

Red Hat
Summit

Connect

Sessione pomeridiana
a cura di Red Hat

Agenda

Attenzione: le demo presentate in sessione non sono disponibili in questo file. Si rimanda al [canale YouTube di ImpresaCity](#) sul quale, nel mese di dicembre 2024, verrà pubblicato un reportage con tutte le riprese della sessione plenaria e di Red Hat.

- 14:00 - 15:00** **Intelligenza Artificiale: modelli aperti, sviluppo, rilascio e gestione in ambienti cloud ibridi**
- 15:00 - 15:30** **Il Machine Learning incontra Ansible Automation Platform: Un nuovo livello di automazione ITSM**
- 15:30 - 16:00** **Trusted Software Supply Chain.
Come rendere sviluppo applicativo e MLOps sicuri e tracciabili**
- 16:00 - 16:30** **Virtualizzazione Cloud Native, approccio dichiarativo e automazione del rilascio di workload virtualizzati**
- 16:30 - 17:00** **Dalla Strategia all'Azione: guidare la trasformazione digitale tramite la modernizzazione applicativa**

Red Hat
Summit

Connect

Intelligenza Artificiale

Modelli aperti, sviluppo, rilascio e gestione in ambienti cloud Ibridi

Daniele Zonca

Senior Principal
Software Engineer

Marco Caimi

Account Solution Architect


Francesco Rossi

Senior Specialist
Solution Architect

Growing demand for AI solutions and services

25%

of the overall tech spending will be dedicated to AI within the next 12 months



Red Hat's AI/ML engineering is 100% open source

Contributing to AI community projects since 2019





Integrated AI platform

Create and deliver gen AI and predictive models at scale across hybrid cloud environments.

Available as

- Fully managed cloud service
- Traditional software product on-site or in the cloud!



Model development

Bring your own models or customize Granite models to your use case with your data. Supports integration of multiple AI/ML libraries, frameworks, and runtimes.



Model serving and monitoring

Deploy models across any OpenShift footprint and centrally monitor their performance.



Lifecycle management

Expand DevOps practices to MLOps to manage the entire AI/ML lifecycle.

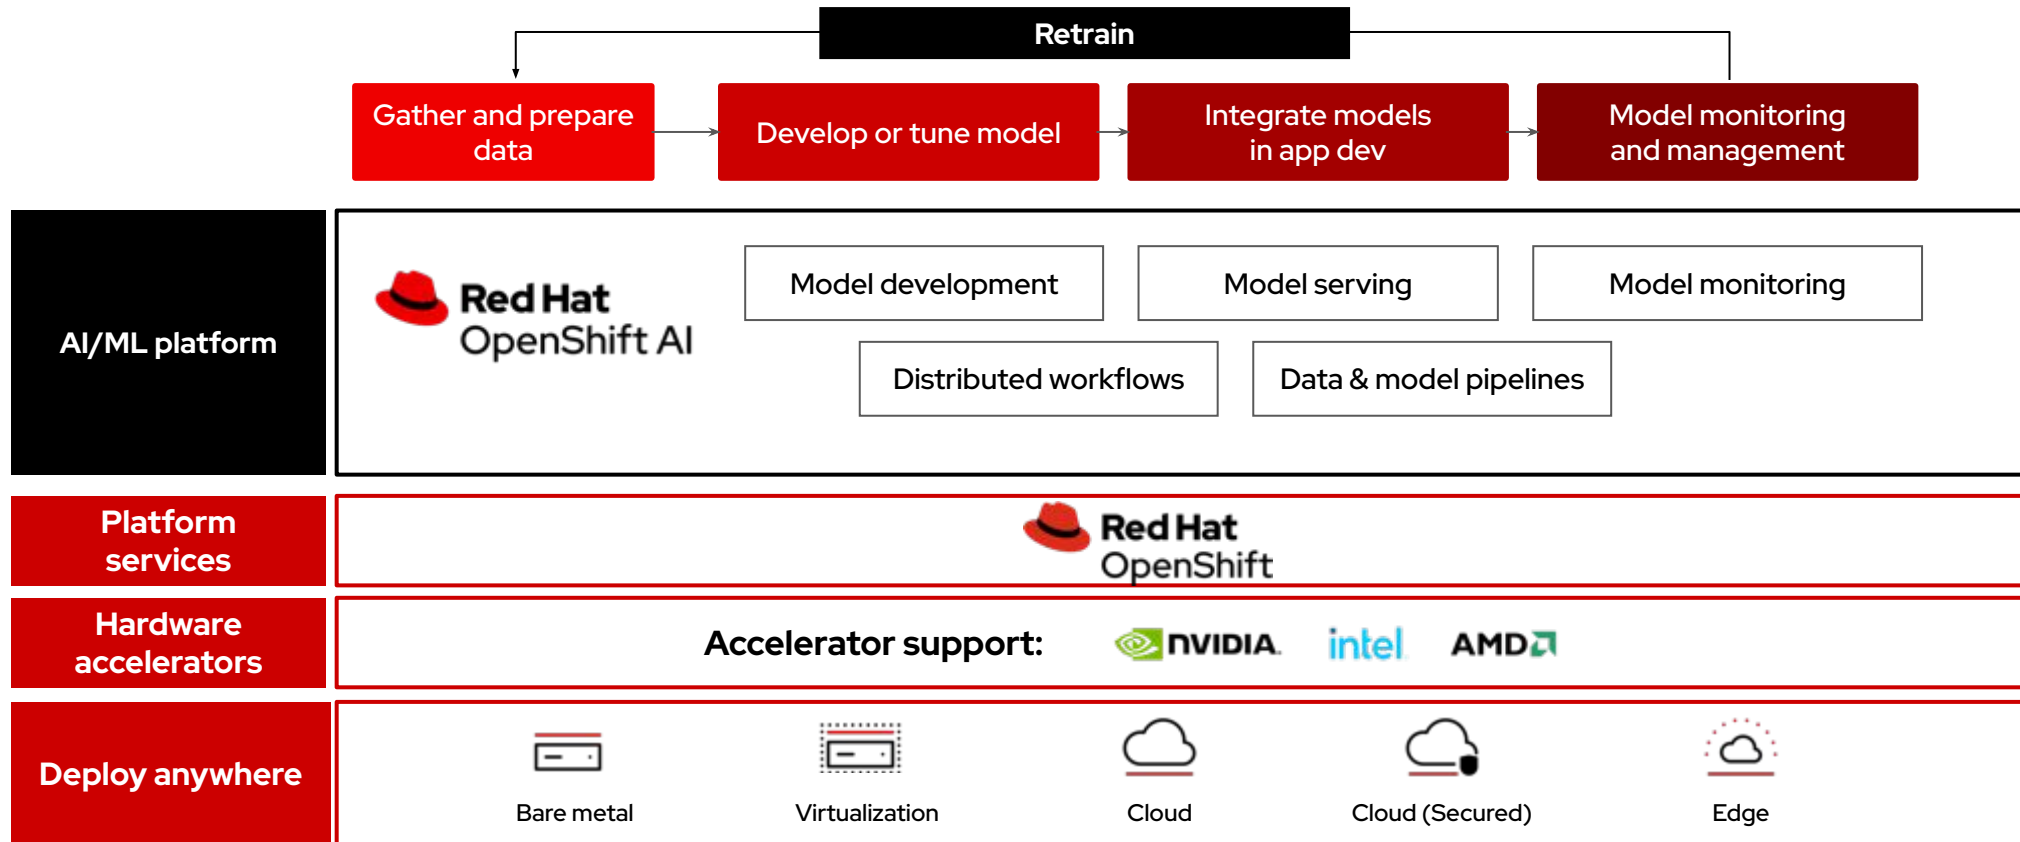


Resource optimization and management


Scale to meet workload demands of gen AI and predictive models. Share resources, projects, and models across environments.

Red Hat OpenShift AI

Red Hat's AI/ML platform for predictive and gen AI applications



Why containers, Kubernetes, and DevOps for AI/ML?



Agility

Respond quickly with automated compute resource management.



Flexibility

Provision AI/ML environments as and when you need them.



Portability

Develop and deploy ML models consistently across datacenter, edge, and public clouds.



Scalability

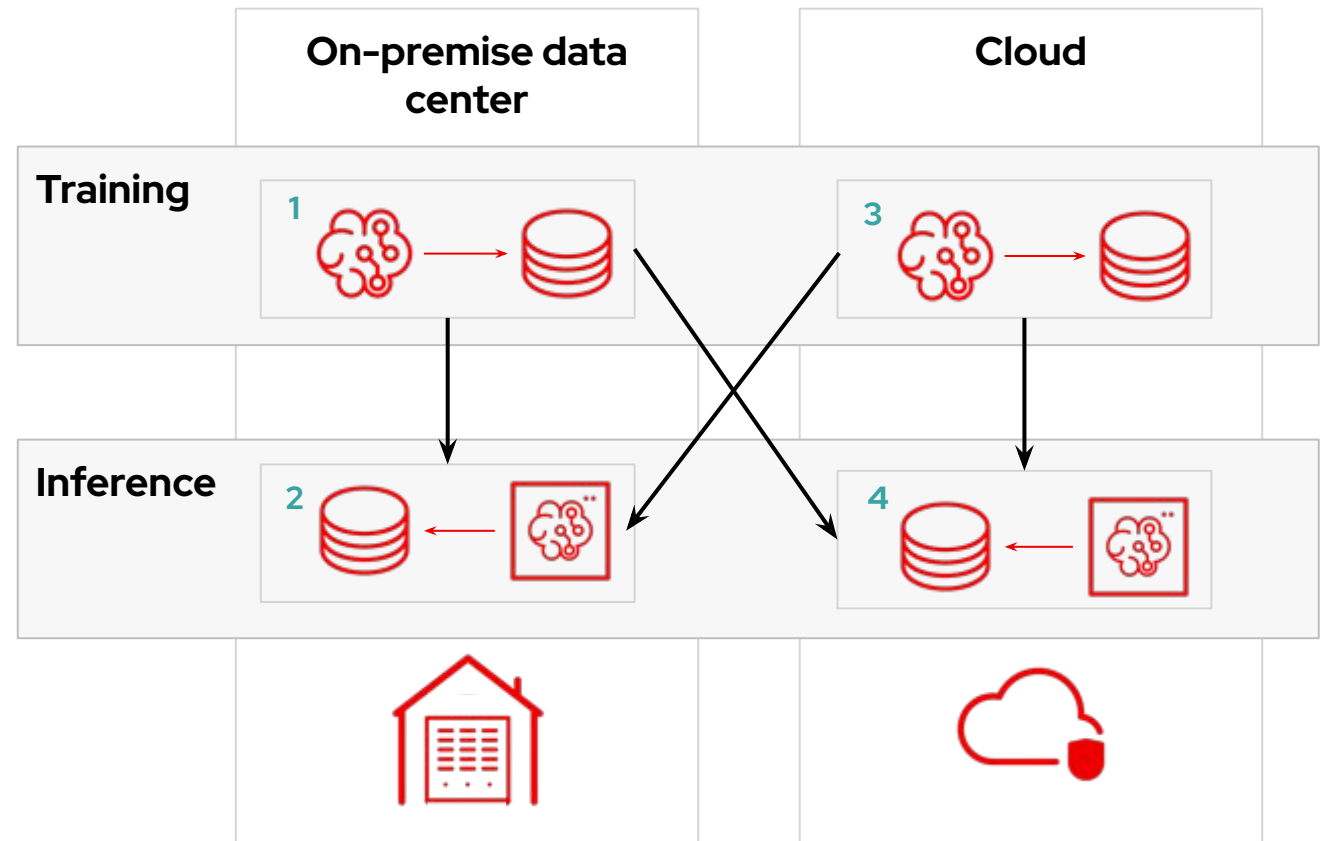
Autoscale and high availability of the AI/ML solution stack.

Why is Kubernetes a platform for AI?


