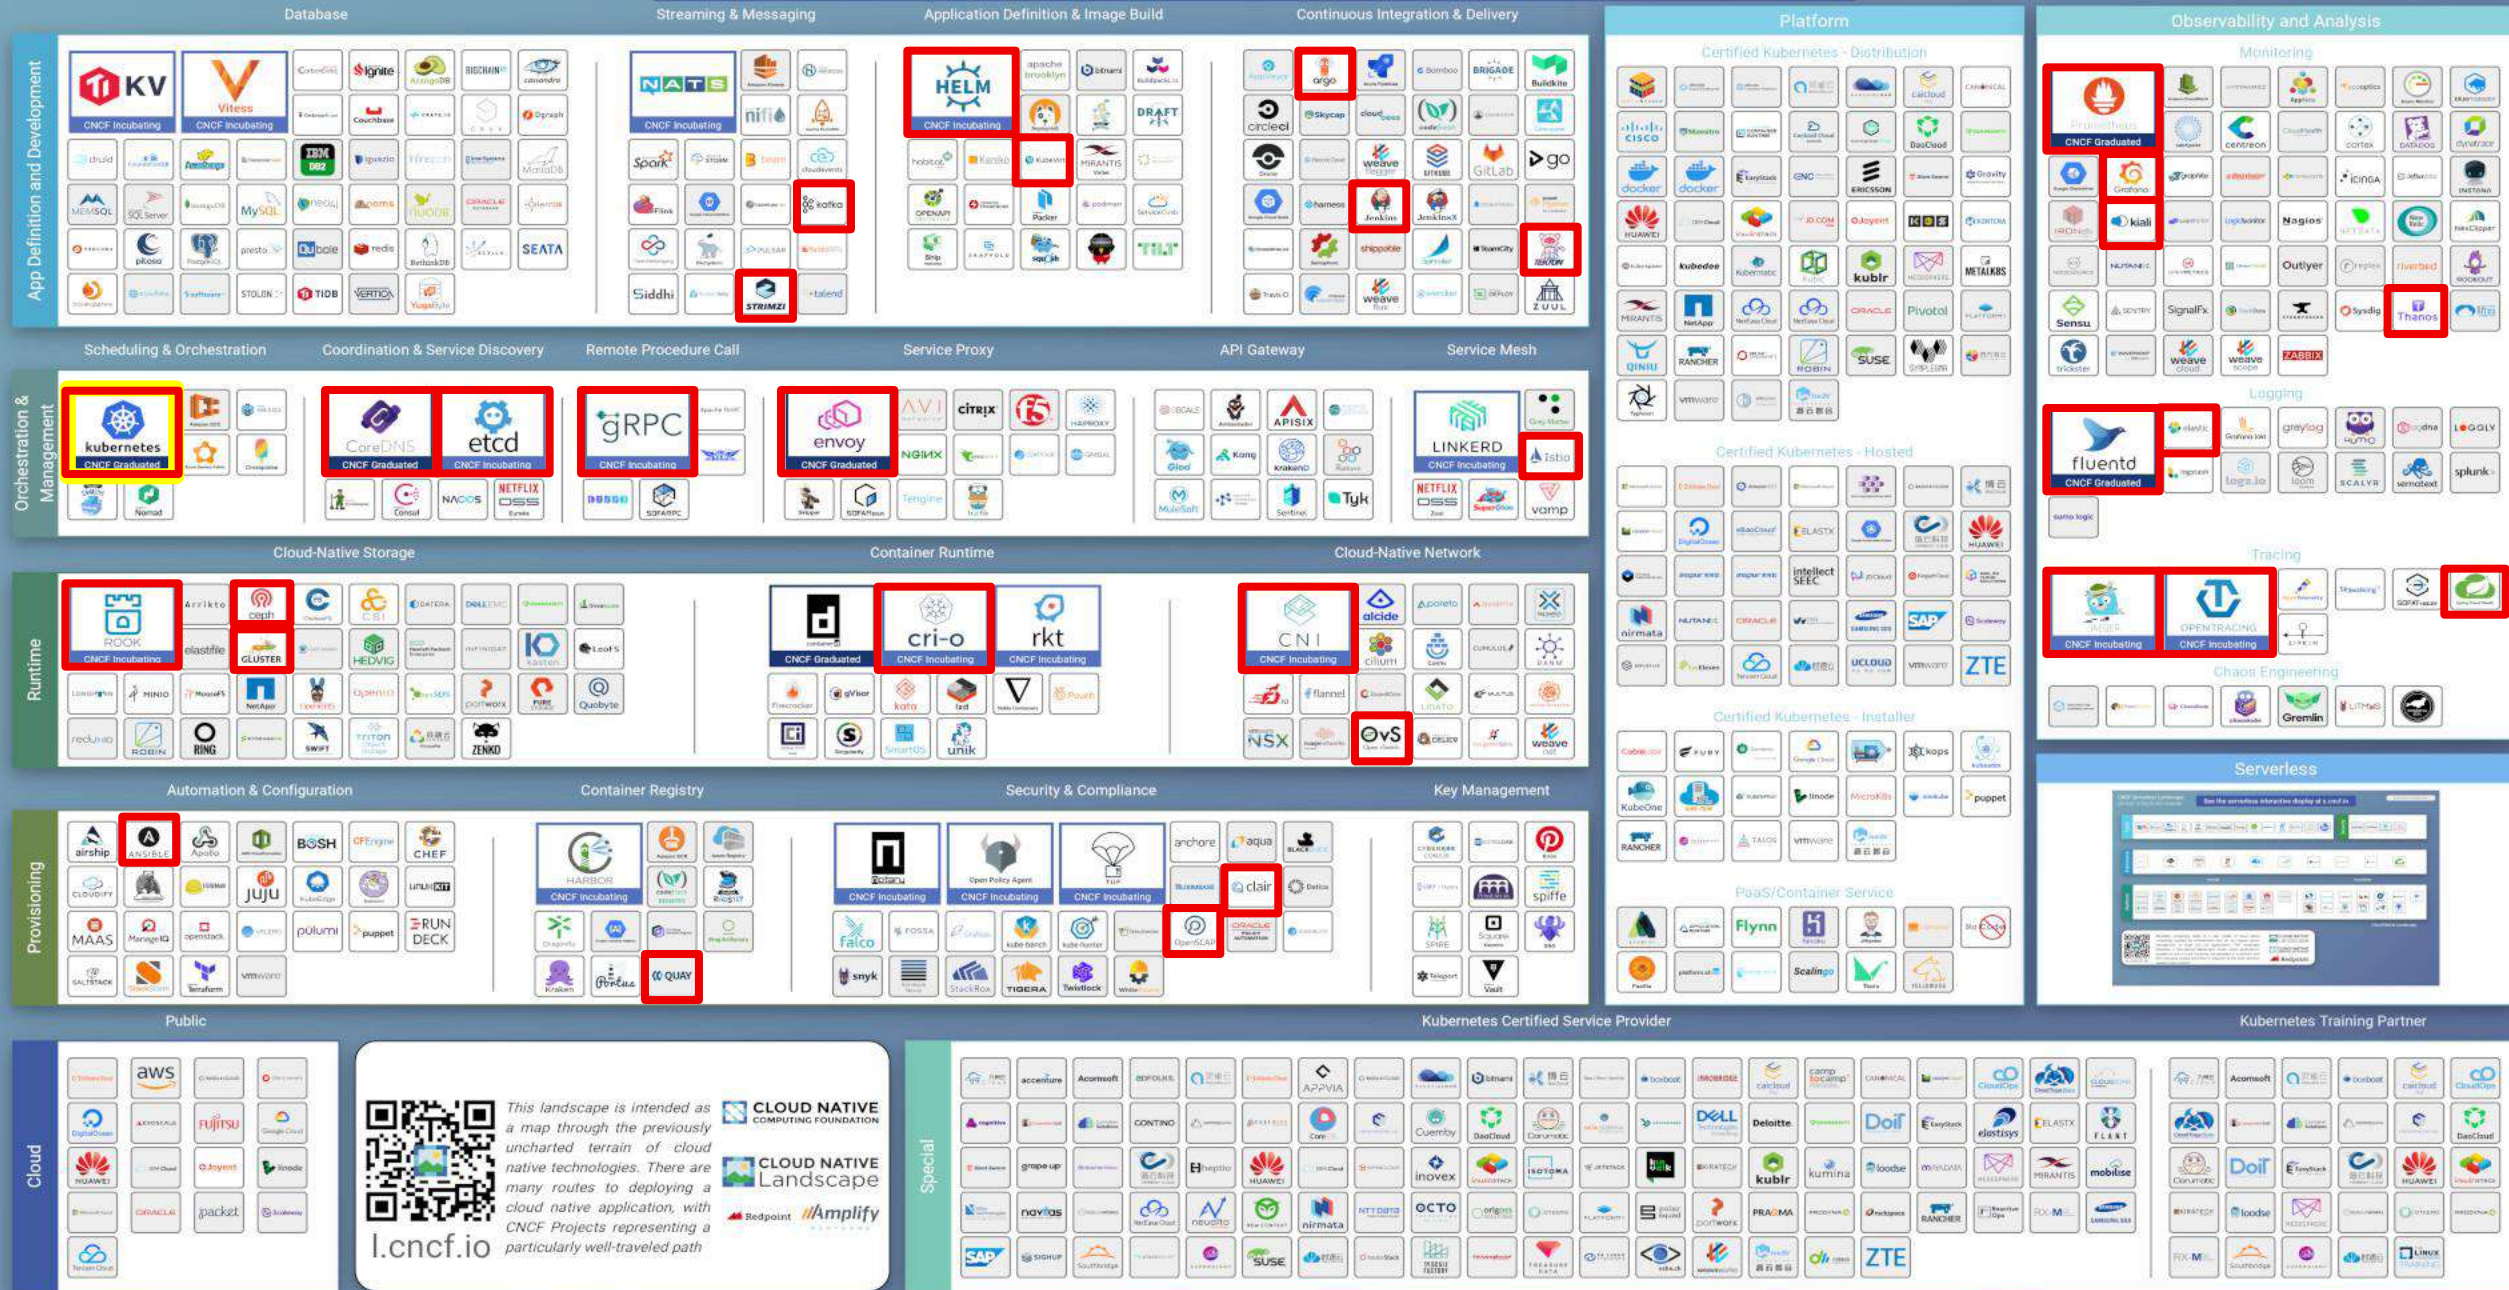
Linux

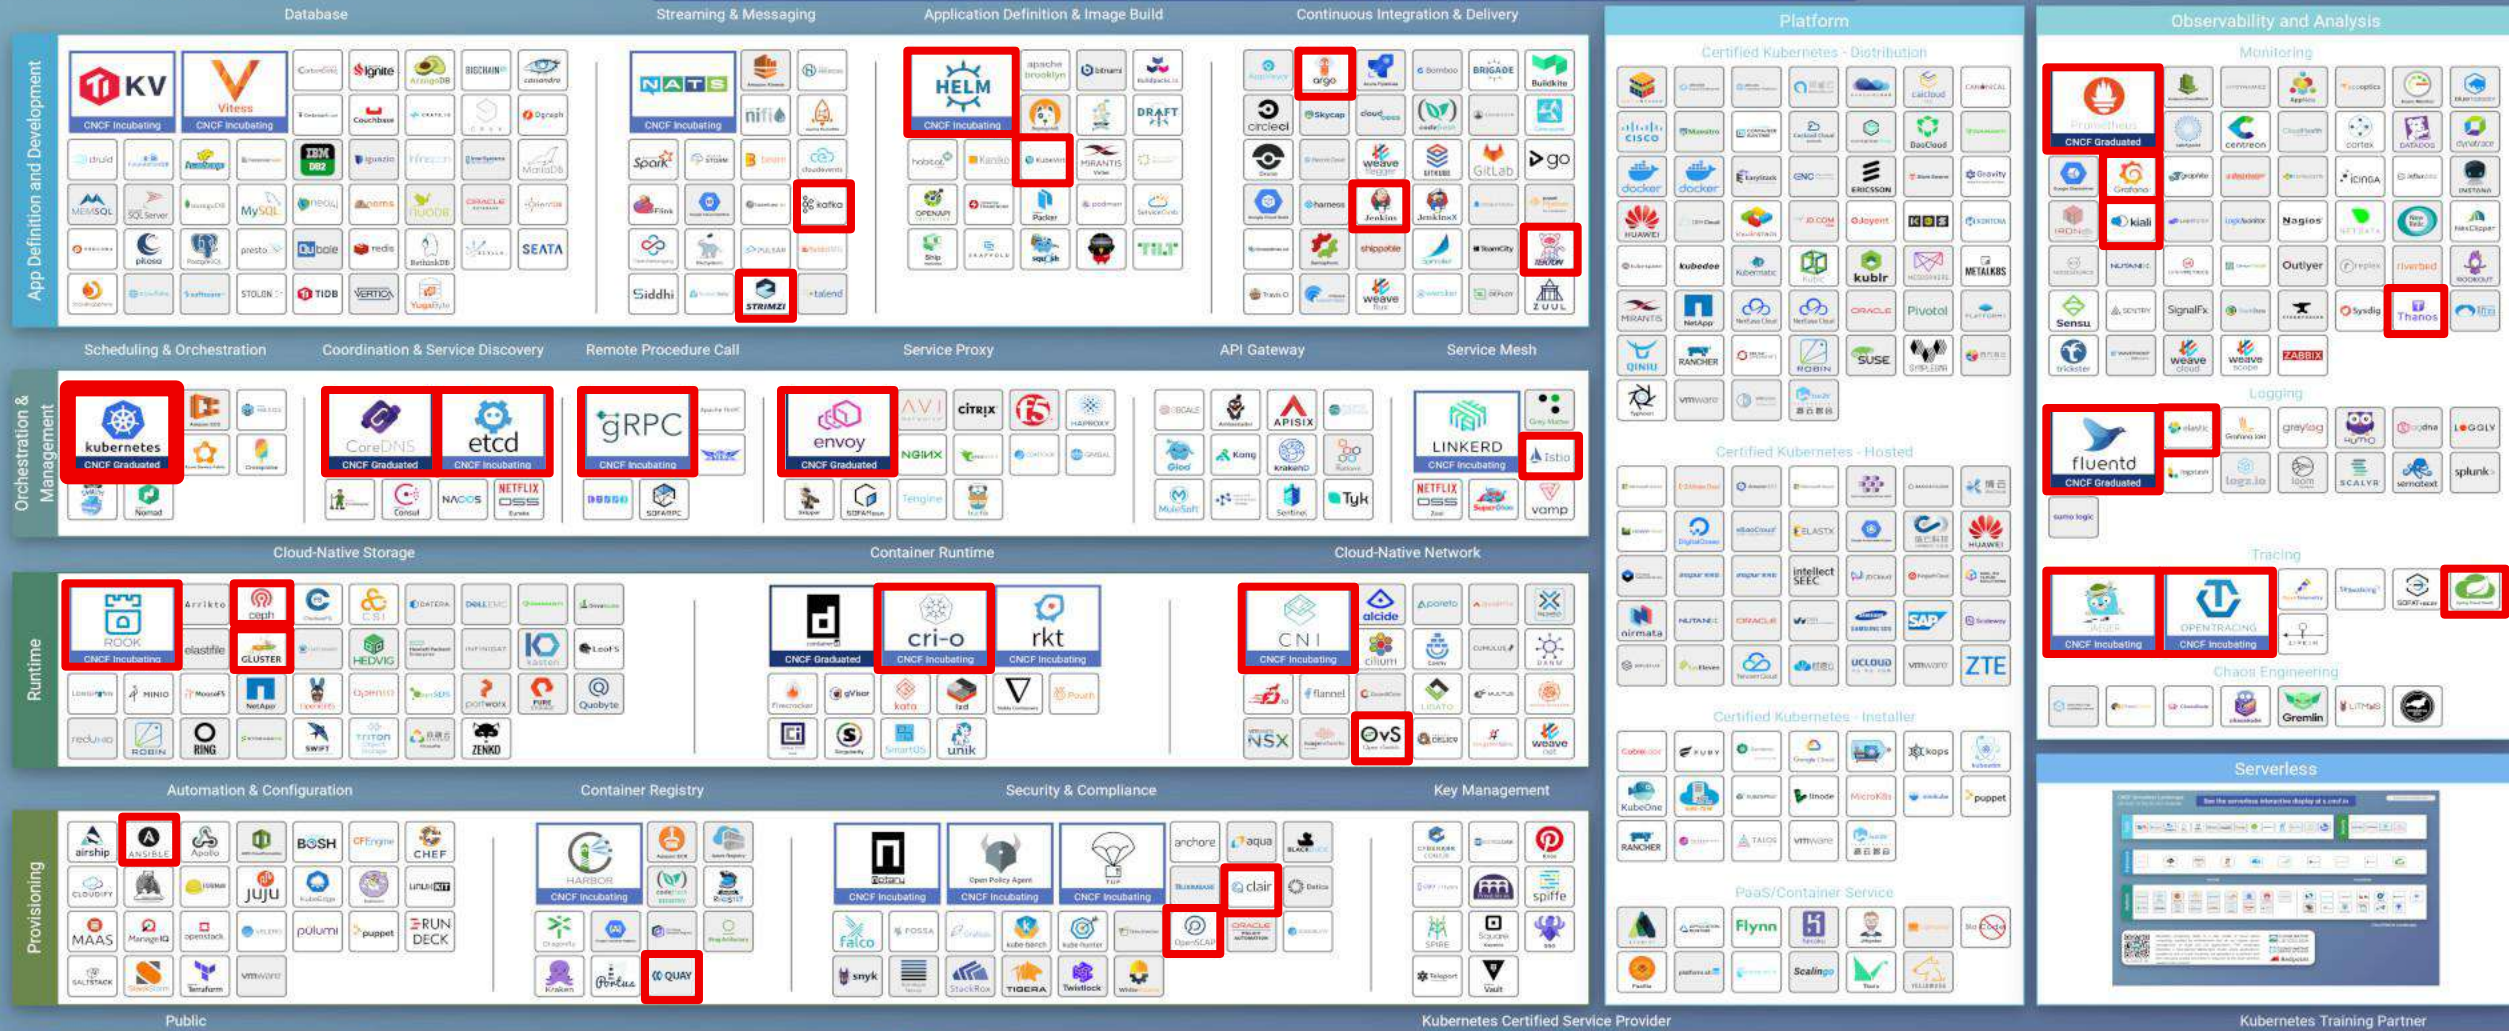
x%

The Challenge of Containers: **they require many components**









Red Hat manages and MAINTAINS this complexity for its customers

A consistent platform no matter how or where you run

Red Hat OpenShift cloud services—Fully managed, start quickly



Red Hat OpenShift Service
on AWS



Azure Red Hat
OpenShift



Red Hat OpenShift
on IBM Cloud



Google Cloud

Red Hat OpenShift
Dedicated

Self-Managed Red Hat OpenShift—Customer managed, for control and flexibility