Addressing data sovereignty, privacy and gravity

What you do should not dictate **where** you do it

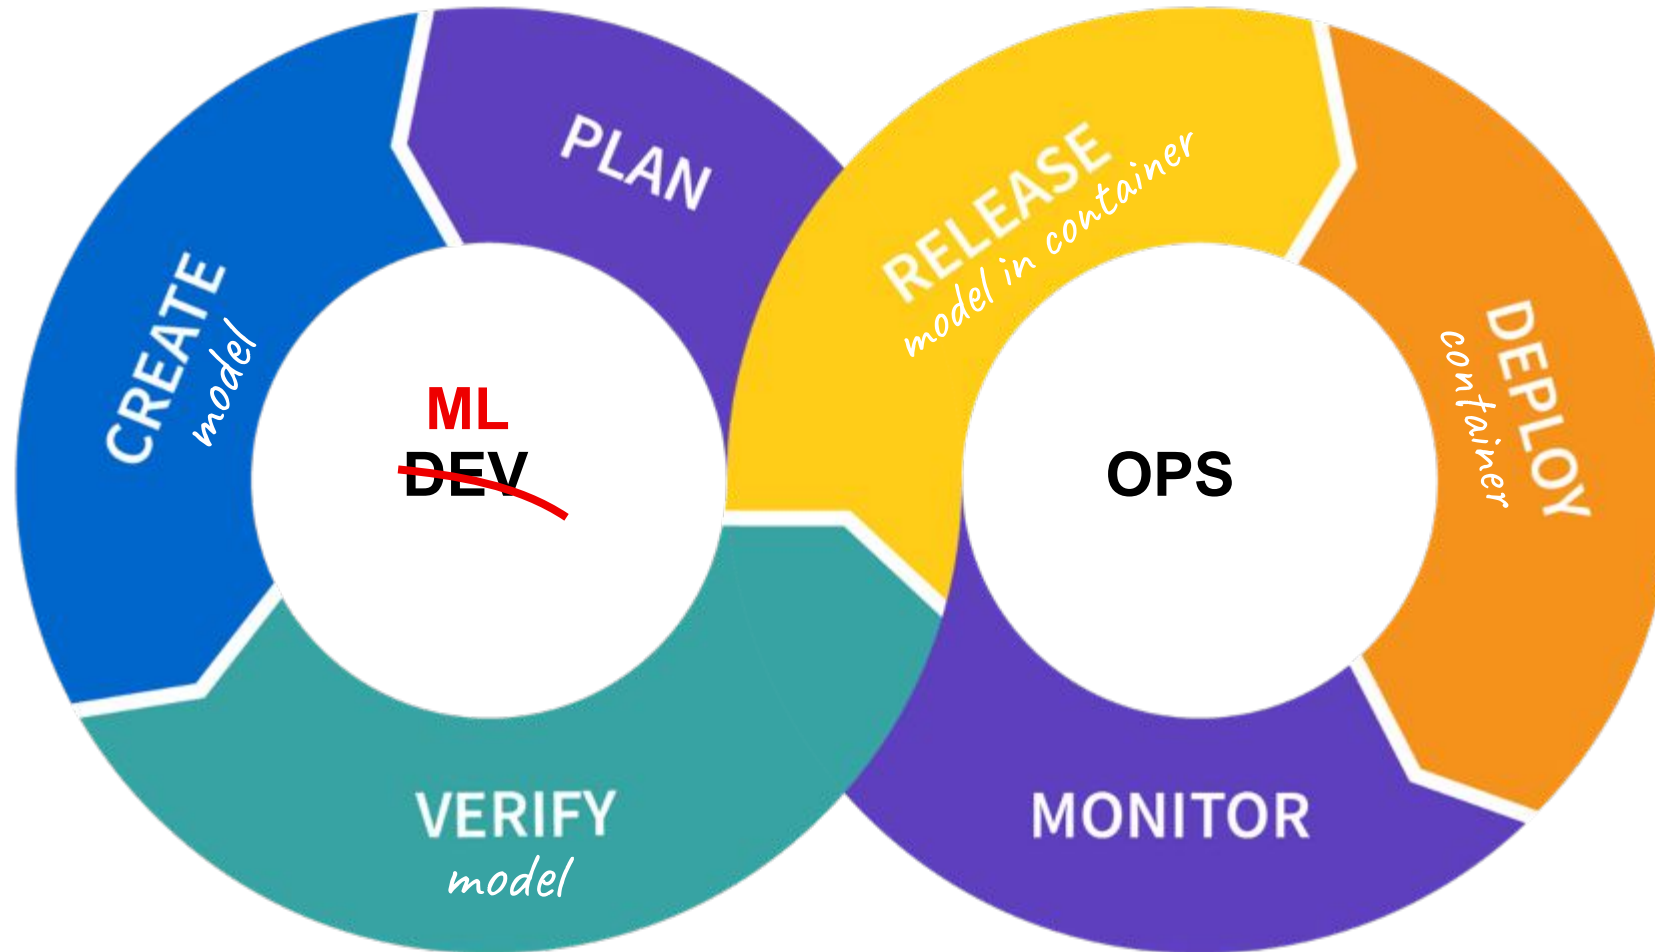
1. Data on-prem = Train on-prem
2. Data on-prem = Inference on-prem
3. Data in the cloud = Train on cloud
4. Data in the cloud = Inference on cloud



Kubernetes - A DevOps platform




Kubernetes - A ~~DevOps~~ MLOps platform



MLOps Workflow

From model development to serving through an API



Demo

Introduction

Part 1



Business Context

Insurance company that needs to improve claims processing

Proposed Improvements:

- Use various AI/ML tools and techniques to assist the claim adjusters
- Provide support for low-level, repetitive tasks
 - Point out areas in need of review
 - Help with parsing and data extraction
 - Reduce repetition fatigue

Using an LLM for text summarization

Allows for faster reading by the claims adjuster

Hi there, XYZ Insurance Company, I hope this email is okay and finds you okay. I had an accident, and I'm not exactly sure how to go about this, but I think it's something to do with a car accident claim, and my policy number is ABC12345, I think.

Okay, so here's what happened:

Accident Stuff:

Date and Time: Um, so this accident thing happened on, like, October 15th, 2023, at, um, 2:30 PM, I think.

Location: So, it happened at this place, um, the intersection of Elm Street and Maple Avenue, near Smith Park in Springfield, Illinois. I heard you might need some coordinates? They're like 39.7476° N and 89.6960° W or something. Hope that helps.

The Accidenty Part:

Weather Conditions: Well, the weather was kinda not great, I guess. It was like, cloudy and a bit rainy. And the road was wet, you know?

Traffic Conditions: There were some cars around, like, moderate traffic, I guess. And I was driving, like, the speed limit, which is, um, 35 mph, I think.

Car Details: So, my car is a Honda Accord, I think, and the other car involved was a Ford Escape. Yeah, that's right.

What Happened: So, I had the green light, and I was driving through the intersection, you know? But the other car, coming from the north or something, ran a red light and hit the front of my car on the passenger side. I didn't really have time to react or anything.

Injuries: Good news, no one got hurt really bad, but our cars got pretty messed up. The police came and made a report, and the officer had a badge number, I guess, it's 12345. I can get you the report if you need it.

Witness Stuff: There were a few people who saw this happen, and I got their names.

original, long-winded e-mail

human-readable summary

Summary:

The text is an email from John Smith to XYZ Insurance Company reporting a recent car accident involving his Honda Accord and a Ford Escape. The accident occurred on October 15, 2023, at approximately 2:30 PM at the intersection of Elm Street and Maple Avenue, near Smith Park, in Springfield, Illinois. John sustained minor injuries, but both vehicles sustained significant damage. He has taken photos of the accident scene and has the contact information of witnesses and the other party's insurance information. John is requesting that XYZ Insurance Company initiate a claim under his policy for the damages to his vehicle and is willing to provide any necessary documentation or information to process the claim efficiently.

Using an LLM for information extraction

Extract key pieces of information for better population of database

Hi there, XYZ Insurance Company,
I hope this email is okay and finds you okay. I had an accident, and I'm not exactly sure how to go about this, but I think it's something to do with a car accident claim, and my policy number is ABC12345, I think.

Okay, so here's what happened:

Accident Stuff:

Date and Time: Um, so this accident thing happened on, like **October 15th, 2023**, at, um, 2:30 PM, I think.

Location: So, it happened at this place, um, the intersection of Elm Street and Maple Avenue, near Smith Park in **Springfield, Illinois**. I heard you might need some coordinates? They're like 39.7476° N and 89.6960° W or something. Hope that helps.

The Accidenty Part:

Weather Conditions: Well, the weather was kinda not great, I guess. It was like cloudy and a bit rainy. And the road was wet, you know?

Traffic Conditions: There were some cars around, like, moderate traffic, I guess. And I was driving, like, the speed limit, which is, um, 35 mph, I think.

Car Details: So, **my car** is a Honda Accord, I think, and the other car involved was a Ford Escape. Yeah, that's right.

What Happened: So, I had the green light, and I was driving through the intersection, you know? But the other car, coming from the north or something, ran a red light and hit the front of my car on the passenger side. I didn't really have time to react or anything.

<input checked="" type="checkbox"/>	Date	Location	Item
<input checked="" type="checkbox"/>	15 oct 2023	Springfield, IL	Car

Using an LLM for sentiment analysis

Detect tone of text, and potentially act on it

Sentiment:

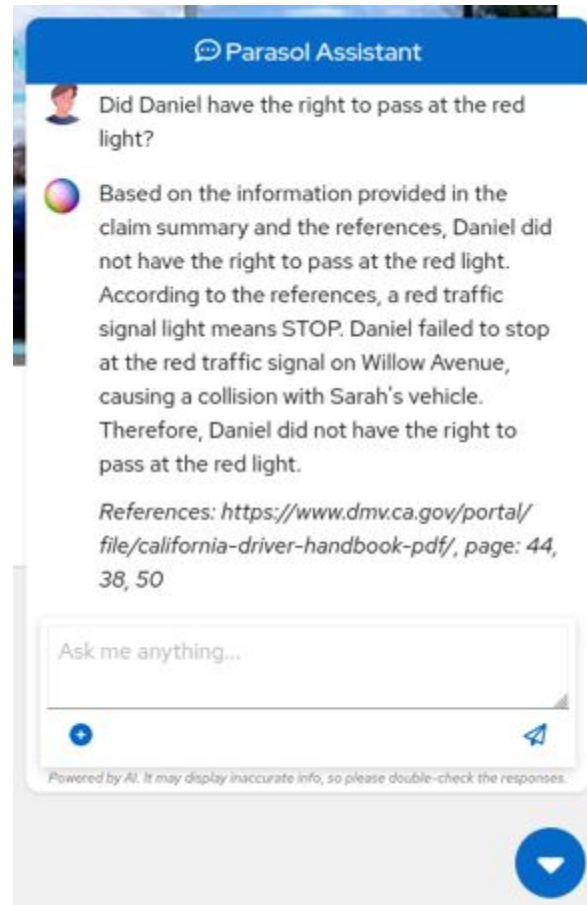
The sentiment of the person writing this text appears to be calm, assertive, and cooperative.

Sentiment:

The sentiment expressed in this text seems to be assertive and frustrated.

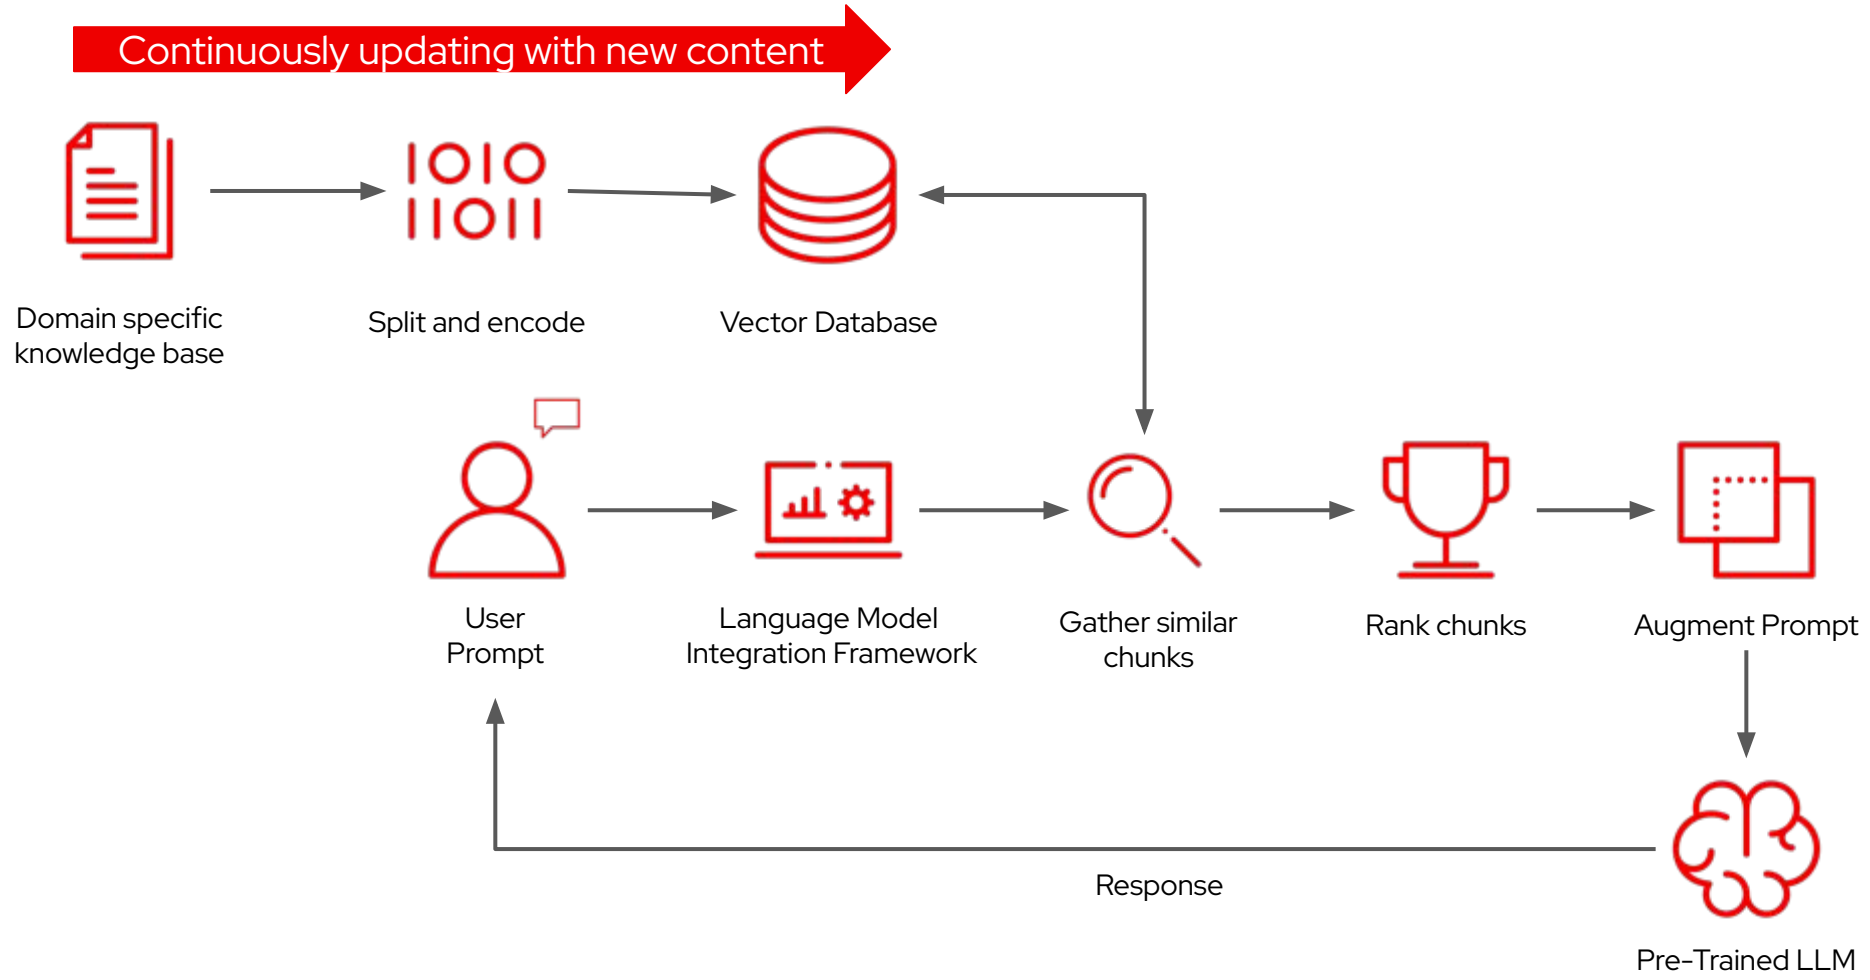
A virtual assistant to help operators

Provide guidance on Claim by consulting Driver Handbook knowledge (RAG)