On **public cloud**, on-premises on **physical** or
virtual infrastructure, and at the **edge**

Azure Red Hat OpenShift is a joint, 1st party Azure solution.



Integrated Dev tools and
cloud native services



Joint
Support & Engineering



Security, resilience
and compliance



Single invoicing, utilize
cloud committed spend


AI on Azure Red Hat OpenShift

What's in it for me?



Accelerate deployments

Build, train, deploy and monitors AI/ML workloads faster with a comprehensive app platform that operationalizes AI faster.



Focus

Focus on AI innovation and minimize risk with a fully managed platform



Consistency and choice

Flexibility to choose the right AI/ML technologies for your business, using the cloud of your choice, with a consistent experience.

Operationalize AI with Red Hat OpenShift AI and Azure Services

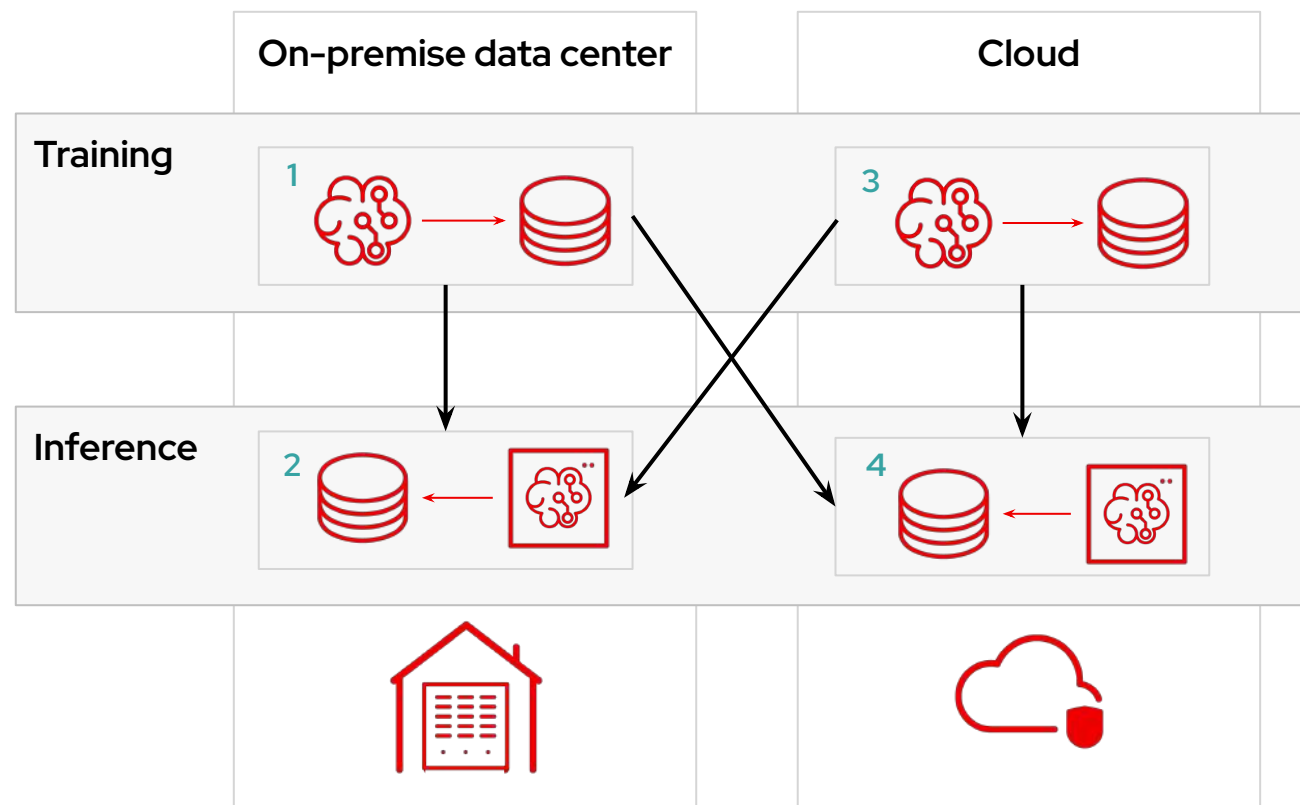


Gain hybrid cloud flexibility

Train and deploy models and AI-enabled apps on-premises, cloud or edge


What you do should not dictate
where you do it

1. Data on-prem = Train on-prem
2. Data on-prem = Inference on-prem
3. Data in the cloud = Train on cloud
4. Data in the cloud = Inference on cloud