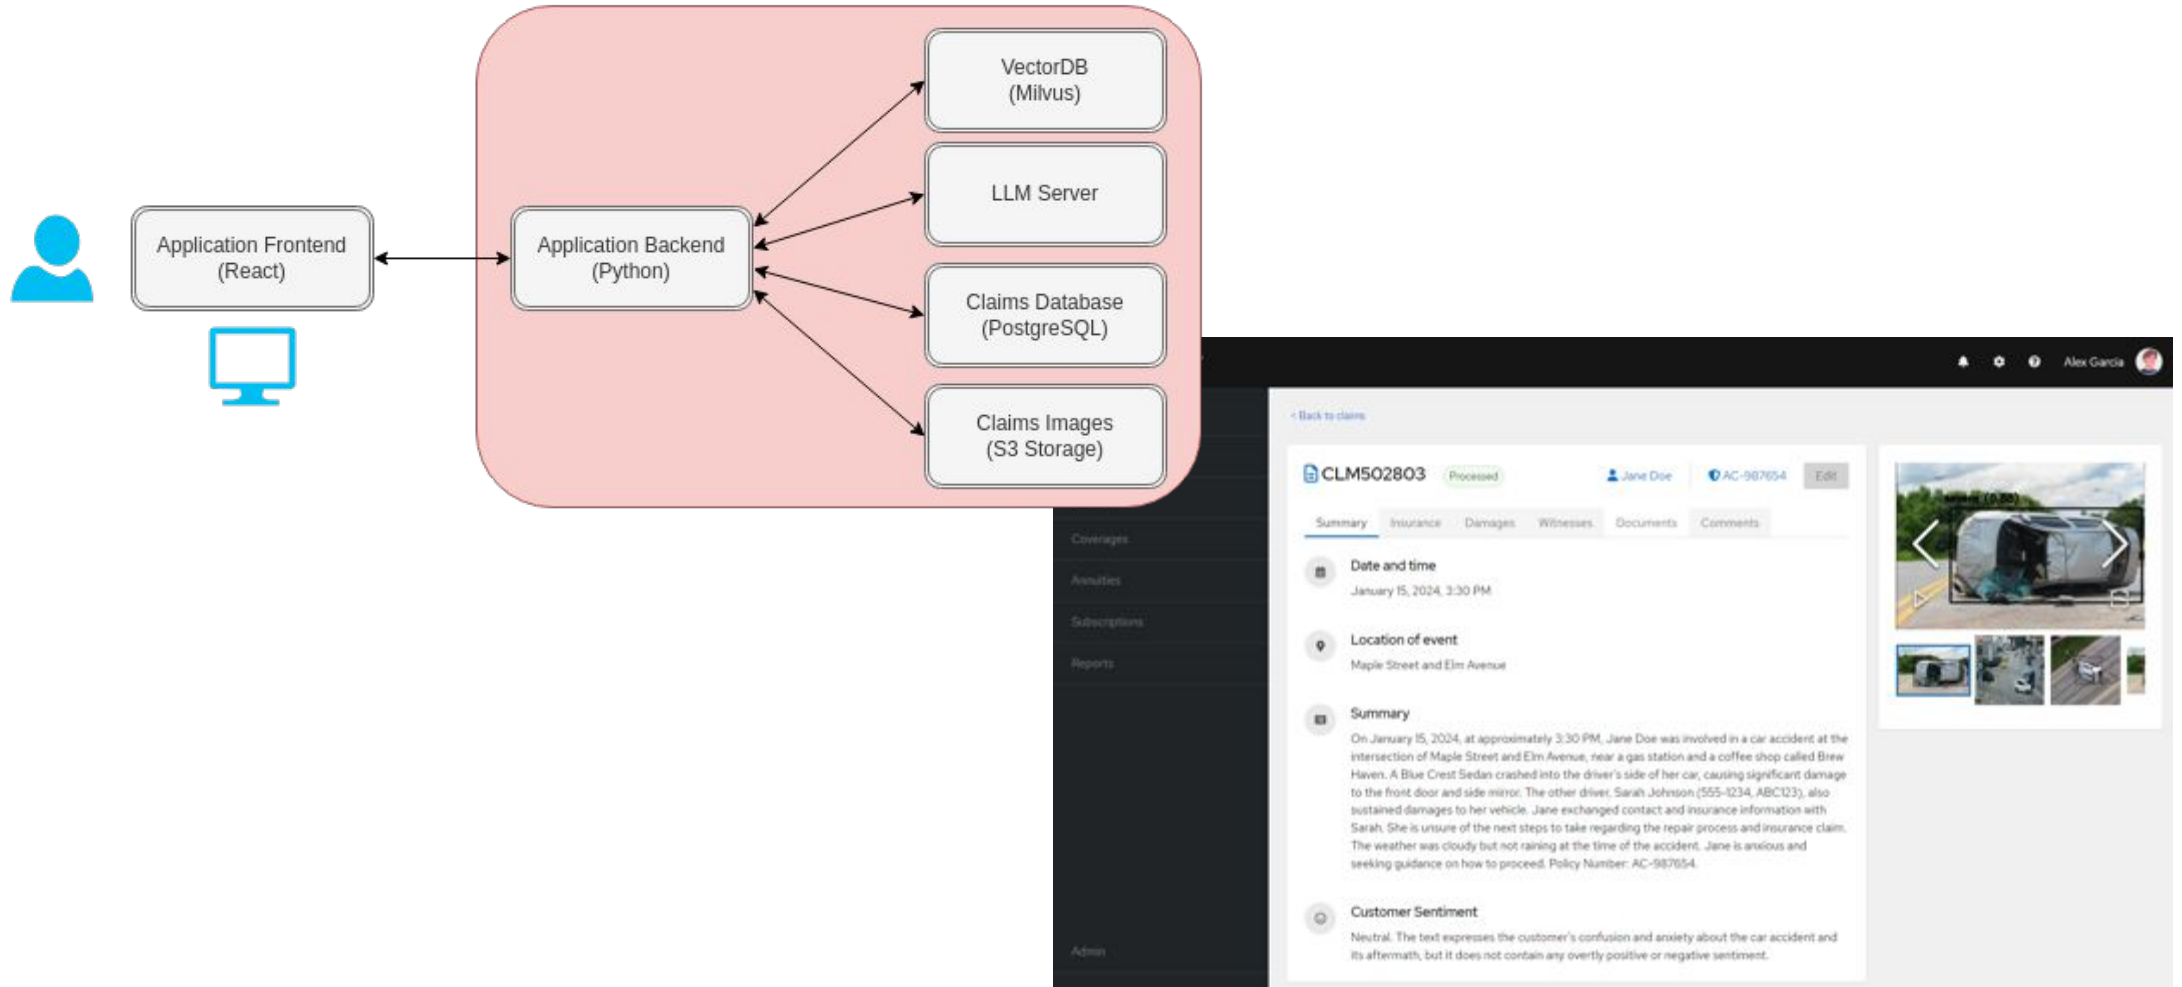


Retrieval Augmented Generation (RAG)

Helps the model to “look up” external information to improve generated text responses



Web Application to review/process claims



Demo

Introduction

Part 2



Demo

Introduction

Part 3



Using image recognition frame vehicle(s) and detect damage

Analyse images provided by customer and assessment of damage based on picture

```
[4]: #write code to get this info (for all of our detected boxes) in a loop, get object type, coords, probability
result = results[0]
for box in result.boxes:
    class_id = result.names[box.cls[0].item()]
    cords = box.xyxy[0].tolist()
    cords = [round(x) for x in cords]
    conf = round(box.conf[0].item(), 2)
    print("Object type:", class_id)
    print("Coordinates:", cords)
    print("Probability:", conf)
    print("----")
```

```
Object type: severe
Coordinates: [1, 22, 186, 218]
Probability: 0.88
----
```

```
[5]: #In the photo place boxes listing name, probability around each car (object type)
from PIL import Image
Image.fromarray(result.plot()[::-1])
```




Wrap-Up

Foundation Models, RAG and Fine-Tuning



Model Size vs. Performance – Large vs. Small LLMs

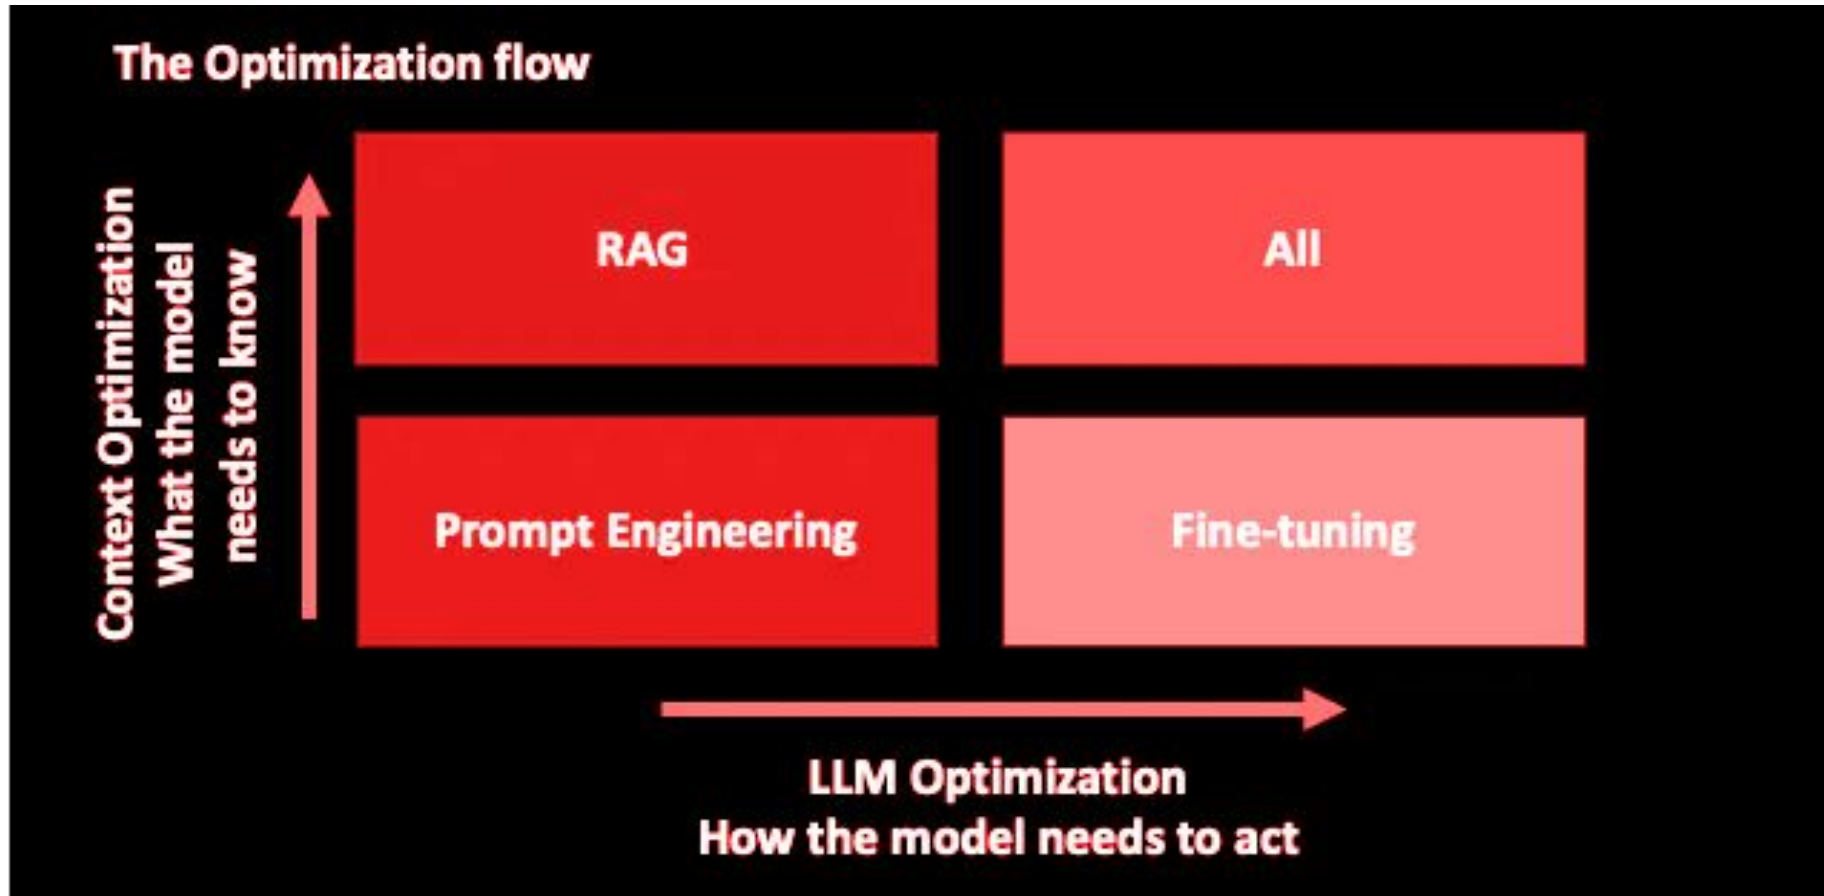
IBM Granite Models target Small LLMs aligned to enterprise data/use case



¹Model performance Calculation: Average of commonly utilized LLM benchmarks – MMLU (Multitask accuracy), HellaSwag (Reasoning), HumanEval (Python coding tasks), BBHard (Probing models for future capabilities), GSM8K (Grade school math)

Source: [LLM Leaderboard 2024 \(vellum.ai\)](https://www.vellum.ai)

Optimizing the performance of LLMs

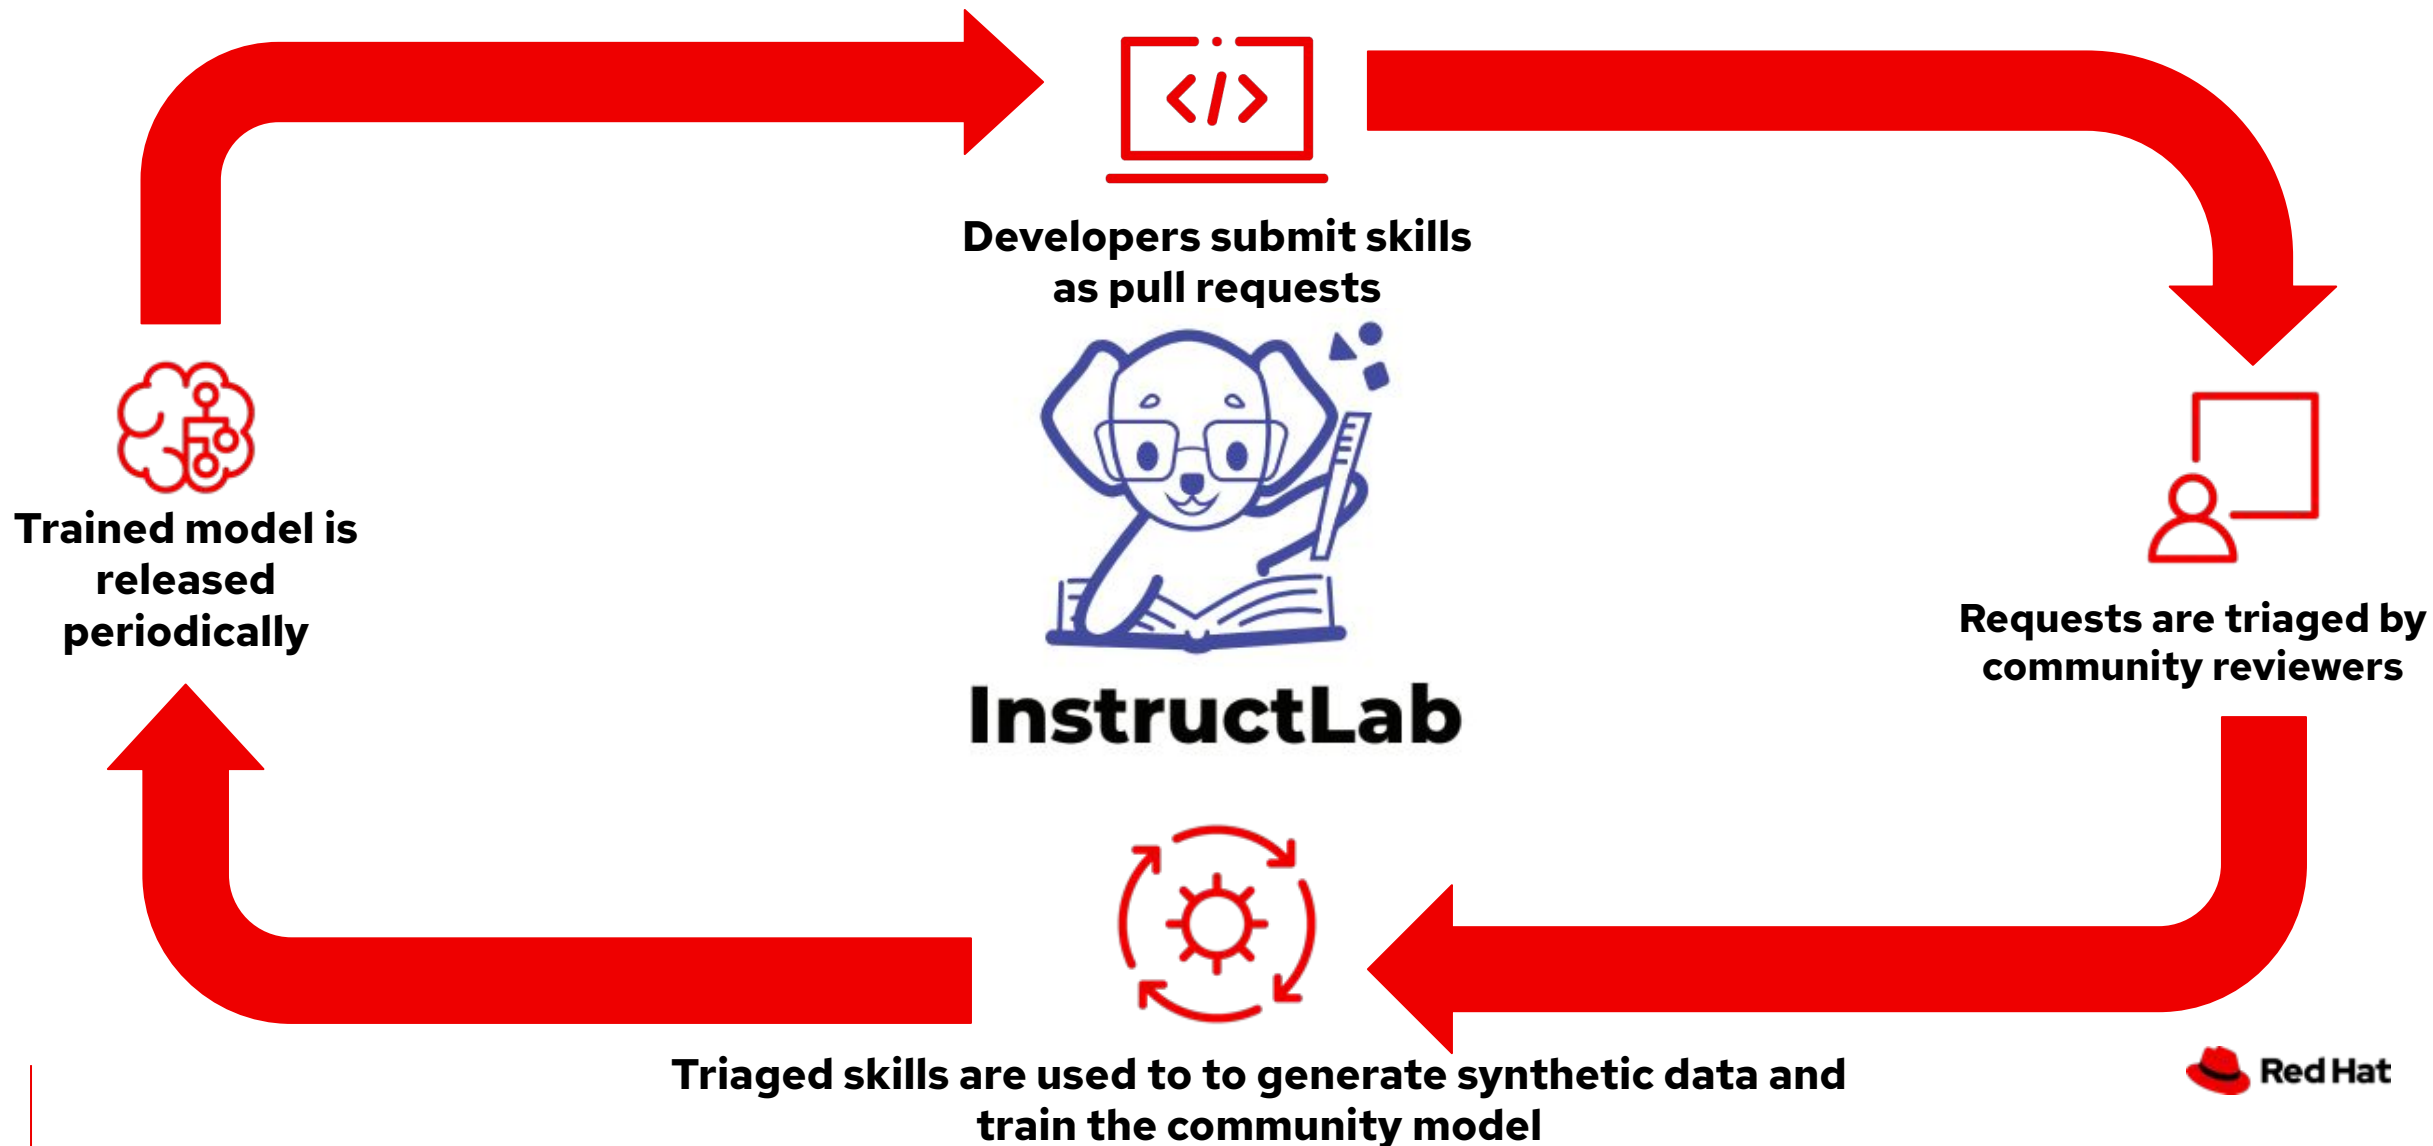


Source:

[Optimizing LLMs: Best Practices](#) (Luv Verma on Medium, January 6, 2024)

Introducing: **InstructLab**

Open source community project for GenAI model development



InstructLab vs. Alternative Model Alignment Approaches

InstructLab provides more accessible fine tuning & compliments RAG (RAFT pattern)

RAG

(Retrieval Augmented Generation)




Enhance Gen AI model generated text by retrieving relevant information from external sources, improving accuracy and depth of model's responses.

NEW

INSTRUCTLAB

(Large-scale Alignment for chatBots)



Leverage a taxonomy-guided synthetic data generation process and a multi-phase tuning framework to improve model performance.

Fine-tuning

(Fine Tuning)



Adjust a pre-trained model on specific tasks or data, improving its performance and accuracy for specialized applications without full retraining.

Red Hat
Summit

Connect

Thank you



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



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



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



twitter.com/RedHat

Red Hat
Summit

Connect

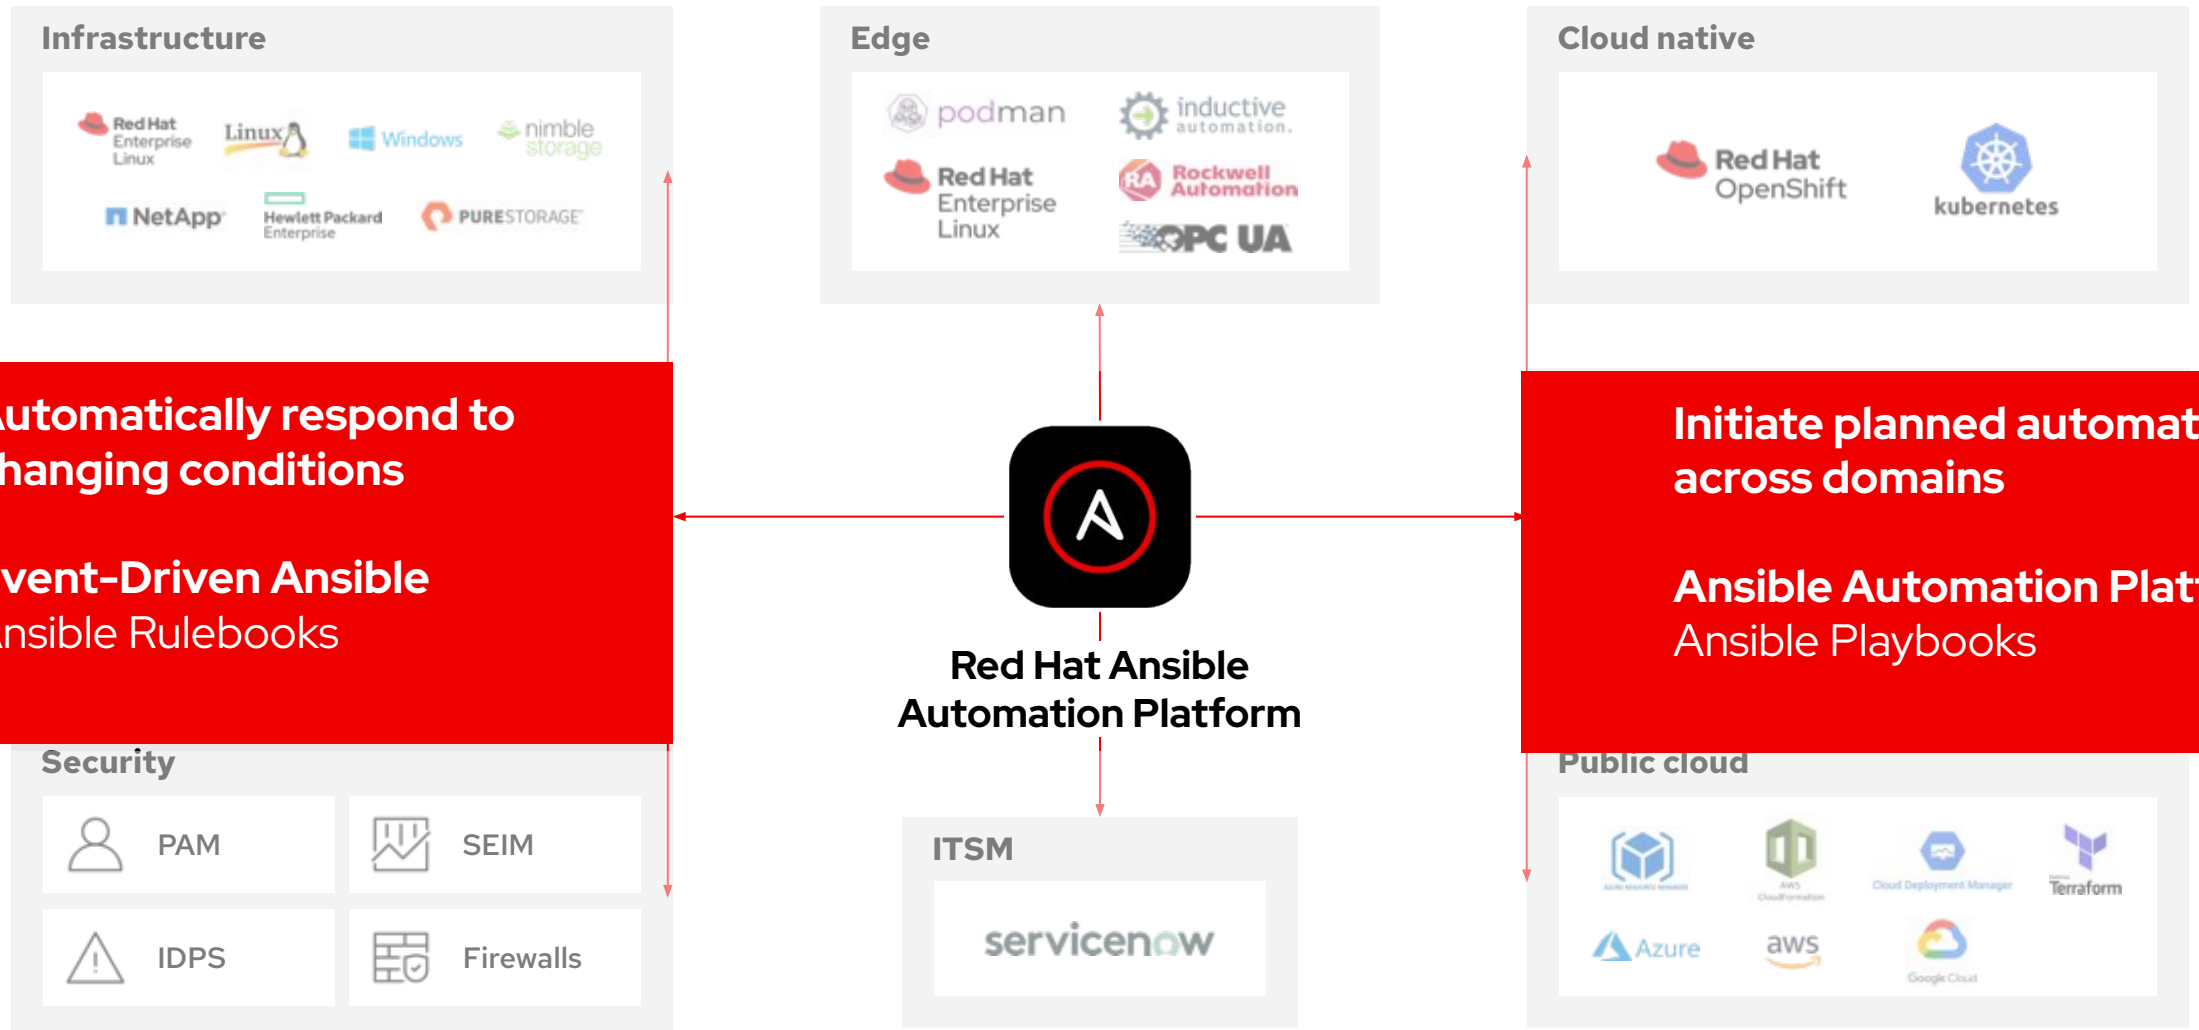
Il Machine Learning incontra Ansible Automation Platform