Why Red Hat OpenShift AI on Cloud


Economic Entry Point with Long-Term Cloud Viability



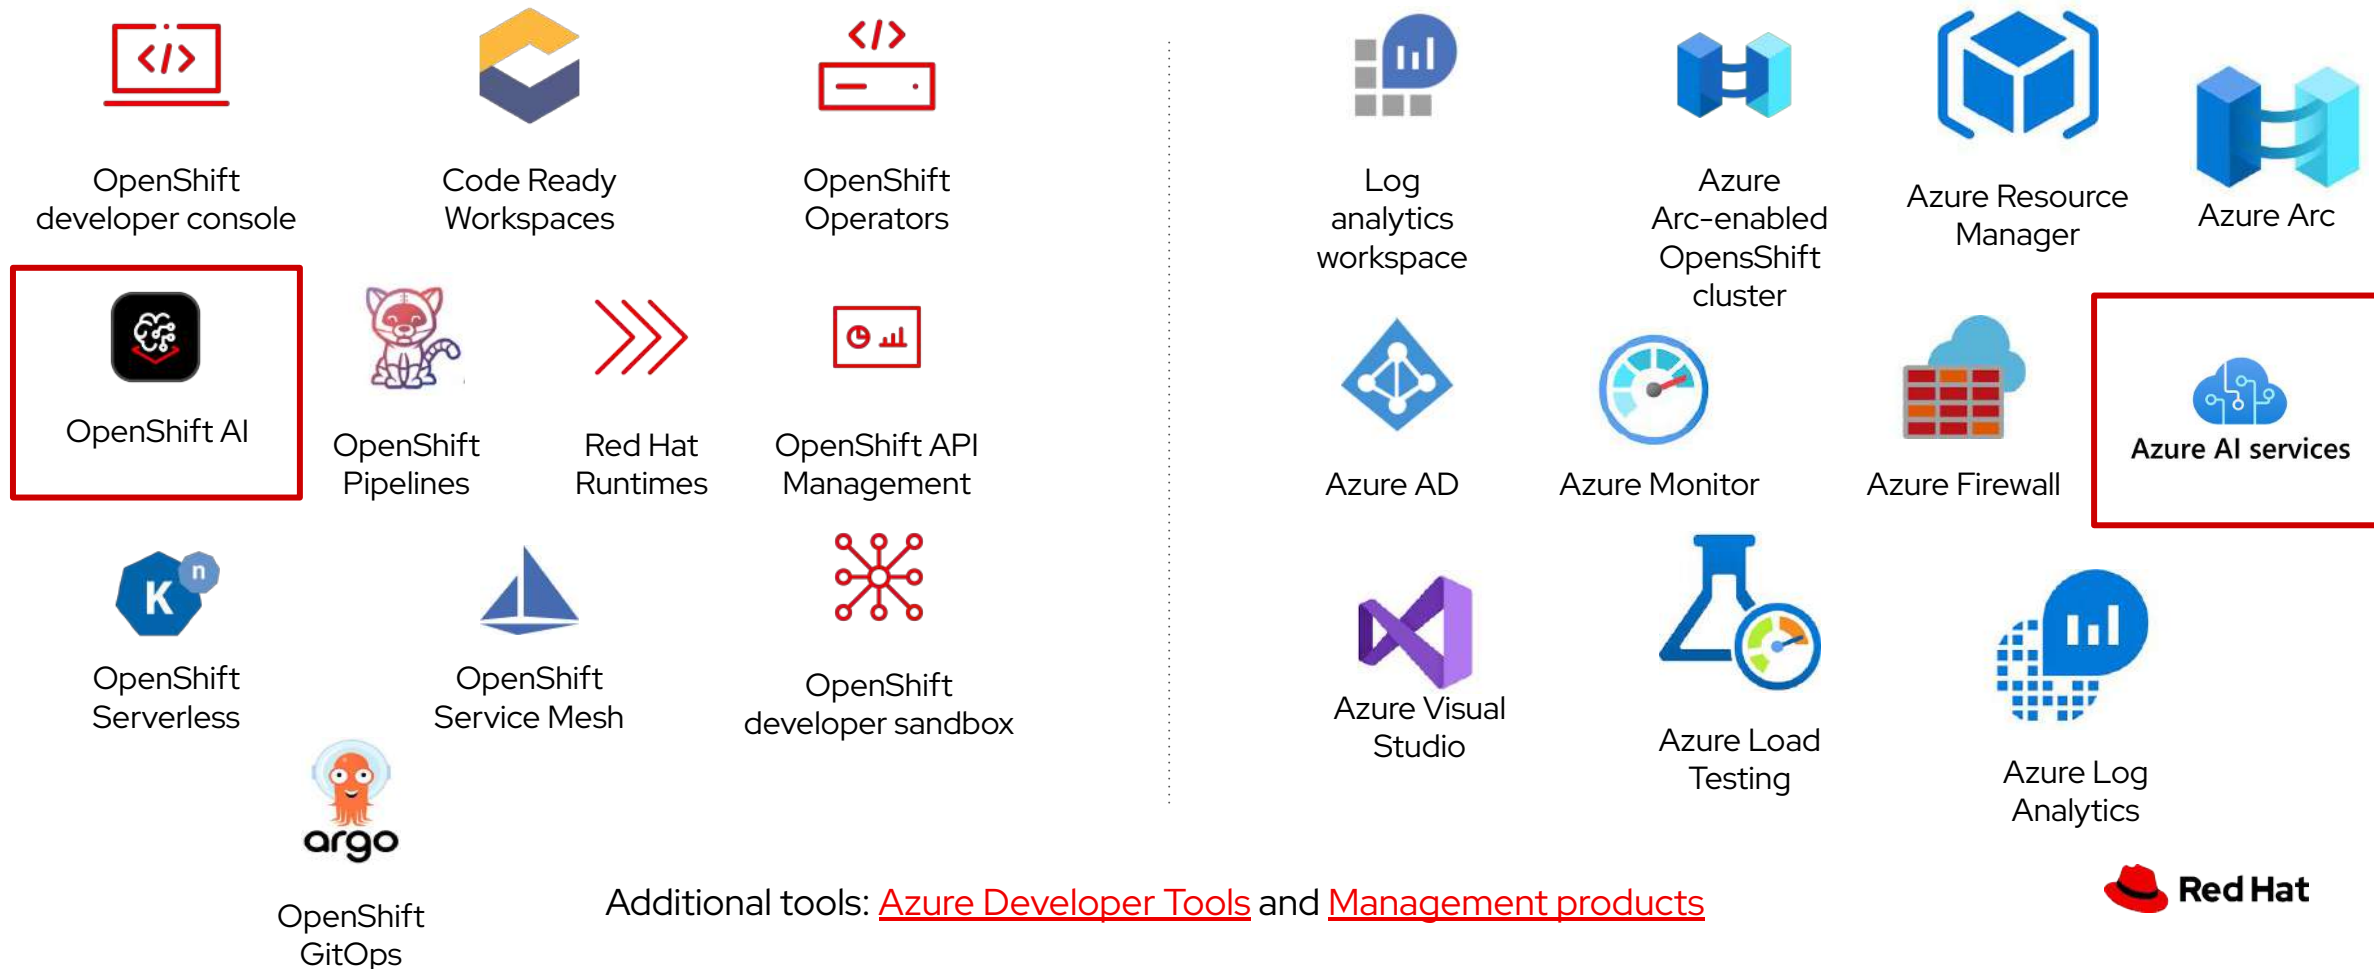
Red Hat OpenShift AI on Cloud

- ▶ Scalability and availability
- ▶ Cost effectiveness
- ▶ Integrated Cloud Environment for AI/ML Collaboration

Build an AI platform for E2E AI lifecycle management



Azure Red Hat OpenShift integrates with OpenShift and Azure Developer and Management Tools



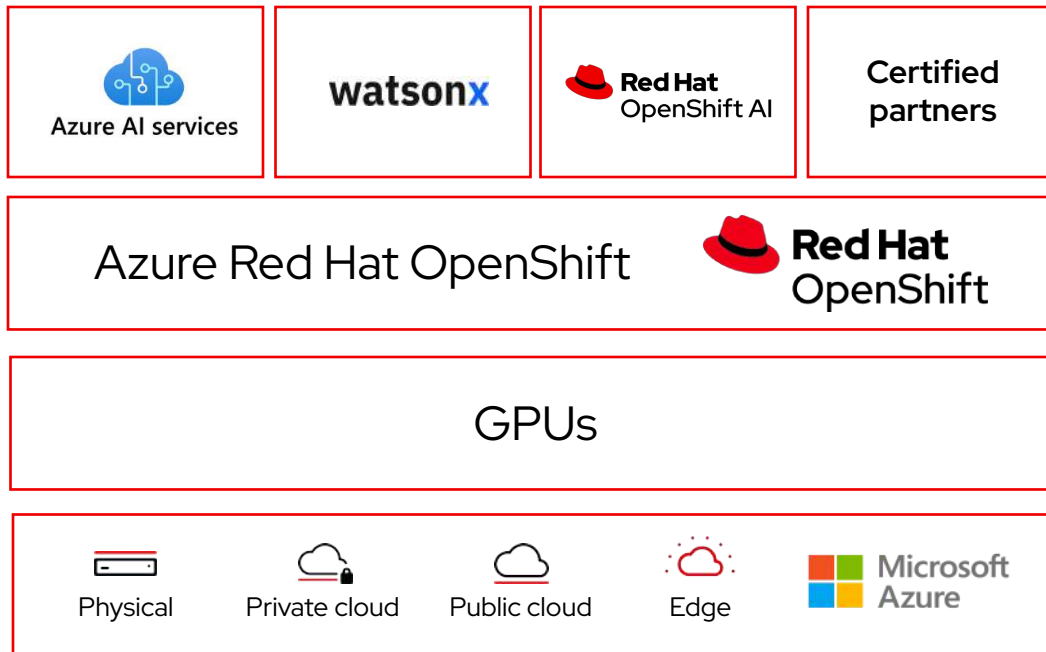
More Choice for Model Development, Tuning and Serving on Azure Red Hat OpenShift

Model
development,
serving and
monitoring

Orchestration,
compute resource
and fleet mgt

GPU acceleration
support

Deploy anywhere



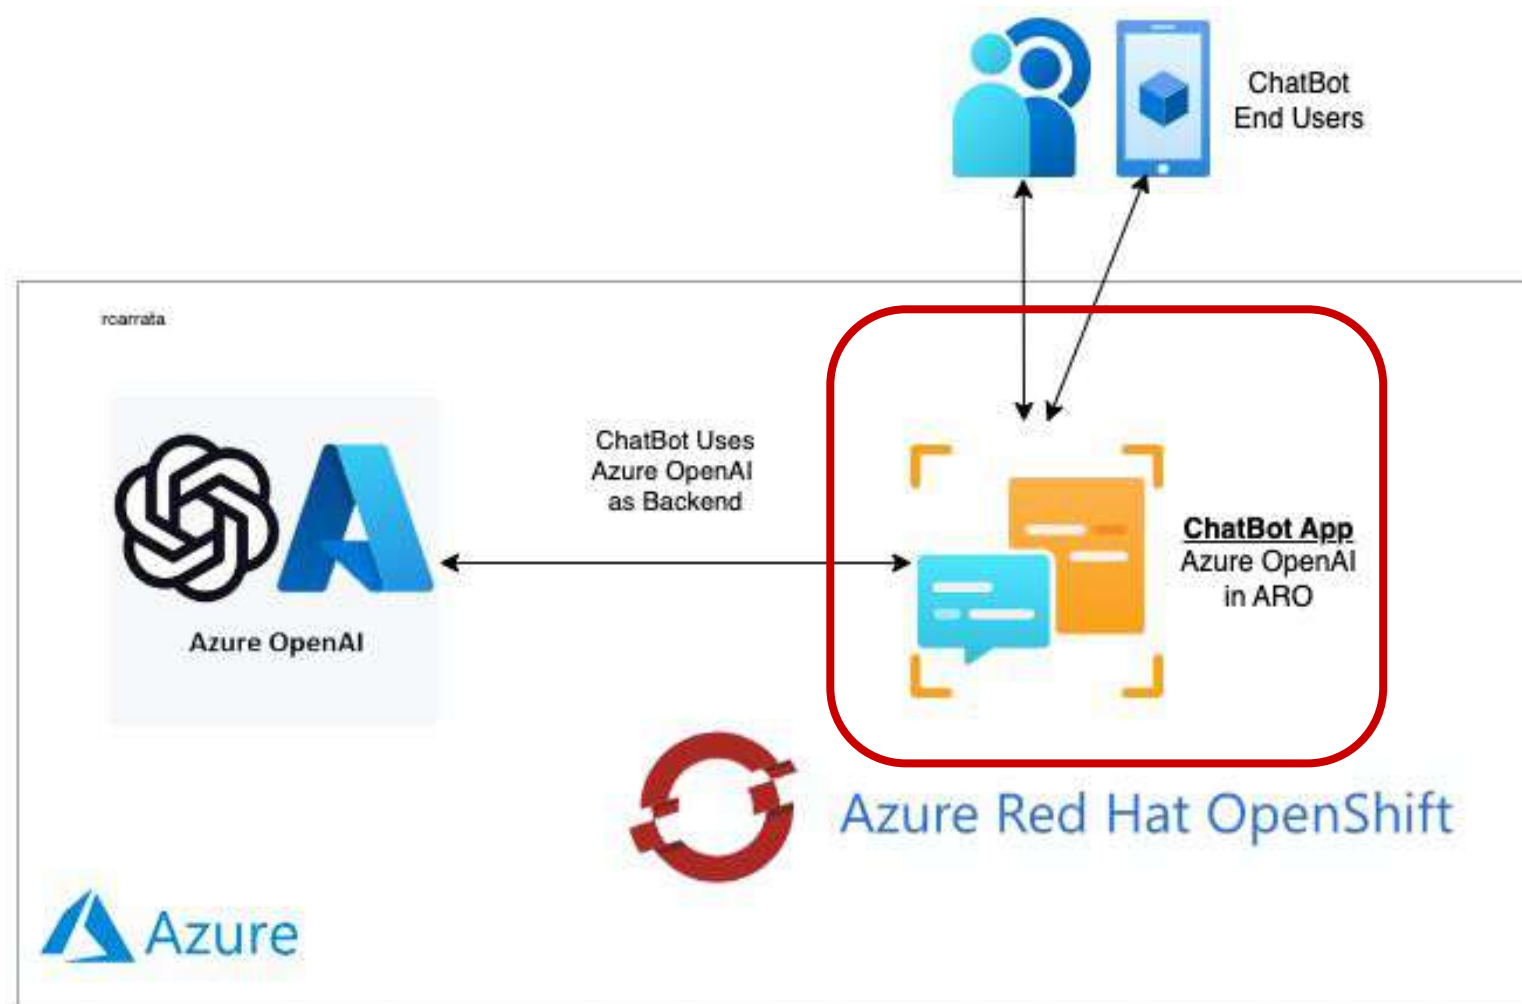
Consistency and choice
across deployment options
and AI/ML tools

Focus on AI Innovation and
minimize complexity with a fully
managed platform

Accelerate deployments
with a comprehensive app
platform that operationalizes AI
faster

Example: AI ChatBot

in Azure Red Hat OpenShift fully integrated with Azure OpenAI



How the **AI/ML landscape** is evolving / MLOps and OpenShift AI Overview



Operationalizing AI is a major challenge



Understanding the nuances

What type of AI model should you invest in?



Building the right infrastructure

How do you choose and build the right environments for development and production?




Realizing value

How do you safely realize the business benefits of AI/ML innovation?