Un nuovo livello di automazione ITSM

Alessandro Arrichiello
Solution Architect
ale@redhat.com

Pietro Bertera
Solution Architect
pbertera@redhat.com

Single enterprise platform now with **more automation options**

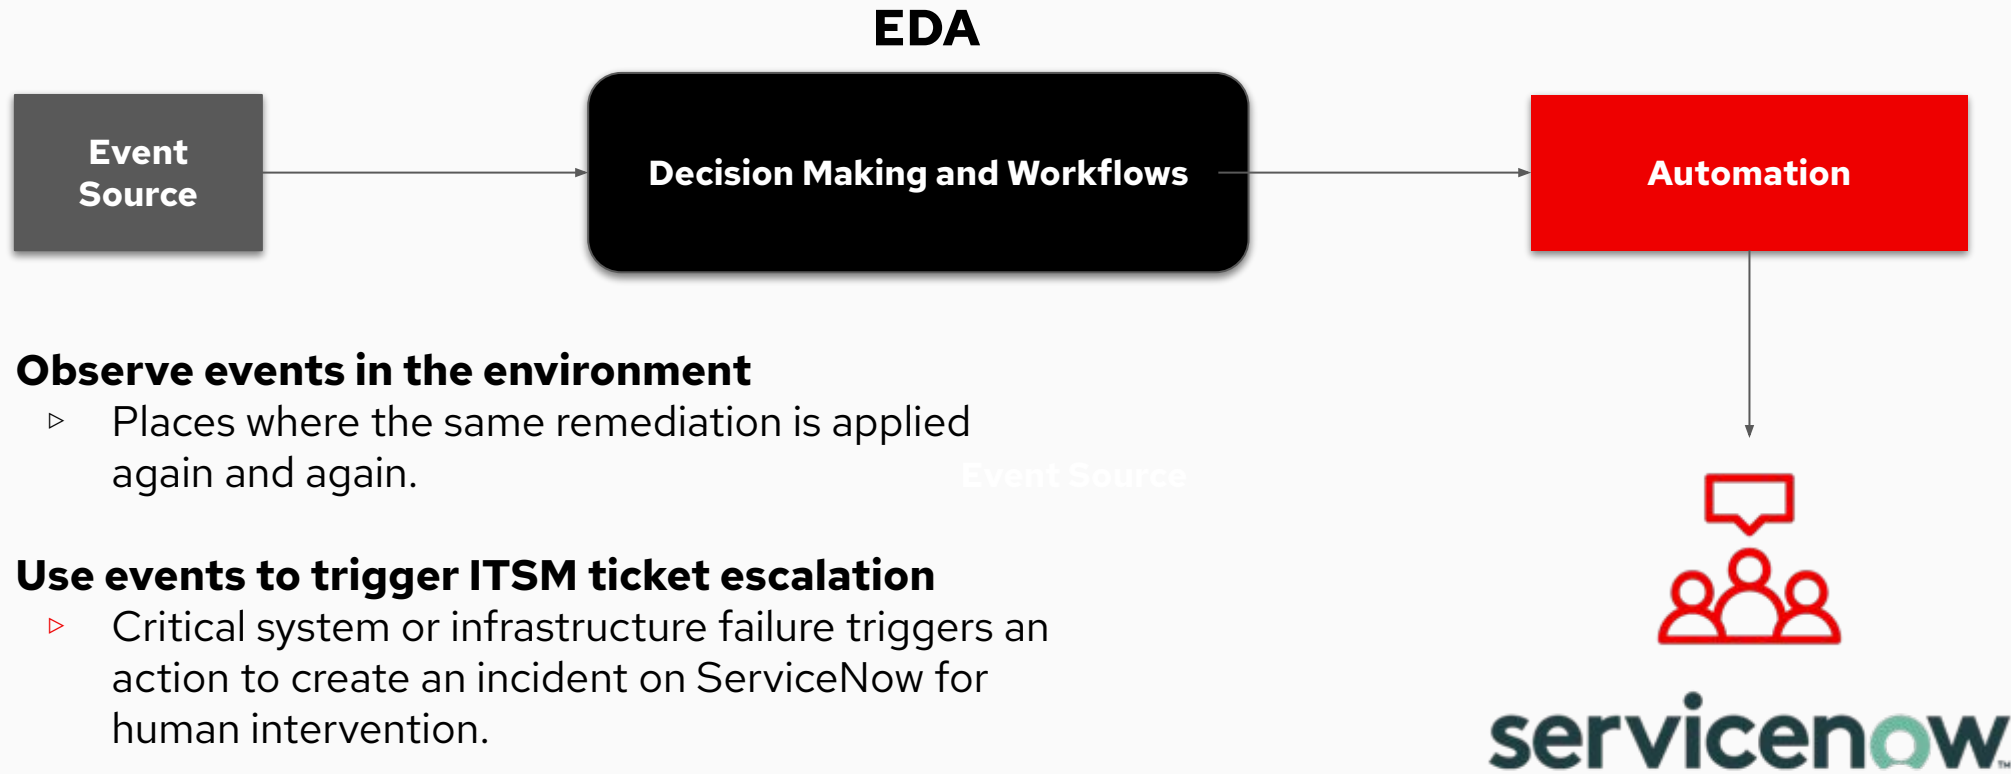


Event-Driven Automation and ITSM Integration



Event-Driven Ansible and ServiceNow ITSM integration


Events to human observation



- ▶ **Observe events in the environment**
 - ▷ Places where the same remediation is applied again and again.
- ▶ **Use events to trigger ITSM ticket escalation**
 - ▷ Critical system or infrastructure failure triggers an action to create an incident on ServiceNow for human intervention.
- ▶ **Update ServiceNOW CMDB**
 - ▷ Infrastructure changes can be observed and used to trigger ServiceNow to update its inventory

A gradual approach to ITSM Automation

That does not require a change to internal business processes



Manual Resolution via ServiceNow

Human operators identify and match the viable automation



- ▶ **Human operators interact with ITSM**
 - ▷ Analyze the informations on ServiceNow and execute a viable automation on AAP
- ▶ **AAP can then execute the automation and report**
 - ▷ After executing the automation Ansible Automation Platform can report back the status on ServiceNow incidents

ServiceNow ITSM integration

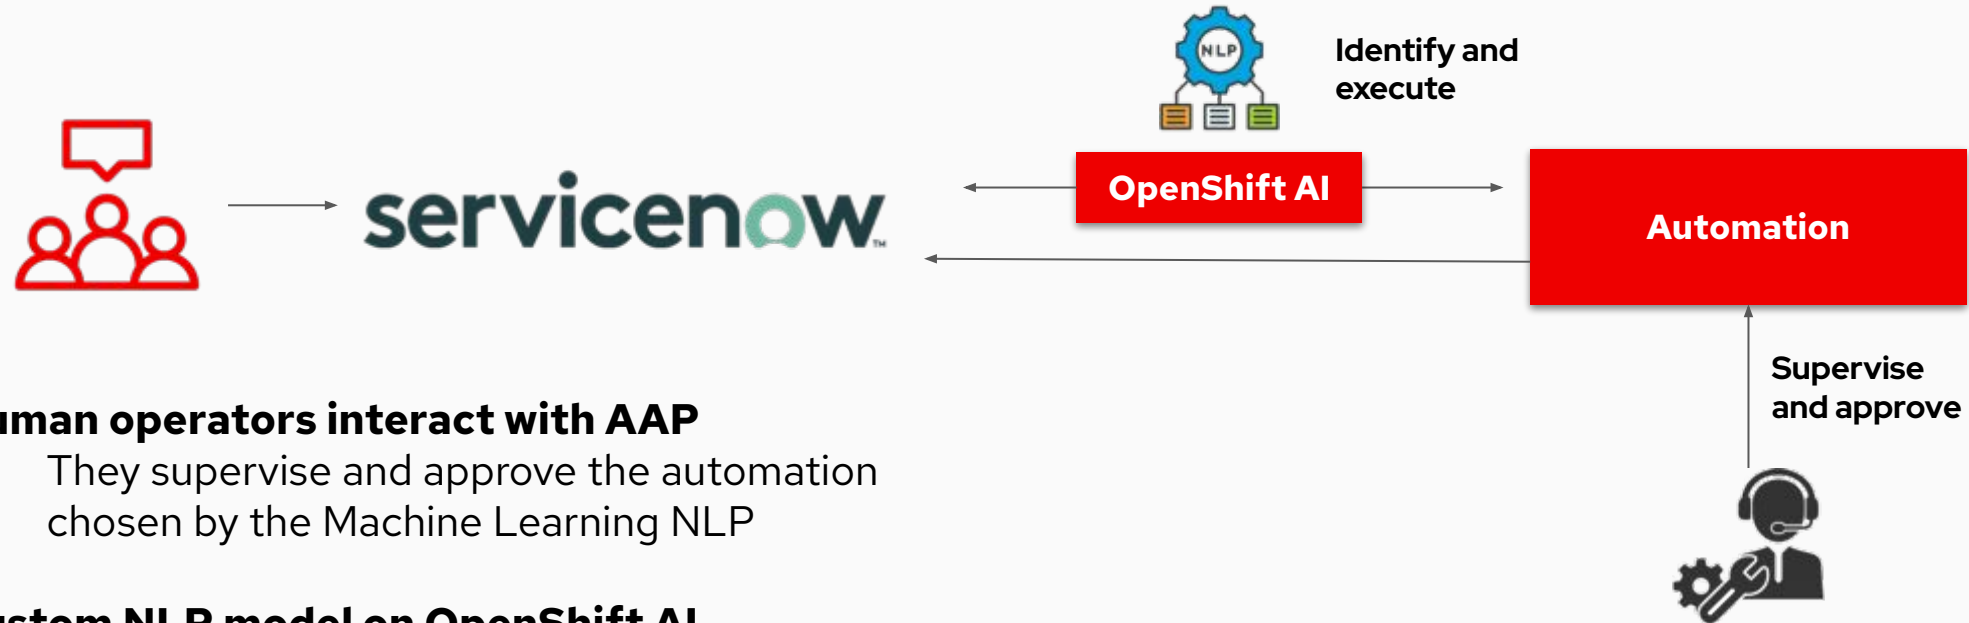
Human operator driven using just ITSM



- ▶ **Human operators interact with ServiceNow**
 - ▷ They work on the ServiceNow incidents and have integrations on the interface to call Ansible Automation Platform (AAP)
- ▶ **AAP can then execute the automation and report**
 - ▷ After executing the automation Ansible Automation Platform can report back the status on ServiceNow incidents


AI/ML Resolution

Natural Language Processing for executing the proper Automation



- ▶ **Human operators interact with AAP**
 - ▷ They supervise and approve the automation chosen by the Machine Learning NLP
- ▶ **Custom NLP model on OpenShift AI**
 - ▷ OpenShift AI is serving the model trained on historical data extracted from ServiceNow (ITSM) to classify the text of the ticket and trying to match a viable automation on AAP
- ▶ **AAP can then execute the automation and report**
 - ▷ After executing the automation Ansible Automation Platform can report back the status on ServiceNow incidents

Key Prerequisites For ITSM Automation



Key Prerequisites are vital for **ITSM Automation Development**



Data Gathering and Categorization

- Historical ServiceNow data must be collected and categorized to understand incident patterns, enabling effective automation development and AI/ML model training.

Ansible Automation Playbook Development

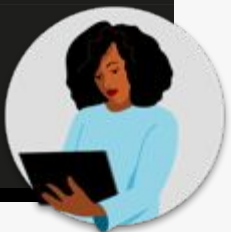
- Based on incident data analysis, Ansible playbooks should be created to automate the most frequent and time-consuming tasks, maximizing the return on automation investment.