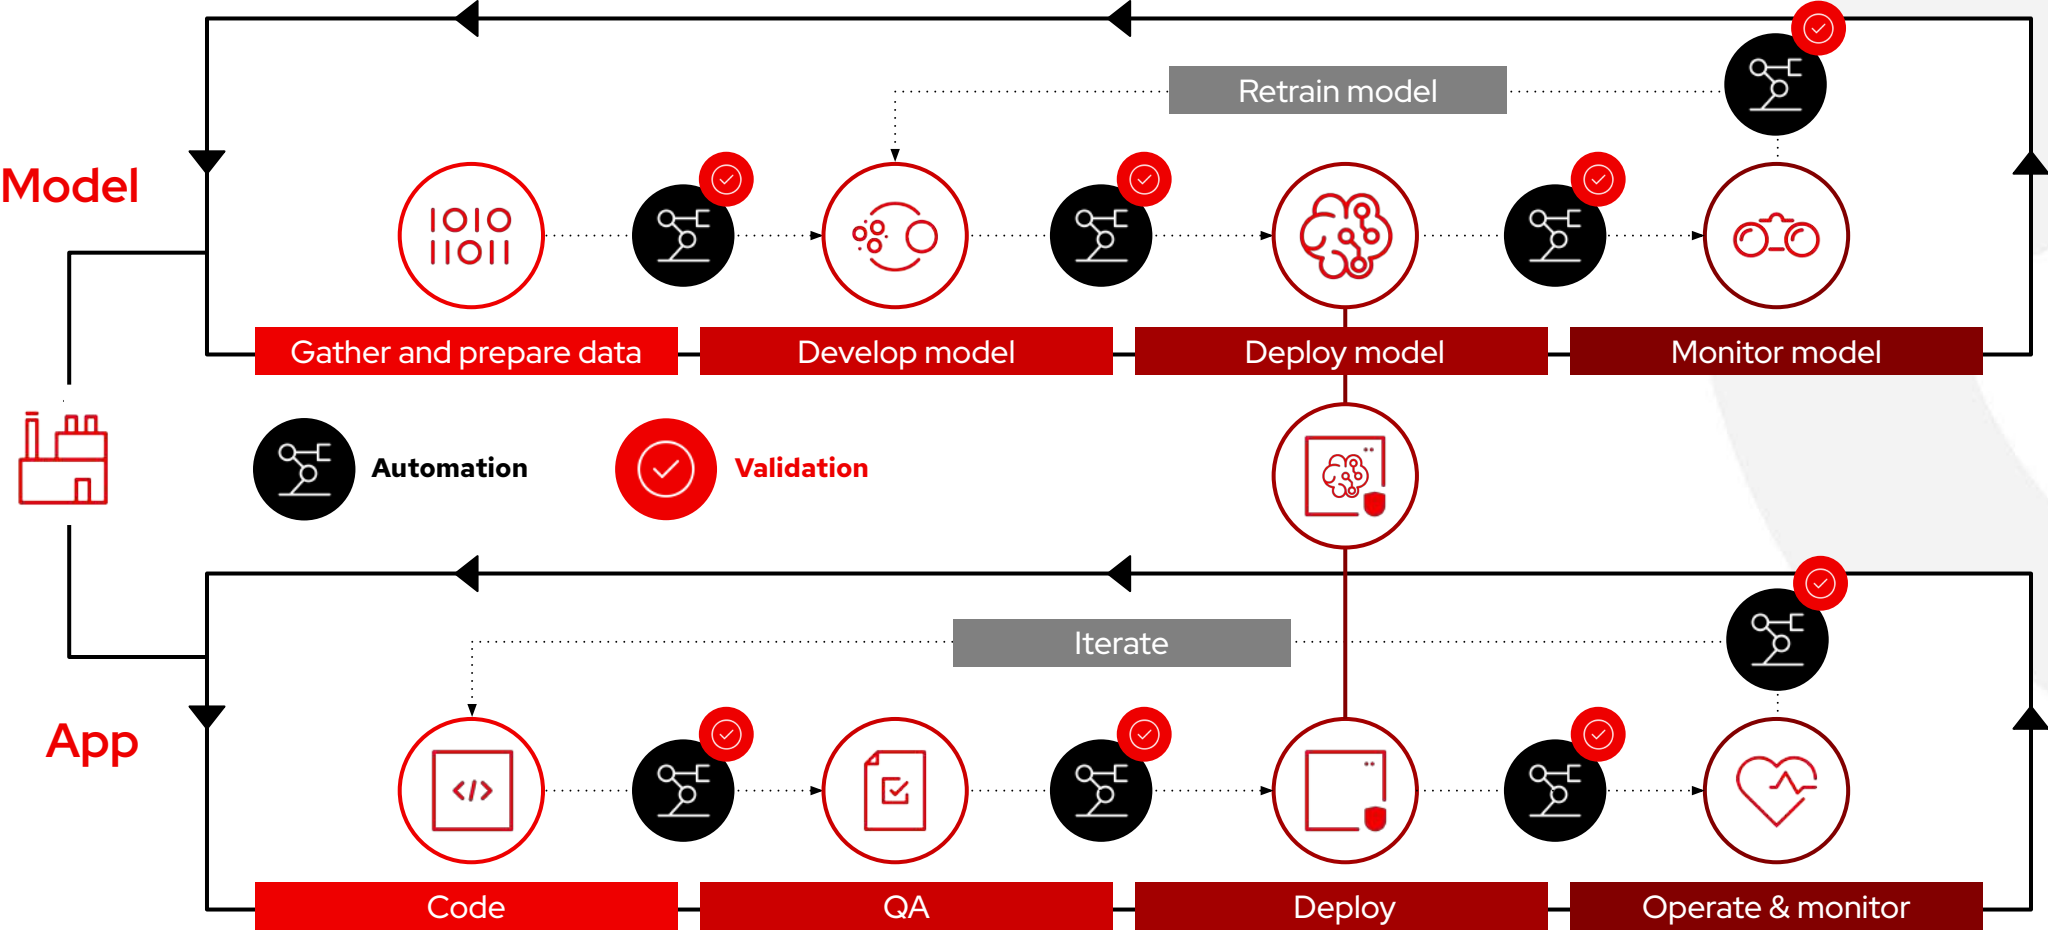


87% of data science projects never make it into production¹




76% of respondents say it takes 7 mo to 1+ year to operationalize a model²

Lifecycle for operationalizing models






Develop, train, serve, monitor,
and **manage** the life cycle of
AI/ML models and applications,
from **experiments** to **production**.



Built on top of OpenShift®

Deliver consistency, **cloud-to-edge** production deployment and monitoring capabilities




Designed for Gen and predictive AI

Scale to meet the workload demands of foundation models and traditional ML.



Empowered collaboration

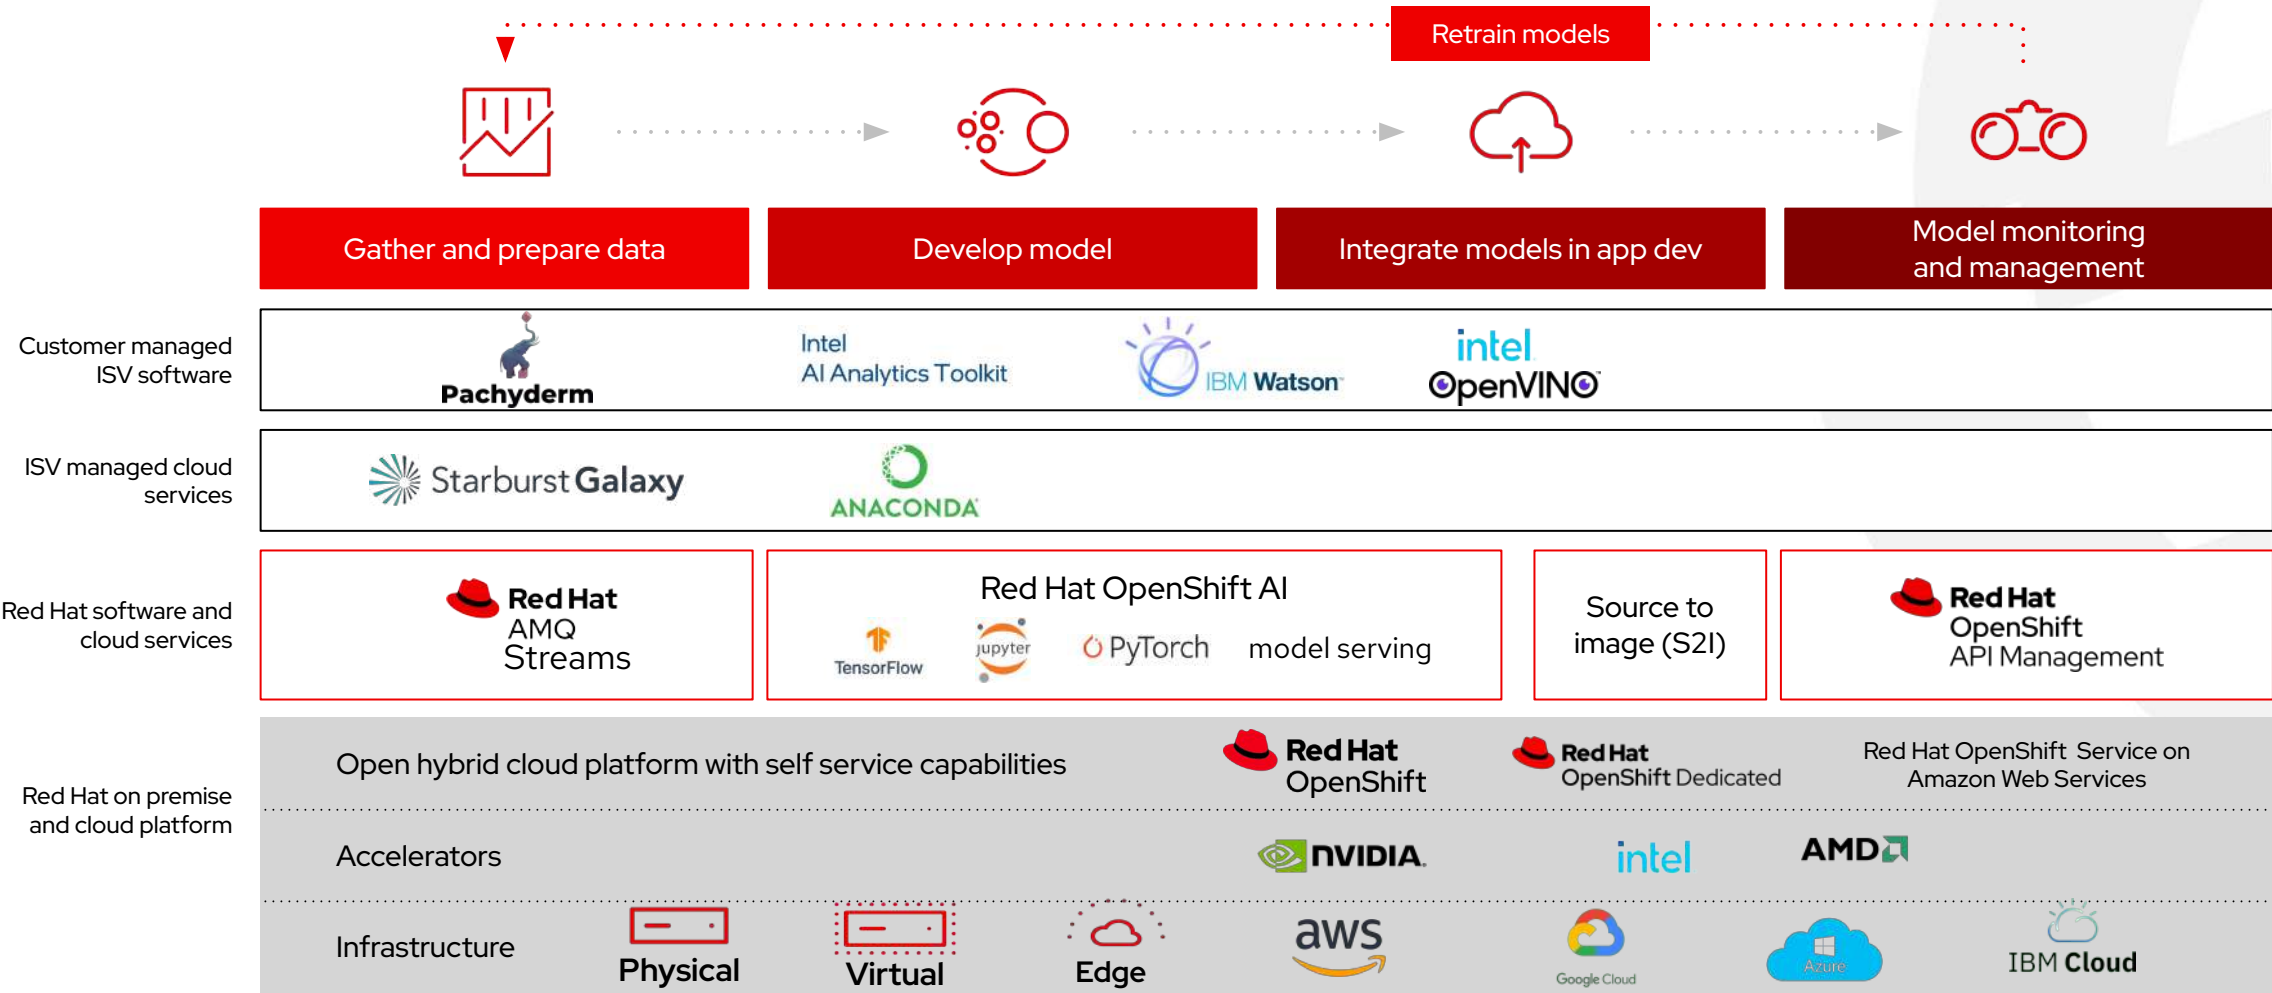
Provide a **unified platform** for data scientists, operations and intelligent application developers