Ansible Lightspeed enhances the automation development experience

Integrated Development Experience

Ansible content creation

```
10 tasks:
11 - name: Include redhat.rhel_system_roles.cockpit
12   ansible.builtin.include_role:
13     name: redhat.rhel_system_roles.cockpit
14
15 - name: Copy files/cockpit.conf to /etc/cockpit/
16   ansible.builtin.copy:
17     src: ./files/cockpit.conf
18     dest: /etc/cockpit/
19     owner: root
20     group: root
21     mode: '0644'
22
23 # - name: Restart cockpit service
24
25 # - name: Allow cockpit through firewall
```



VS Code
extension

Ansible Lightspeed

Best - Practices
Anonymize
Post Processing

IBM Watson Code Assistant

1010 1010 1010
11011 11011 11011


Red Hat OpenShift AI

Ansible Content Tools

An Open and Collaborative Platform for AI and Apps

AI/ML Resolution


Natural Language Processing for executing the proper Automation




- ▶ **Human operators interact with AAP**
 - ▷ They supervise and approve the automation chosen by the Machine Learning NLP
- ▶ **Custom NLP model on OpenShift AI**
 - ▷ OpenShift AI is serving the model trained on historical data extracted from ServiceNow (ITSM) to classify the text of the ticket and trying to match a viable automation on AAP
- ▶ **AAP can then execute the automation and report**
 - ▷ After executing the automation Ansible Automation Platform can report back the status on ServiceNow incidents

Demo Workflow

Training the model, serving it and let tickets to be classified to be resolved by the Automation Platform



AI Models and Automation in the same platform




Ansible workflows: solving complex problems

What is it?

- ▶ Workflows enable the creation of powerful holistic automation, chaining together multiple pieces of automation and events
- ▶ Simple logic inside these workflows can trigger automation depending on the success or failure of previous steps
- ▶ Add approvals to your workflows to enhance governance
- ▶ Integrate other systems, such as ITSM to fit with your existing controls and processes

JOBS / 363 - Wave Workflow Example



Supercharge your Ansible Workflows!

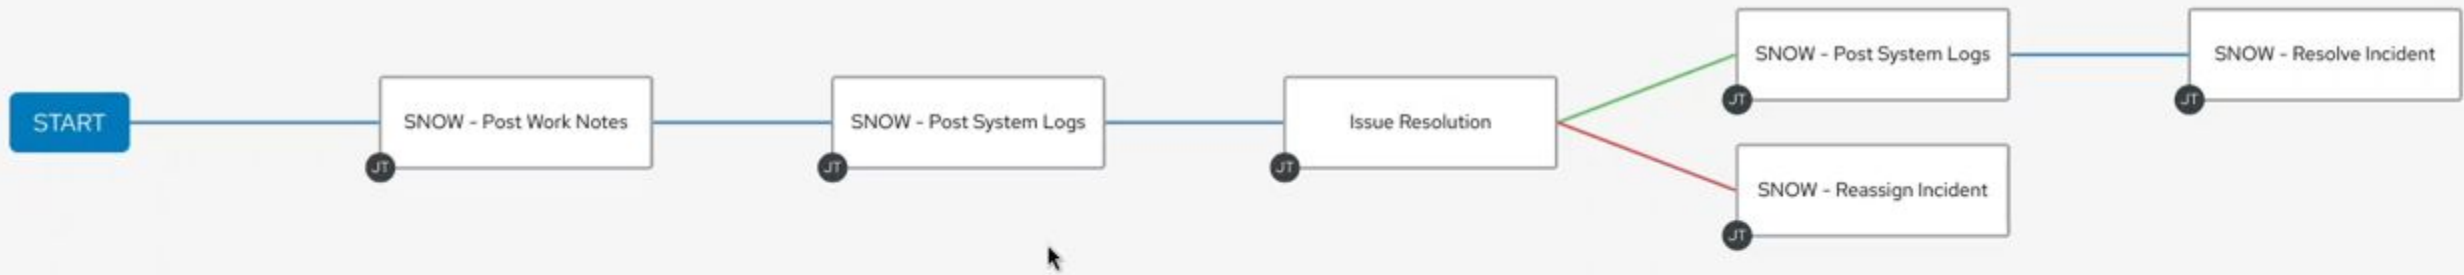
Hybrid Cloud | Red Hat OpenS | 2_save_model | ai2aap-snow-i | ai2aap-snow-i | ServiceNow | Red Hat Demo | Ansible Autom

ansible-1.2rj8f.sandbox1049.opentlc.com/#/templates/workflow_job_template/12/visualizer

Issue Resolution Workflow

Total Nodes 6

Save



Where to go **next**



Learn more

- ▶ Transforming ITSM with Ansible Automation: A Gradual Approach
- ▶ Red Hat Developer Sandbox: Your Free OpenShift AI Playground
- ▶ How to train a BERT machine learning model with OpenShift AI
- ▶ Revolutionize IT automation with the new ServiceNow integration



Get started

- ▶ Self-paced labs
- ▶ Evals
- ▶ console.redhat.com



Get serious

- ▶ Red Hat Automation Adoption Journey
- ▶ Red Hat Training
- ▶ Red Hat Consulting

What's new in Red Hat Ansible Automation Platform 2.5?

Join us for **the webinar** on **November 20th**
to unlock the potential of automation



Red Hat
Summit

Connect

Thank you



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



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



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



twitter.com/RedHat

Red Hat
Summit

Connect

Trusted Software Supply Chain

Come rendere sviluppo applicativo e MLOps sicuri e tracciabili

Matteo Combi

Senior Specialist Solution
Architect


Matteo Grimaldi

Senior Account Solution
Architect

Matteo Mortari

Principal Software Engineer

Why we are here today



Software supply chain attacks: a matter of when, not if

Ransom paid but a mere fraction to the overall
downtime and recovery costs of a data breach



742%

average annual increase in
software supply chain
attacks over the past 3 years¹

20%

data breaches are due to a
compromised software
supply chain²

78%


have initiatives to
increase collaboration
between DevOps and
Security teams³

92%

say enterprise open source
solutions are important as
their business accelerates
to the open hybrid cloud⁴

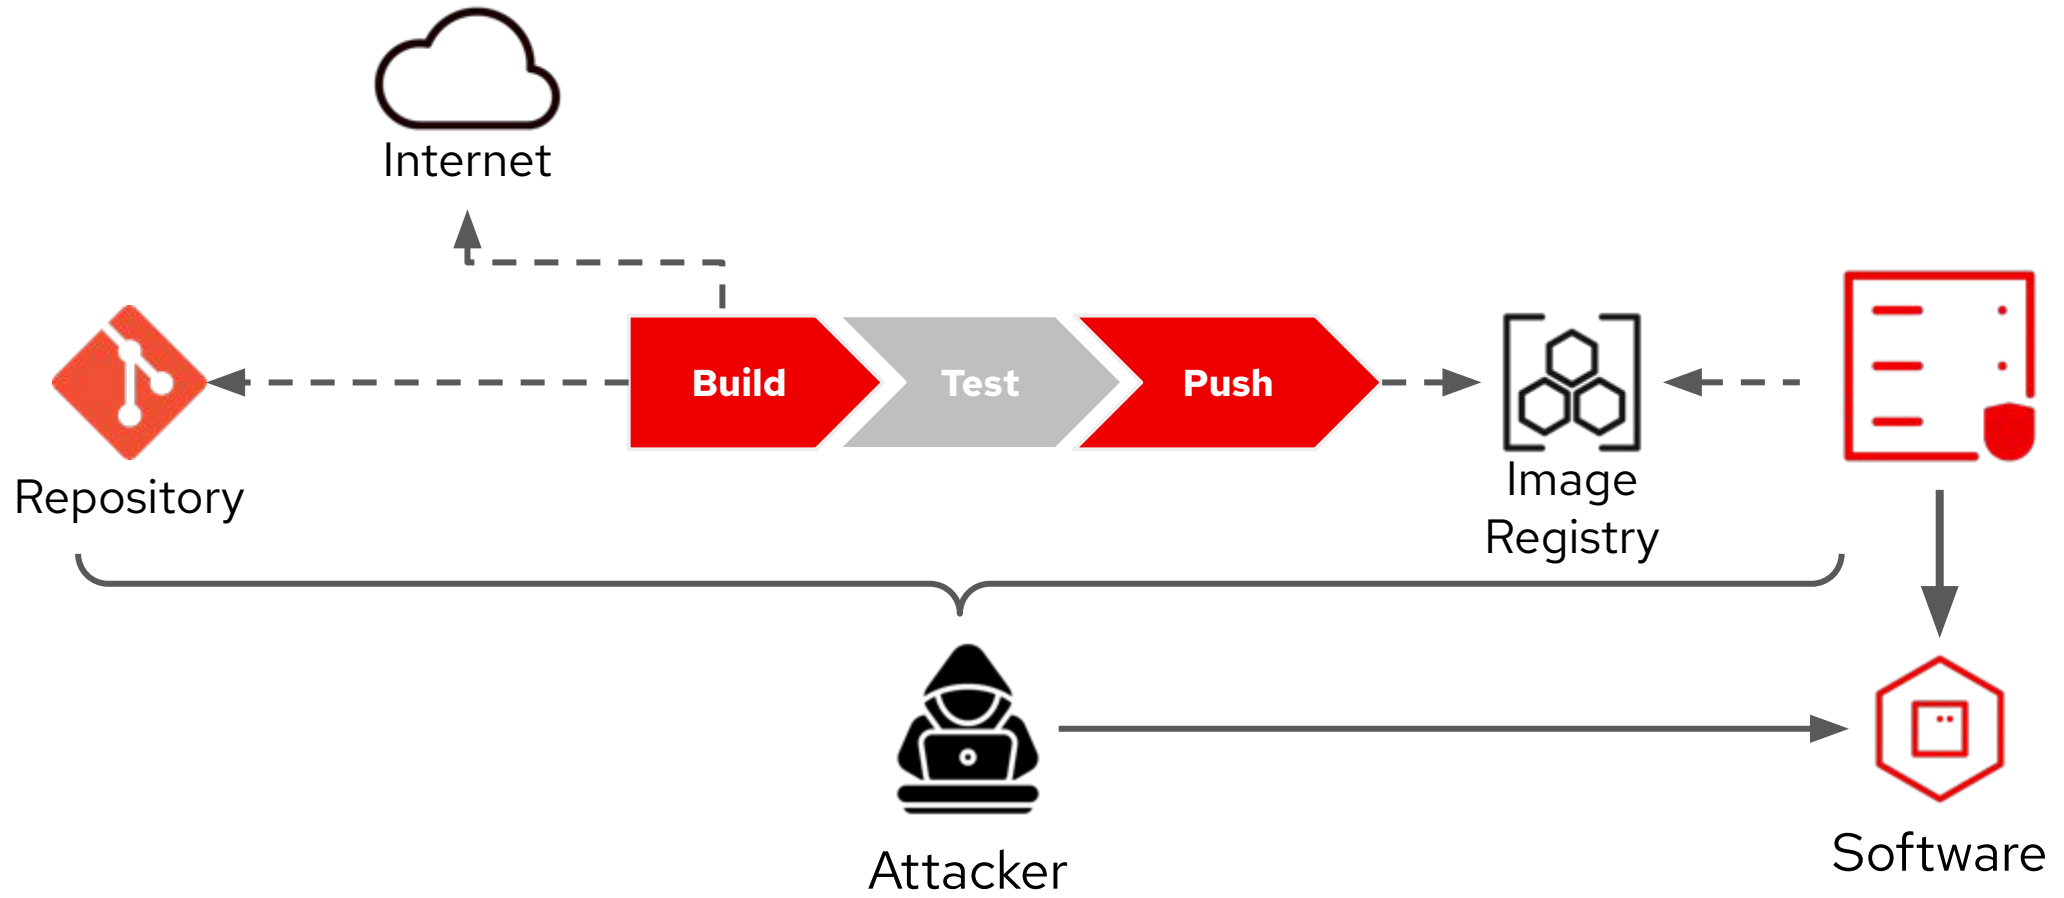
Supply Chain Attack

Hardware



Supply Chain Attack

Software




Growing Attack Surfaces




Supply Chain Attack

Attack surfaces



Supply Chain Attack

Attack surfaces

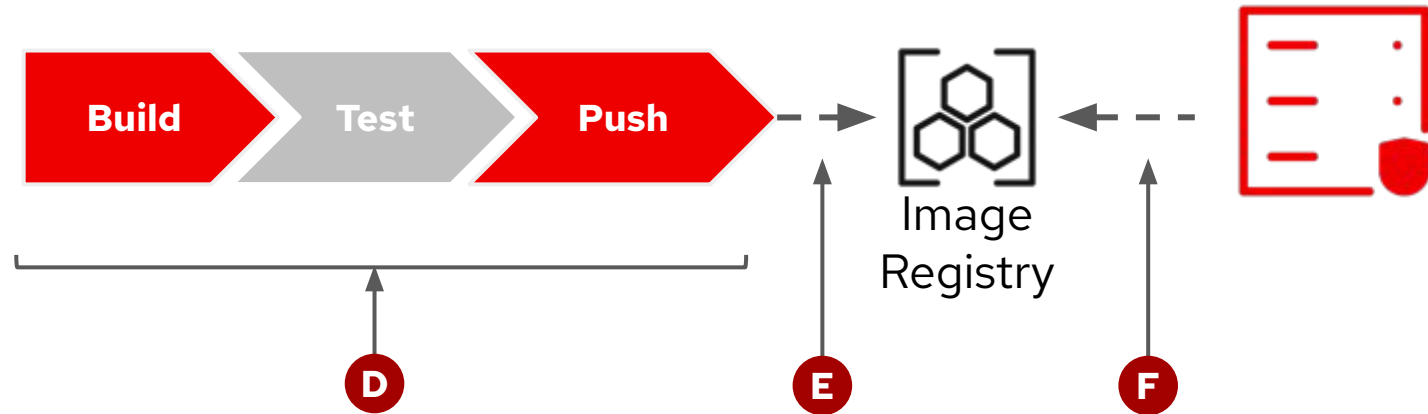


- A** Bypass code review or compromised source control system
- B** Source injection / alteration
- C** Vulnerable dependencies injection

Supply Chain Attack

Attack surfaces

- D** Compromised or bypassed CI/CD system
- E** Source injection / alteration
- F** Container image alteration



What is SLSA ?

Going beyond application security testing



SLSA stands for **Supply Chain Levels for Software Artifacts**.

SLSA is a security framework and a common language for improving software security by **ensuring supply chain integrity**.

It is a cross-industry collaboration, maintained as part of the Open Source Security Foundation, that is based on concepts that have been used **since 2013**.

Development-time controls

Shifting left security and compliance



Preventing Mistakes

Automated Build Process

Generated provenance about source, build process, artifact and dependencies

Preventing tampering after the build

Generated, signed and verifiable provenance

Preventing tampering during the build

Detecting and preventing vulnerabilities at **code time**

Preventing non compliant software at **code time**

Development-time controls

Shifting left security and compliance



Preventing Mistakes

Automated Build Process

Generated provenance about source, build process, artifact and dependencies

Preventing tampering after the build

Generated, signed and verifiable provenance

Preventing tampering during the build

Detecting and preventing vulnerabilities at **code time**

Preventing non compliant software at **code time**

SLSA concepts

How to move forward



SBOM

Or Software Bill of Materials, it lists all the components that went into making a given piece of software

Provenance

It is the recording of origin, history and who made changes

Attestation

Authenticated statement (metadata) about a software artifact or collection of software artifacts

Accelerate Innovation that Safeguards User Trust

Delivered with integrated security guardrails at every phase of the software development lifecycle



Red Hat
Trusted Software
Supply Chain


Developer self-service hub with pre-integrated security guardrails: artifact signatures, attestations & SBOMs

Accelerate Innovation that Safeguards User Trust

Delivered with integrated security guardrails at every phase of the software development lifecycle

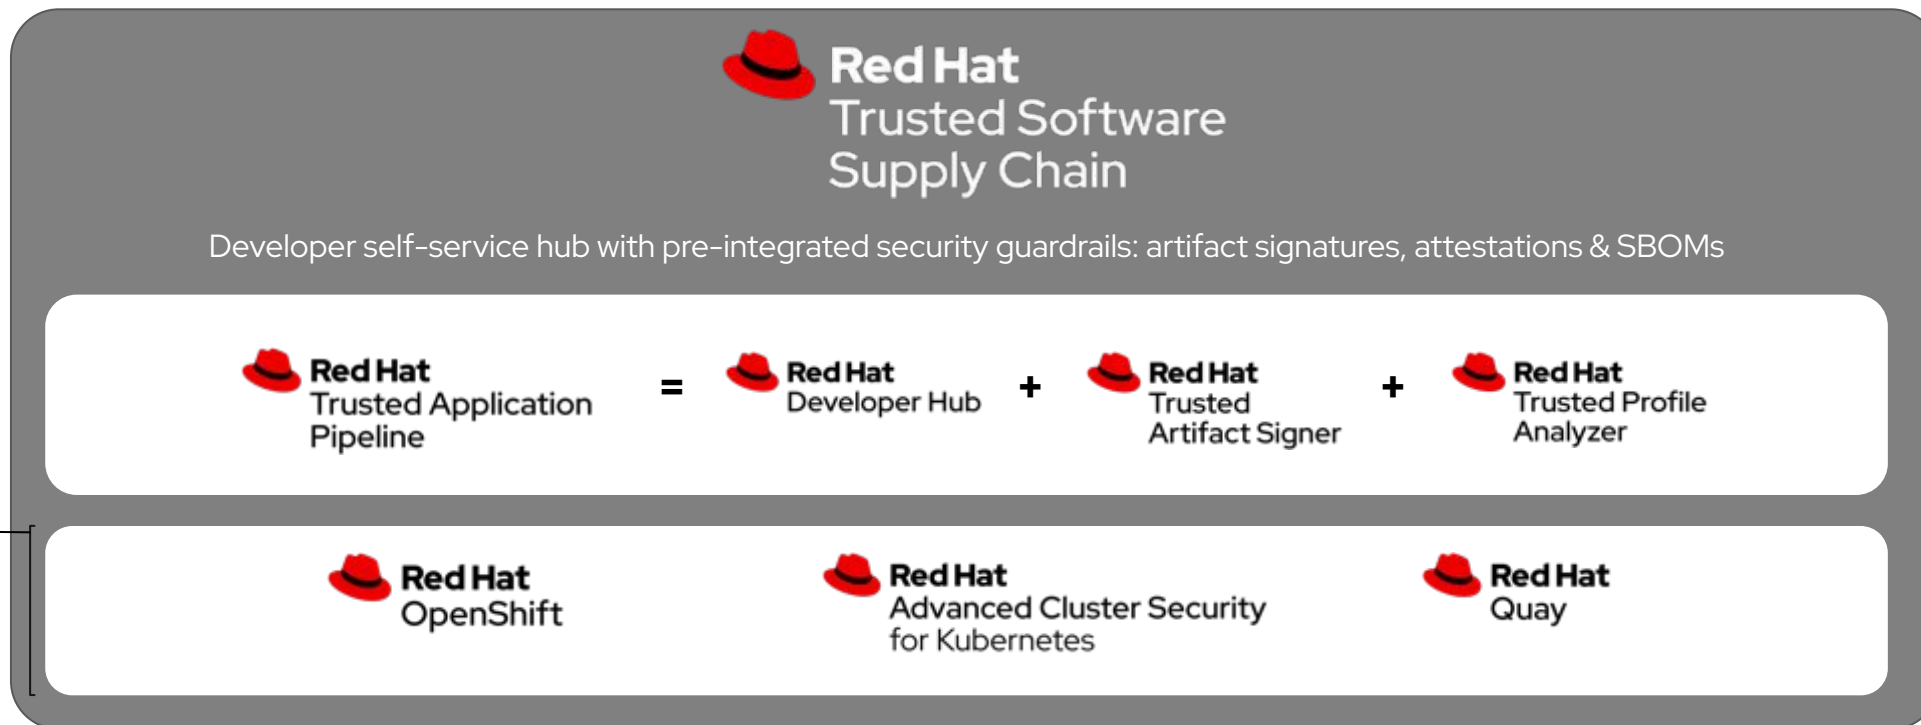


The diagram features a large grey rounded rectangle representing the TSSC hub. At the top center is the Red Hat logo (a red hat) followed by the text "Red Hat Trusted Software Supply Chain". Below this, centered text reads "Developer self-service hub with pre-integrated security guardrails: artifact signatures, attestations & SBOMs". At the bottom of the grey rectangle is a white rounded rectangle containing three product logos: "Red Hat OpenShift", "Red Hat Advanced Cluster Security for Kubernetes", and "Red Hat Quay". A line from the text "Included in" on the left points to the white box.

Included in  **Red Hat OpenShift Platform Plus**
but also available separately

Accelerate Innovation that Safeguards User Trust

Delivered with integrated security guardrails at every phase of the software development lifecycle



Included in **Red Hat OpenShift Platform Plus** but also available separately

Accelerate Innovation that Safeguards User Trust

Delivered with integrated security guardrails at every phase of the software development lifecycle



Red Hat
Trusted Software
Supply Chain

Developer self-service hub with pre-integrated security guardrails: artifact signatures, attestations & SBOMs



sigstore



TEKTON



TEKTON
CHAINS



argo



StackRox



Open Policy Agent



clair

Relevant Upstream Projects



Guac

Graph for Understanding Artifact Composition (GUAC) provides insights into artifact relationships and dependencies by aggregating SBOMs dependencies



Sigstore

A combination of technologies to handle keyless signing (**cosign**), transparency log and verify signed artifacts for integrity and provenance.



Tekton Chains

A Kubernetes Custom Resource Definition (CRD) controller to manage signing task run, task run result and OCI registry image using tools such as Sigstore cosign and securely store such signatures



Enterprise Contract

Workflow for verifying provenance by checking image signatures and attestations of OCI images

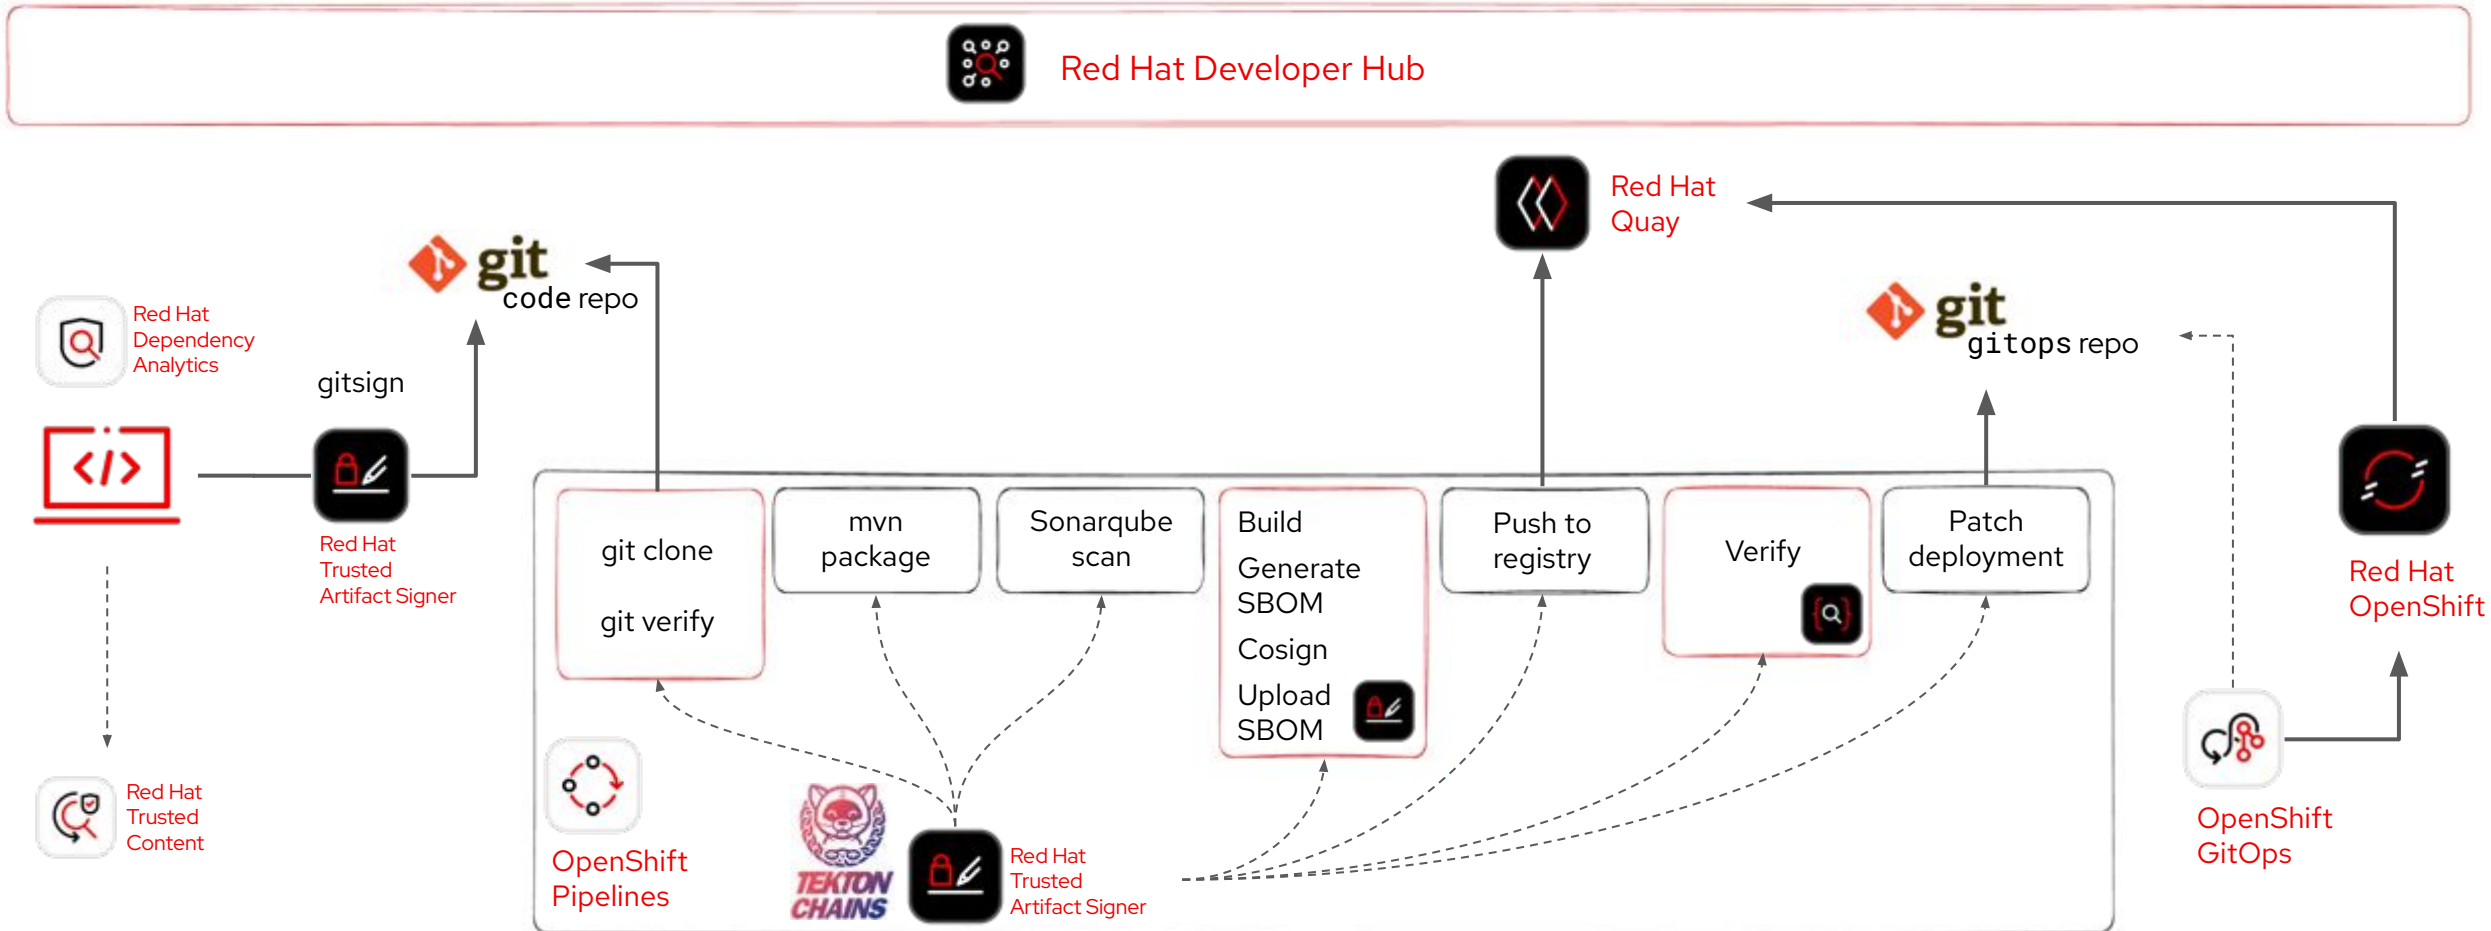
TSSC: Come rendere sviluppo applicativo e MLOps sicuri e tracciabili

Demo

Traditional application



Hands-on Scenario



Search

Search

Clear

Home

My Group

Catalog

APIs

Learning Paths

Create...