DevOps applied to ML

Set up rigorous **pipelines** and **workflows** to take you from **development** to **production**.

Red Hat OpenShift AI and MLOps



Hands-on Workshop

Digital Transformation at Parasol Insurance



Hands-on lab: Transparency Guide

[Document](#)

- ▶ You are accessing Azure Red Hat OpenShift, Red Hat OpenShift AI and Azure OpenAI via the Red Hat Demo Platform that is operated by Red Hat. **You are not using your own environment or applications during this demo.**
- ▶ Access to this demo is restricted by invitation only.
- ▶ Only sample text will be used in the demo; you will not bring your own data. **Do not enter personal or confidential information while using the demo.**
- ▶ Usage metrics from the demo environment are tracked for reporting purposes. Red Hat does not retain any customer entered data.
- ▶ Access to the lab will be turned off after the workshop is complete and the demo will be deleted.
- ▶ You can provide feedback about the demo environment functionality to your Red Hat Account Team who invited you to this session.
- ▶ You can provide feedback about the content of the demo to brjackso@redhat.com.

If you have any questions or concerns about the transparency guidelines listed above, please do not move forward with the workshop.

Hands-on lab: Use AI/ML to help insurance adjusters be more productive

- ▶ Work with data projects and workbenches.
- ▶ Use a large language model (LLM), compare different model sizes, and test accuracy.
- ▶ Use the Retrieval Augmented Generation, or RAG, technique to get more accurate responses.
- ▶ Fine tune an image recognition model and test accuracy.
- ▶ Serve & deploy the models.
- ▶ Integrate a frontend application with the models.
- ▶ Test the application as a whole.

Hands-on lab: Access the Workshop

► Link:


- <https://red.ht/aro-ai-bucharest>


No place? Alternative:

<https://red.ht/aro-ai-bucharest-i-am-lucky>

► Password:

- **aro**






AI Workshop on ARO – Bucharest 2025 April



Access to AI Workshop on ARO – Bucharest 2025 April

Email * ⓘ


Workshop Password * ⓘ

[Access this workshop →](#)

Copyright © 2024 Red Hat, Inc. | [Privacy statement](#) | [Terms of use](#) | [All policies and guidelines](#) | [Cookie preferences](#) | [Babylon](#)



OpenShift AI Unleashed: Transforming Claims Processing for Maximum Efficiency!

You are user3 [Links](#) 

[1. Background](#)

[2. Connection and Setup \(user3\)](#)

[3. Working with an LLM](#)


[4. Image Processing](#)

[5. Web App Deployment](#)

[6. Productization and Extrapolations](#)

[7. End of Lab](#)

[Dev Mode](#)

 OpenShift AI on Azure Unleashed: Transforming Claims Processing for Maximum Efficiency! [Edit this Page](#)

OpenShift AI on Azure Unleashed: Transforming Claims Processing for Maximum Efficiency!

Parasol Insurance lab Overview

This lab will illustrate how the use of various AI/ML technologies can be combined to produce a valuable solution to a business problem.

The information, code, models and techniques it contains are illustrations of what a first prototype could look like. It is not the definitive way of addressing the stated requirements.

Azure and Azure AI Services

In this lab we are using an **ARO (Azure Red Hat OpenShift)** cluster, a native Azure service, with **Red Hat OpenShift AI** installed. We are going to use **Azure OpenAI GPT-4**, GPT-4 is a large language model that we will use to summarize and extract information from fictive insurance claims. We are then also going to use **Azure AI Search**, it is a vector database that can be used to store documents for retrieval when implementing RAG (Retrieval Augmented Generation). For our test application we are going to use **Azure Databases for PostgreSQL** to store our fictive insurance claims data.

Disclaimer

This lab is an example of what a customer could build using Red Hat OpenShift AI. Red Hat OpenShift AI itself has no specific feature related to Insurance Claim Processing.

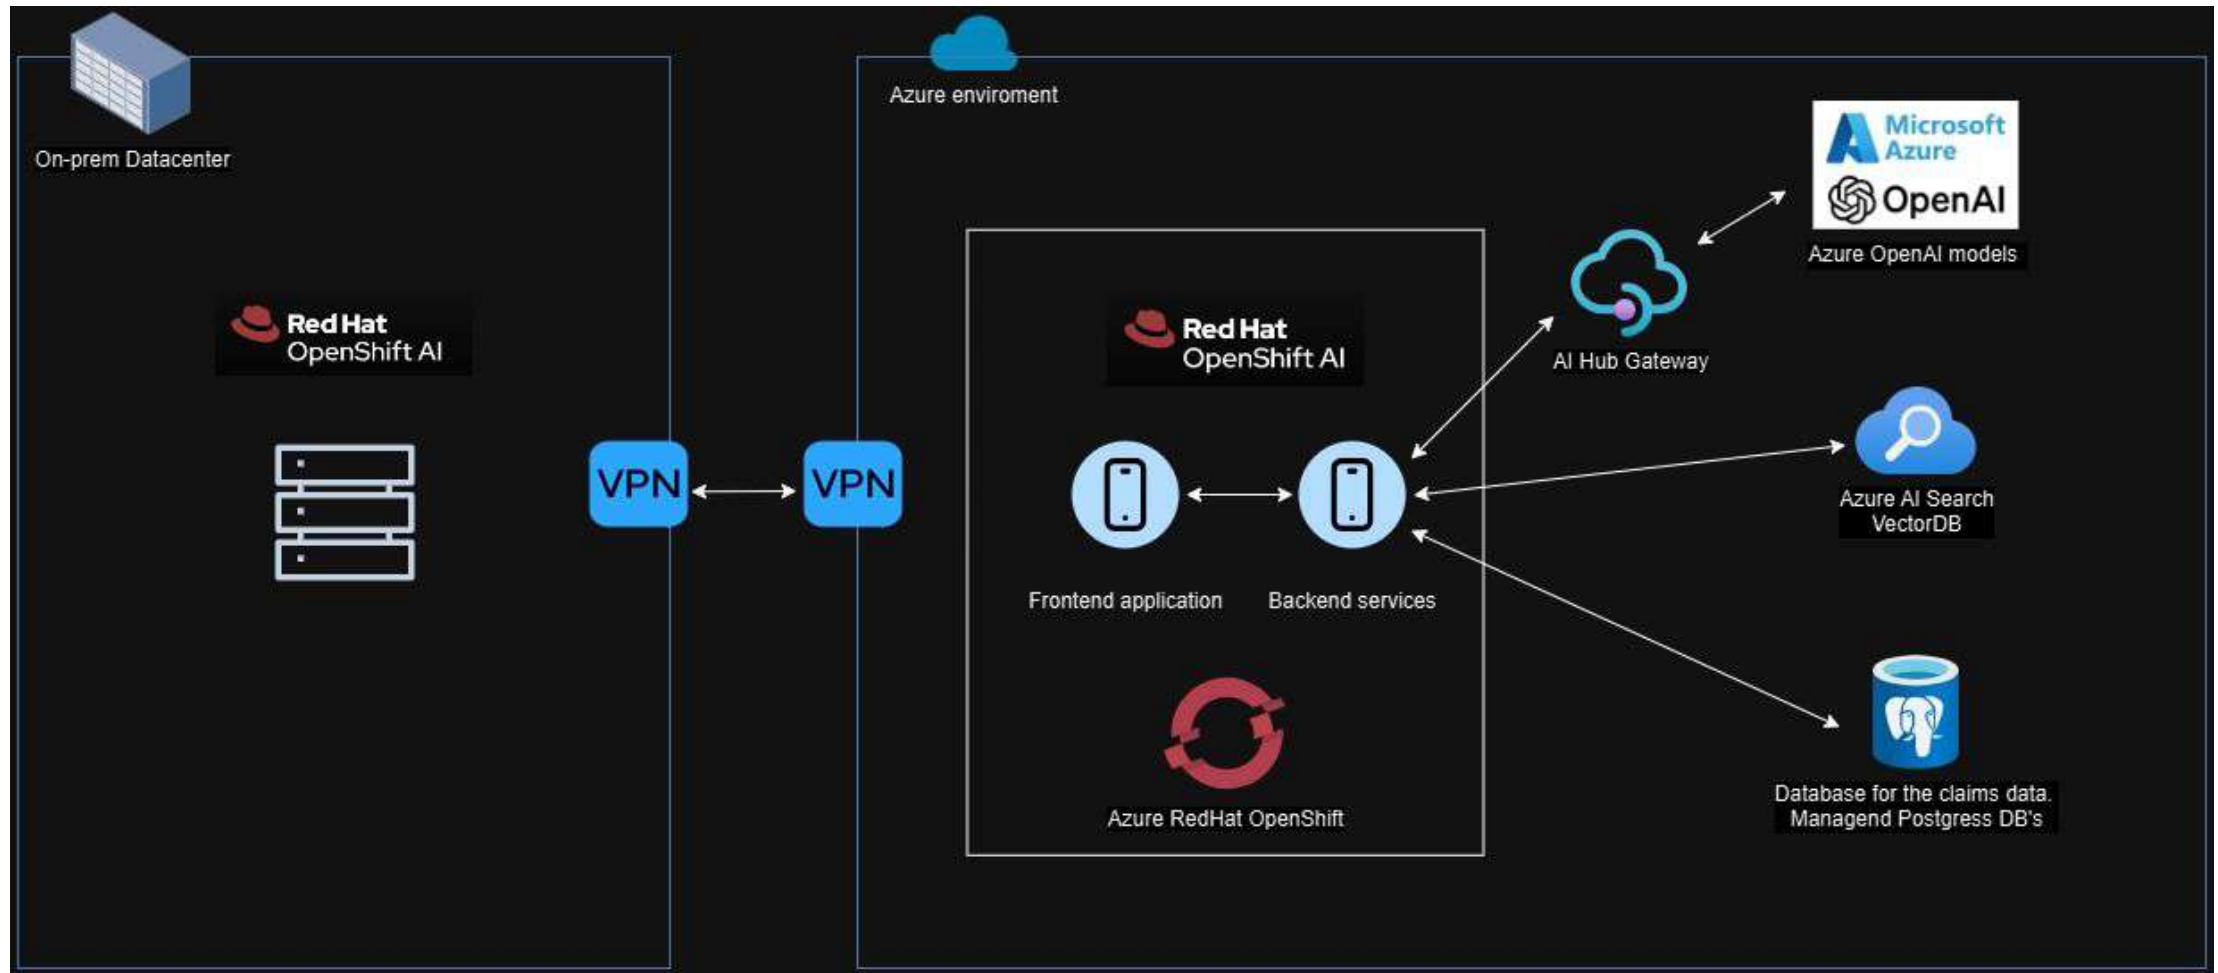
This lab makes use of large language models (LLM) and image processing models. These models are not included in the Red Hat OpenShift AI product. They are provided as a convenience for this lab.

The quality of these models is enough for a prototype. Choosing the right model to use in a production environment is a complex task that requires a lot of experimentation and tuning. This lab does not cover this aspect.

Contents

- [Parasol Insurance lab Overview](#)
- [Azure and Azure AI Services](#)
- [Disclaimer](#)
- [Timetable](#)

Azure Red Hat OpenShift AI Workshop Architecture




Hands-on Workshop

MCP Protocol



Model Context Protocol (MCP)

Say what now?



[Video Link](#)

Earn your digital credential today

Complete the quiz and survey



Red Hat Technical Workshop

AI on Azure Red Hat
Openshift

2025 Recipient



red.ht/AlonARO_DigitalBadge.

- Complete the quiz and survey and pass it over 50 %
- Will be eligible for a digital badge.
- Successful recipients will receive an email from Credly.com within 2-3 weeks
- Promote your credential on LinkedIn and Tag one of your speakers (Red Hatters) on LinkedIn



Wrap up and Q&A





Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



[facebook.com/redhatinc](https://www.facebook.com/redhatinc)



twitter.com/RedHat

Appendix

Operationalize AI with Red Hat Services



Red Hat Consulting Services for your AI/ML journey

Accelerated expertise, hands-on assistance and mentoring
drives positive outcomes and quick time-to-value

Red Hat OpenShift AI Pilot


- Build a scalable unified AI/ML platform with Red Hat OpenShift and OpenShift AI and integrate all necessary tools.
- Enable entire organization to access tools and compute necessary to build and operationalize high quality AI/ML projects.

MLOps Foundation

- Provide critical tools to operationalize AI/ML models through their full lifecycle..
- Assist operations team to automate full ML pipelines and roll them out to their users.

Generative AI Incubator

- Create a consistent flexible platform focused on addressing Generative AI use cases.
- Introduce automation to assist with data ingestion, processing and model fine-tuning.





Training and Certification

Close skills gaps and hone your teams' Red Hat product expertise;
flexible learning options to meet your business needs

coming soon

Introduction to Red Hat OpenShift AI (AI262)	Red Hat OpenShift AI Administration (AI263)	Creating Machine Learning Models with Red Hat OpenShift AI (AI264)	Deploying Machine Learning Models with Red Hat OpenShift AI (AI265)	Automation using Data Science Pipelines (AI266)
Learn about the general architecture and main features of RHOAI and create some basic projects.	Get hands-on administration experience with RHOAI including install, upgrades, management and configuration.	Learn basic Machine Learning concepts and create and train Machine Learning models.	Deploy and serve Machine Learning models on RHOAI and troubleshoot deployed models.	Learn about and get hands-on experience with KubeFlow and Elyra pipelines.
Developing and Deploying AI/ML Applications on Red Hat OpenShift AI (AI267): 4 day instructor-led virtual training				
Red Hat Certified Specialist: OpenShift AI (EX267): 3 hour exam				



Technical Account Management

Proactive, trusted and expert technical advisor helps you
operate resilient and secure platforms

Enterprise TAM

Senior specialist with expertise across multiple Red Hat products ensure solutions meet and evolve with unique needs of AI/ML

OpenShift TAM

Highly technical OpenShift product specialist who proactively partners with your organization to help achieve business outcomes

OpenShift AI TAM **coming soon**

Technical OpenShift AI specialist who proactively partners with your organization to help achieve business outcomes