Tech Radar

Docs

Clusters

Orchestrator

Notifications

Administration

Settings

Quick Access

COMMUNITY

DEVELOPER TOOLS



Podman
Desktop

CI/CD TOOLS



ArgoCD



SonarQube




Quay.io

OPENSIFT CLUSTERS

Your Starred Entities

Click the star beside an entity name to add it to this list!

Hands-on Scenario

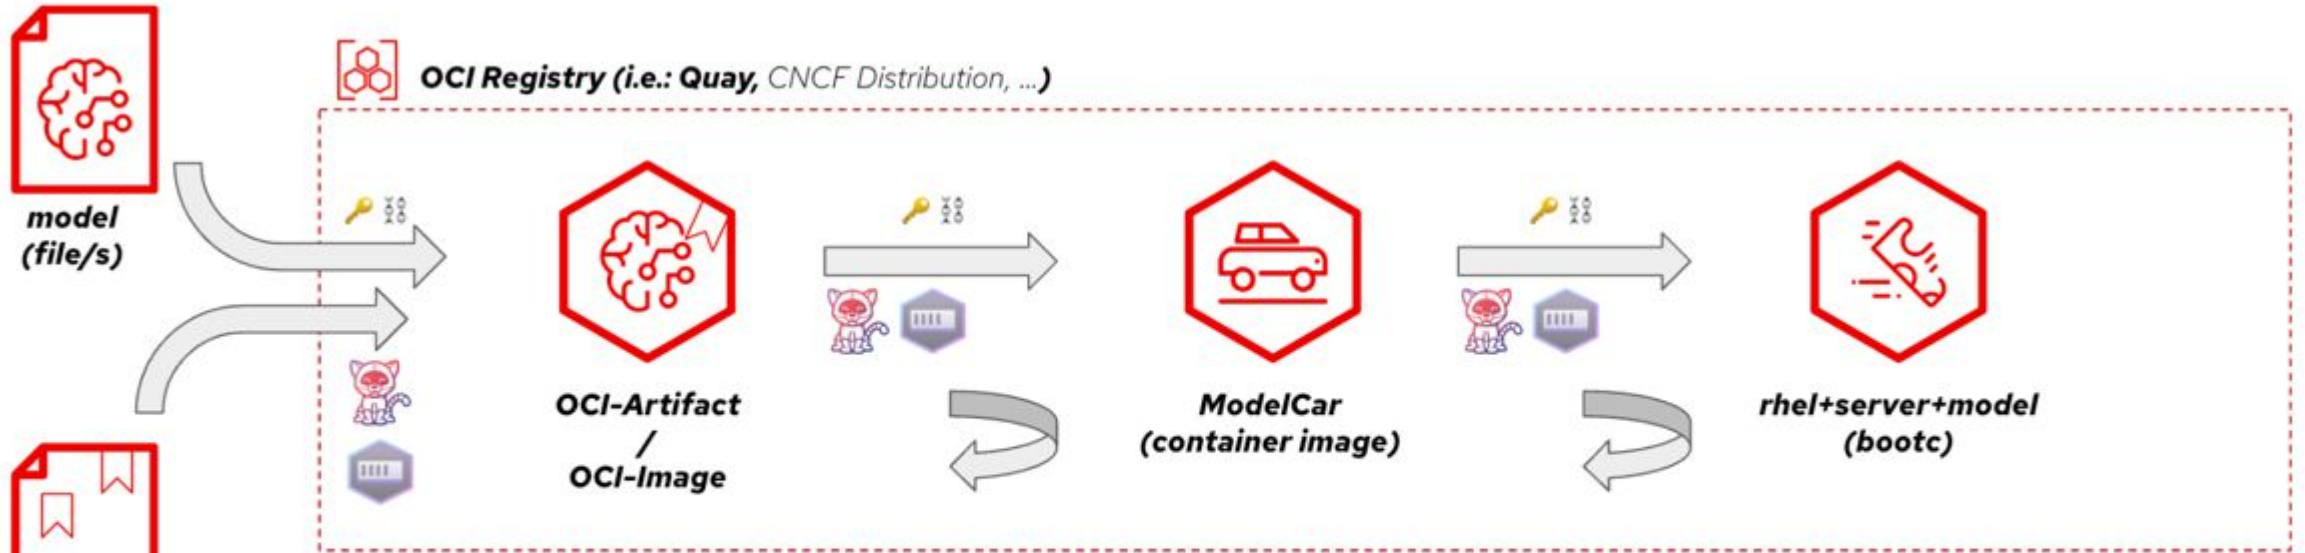


TSSC: Come rendere sviluppo applicativo e MLOps sicuri e tracciabili

Demo

MLOps

TSSC: Come rendere sviluppo applicativo e MLOps sicuri e tracciabili



- ▶ ML model & metadata distribution mechanism using existing tooling
- ▶ As it's a OCI container too, can be signed using existing tooling
- ▶ KEP-4639 would enable direct consumption in K8s

- ▶ Can be used with KServe
- ▶ Could *also* be used as initContainer in bootc (see later)
- ▶ Could *also* be used in other deployment scenarios

- ▶ Combines kernel + server + model using previous steps in 1 single container
- ▶ Could be "lift & shift-ed"
- ▶ ...but also as it composes the previous steps, could be decomposed as needed

Red Hat
Summit

Connect

Thank you



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



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



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



twitter.com/RedHat

Red Hat
Summit

Connect

Virtualizzazione Cloud Native

Approccio dichiarativo e automazione del rilascio di workload virtualizzati

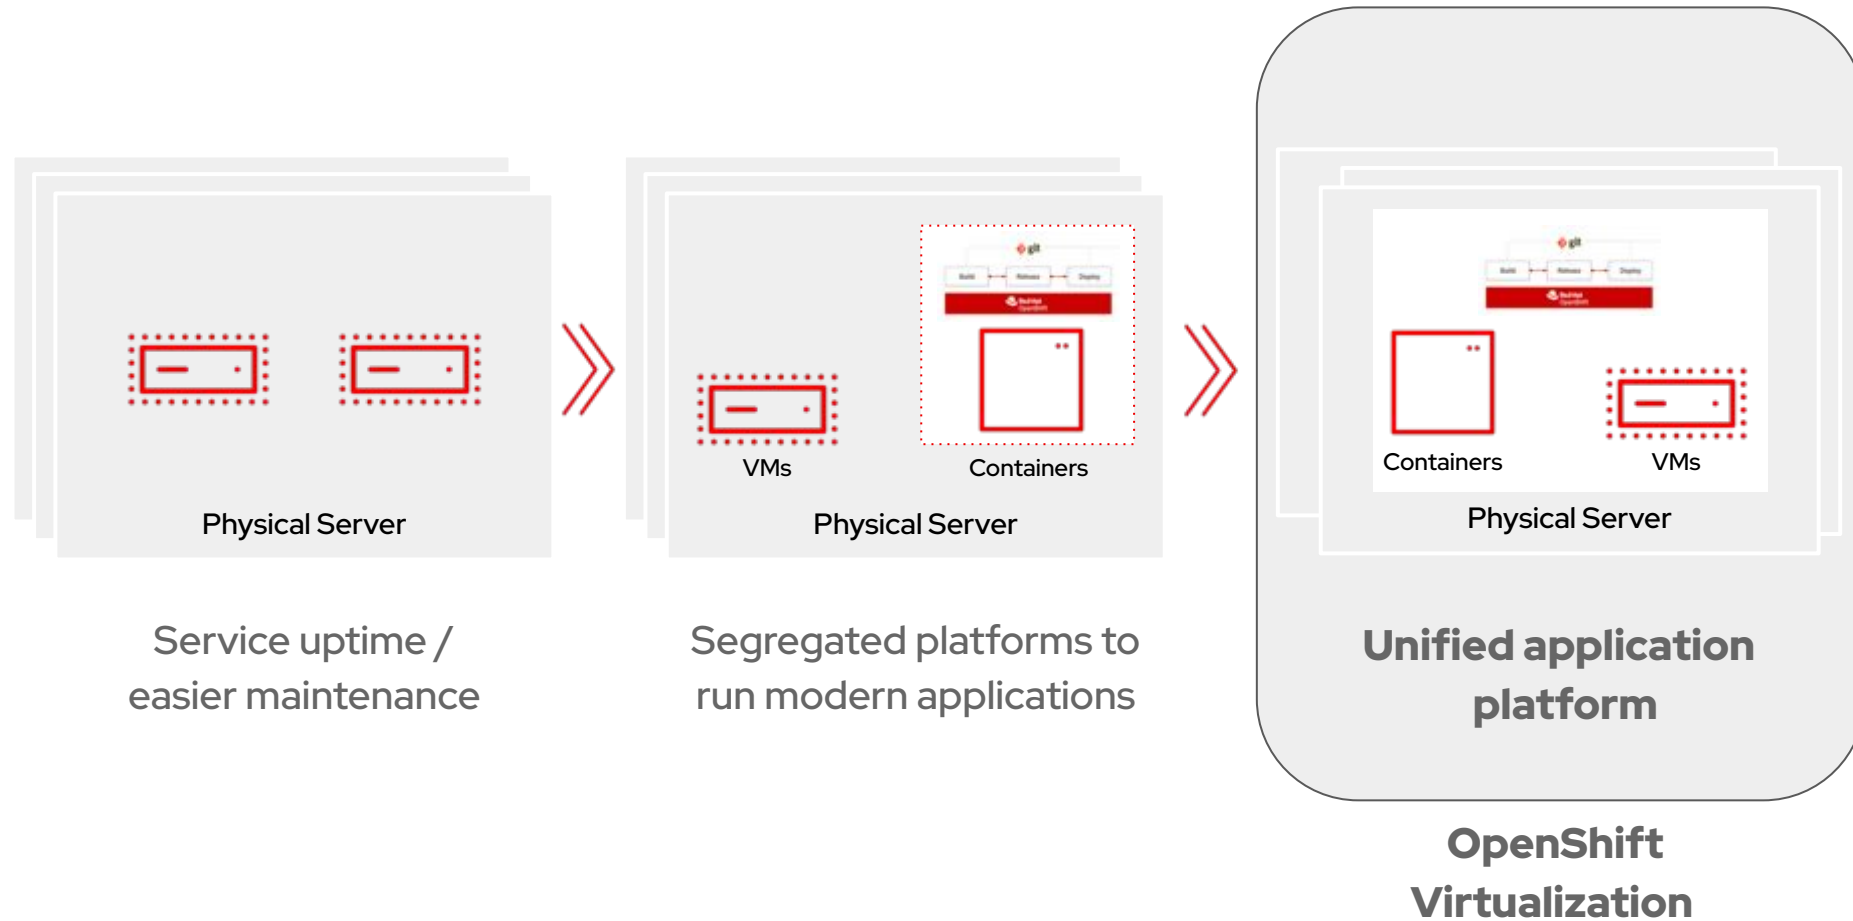
Valentino Uberti

Specialist Solution Architect

Gianni Salinetti

Senior Account Solution
Architect

Virtualization Evolution




Managing both VMs and containers



Virtual machines

VMs have been built for decades, and they will not go away overnight.



Containers

Containers solve certain use cases and will continue to rise, but some VMs will remain.



Applications

VMs and containers will be used to build applications, and some might even build on both.

Managing both VMs and containers

The screenshot shows the Red Hat OpenShift console interface. The top navigation bar includes the Red Hat OpenShift logo, the cluster name 'local-cluster', and user information 'vale'. The left sidebar contains a navigation menu with 'Virtualization' expanded to 'Catalog'. The main content area is titled 'Create new VirtualMachine' and shows a list of default templates. The templates are organized into two rows: RHEL and Windows. The RHEL row includes four templates: Red Hat Enterprise Linux 6.0+ VM, Red Hat Enterprise Linux 7 VM, Red Hat Enterprise Linux 8 VM, and Red Hat Enterprise Linux 9 VM. The Windows row includes four templates: Microsoft Windows 10 VM, Microsoft Windows 11 VM, Microsoft Windows Server 2012 R2 VM, and Microsoft Windows Server 2016 VM. The RHEL templates have a 'Source available' badge. The left sidebar also shows a filter section for 'Default templates' with a search bar and a list of operating systems and workloads.

Project: All Projects

Create new VirtualMachine

Select an option to create a VirtualMachine from.

Template catalog InstanceTypes

Template project: All projects

Default templates

Filter by keyword...

10 items

- Boot source available
- Operating system
 - CentOS
 - Fedora
 - Other
 - RHEL
 - Windows
- Workload
 - Desktop
 - High performance
 - Server

Operating System	Template Name	Boot Source	Workload	CPU	Memory	Source Available
Red Hat Enterprise Linux 6.0+ VM	rhel6-server-small	Project openshift	Workload Other	CPU 1	Memory 2 GiB	No
Red Hat Enterprise Linux 7 VM	rhel7-server-small	Project openshift	Workload Server	CPU 1	Memory 2 GiB	No
Red Hat Enterprise Linux 8 VM	rhel8-server-small	Project openshift	Workload Server	CPU 1	Memory 2 GiB	Yes
Red Hat Enterprise Linux 9 VM	rhel9-server-small	Project openshift	Workload Server	CPU 1	Memory 2 GiB	Yes
Microsoft Windows 10 VM						No
Microsoft Windows 11 VM						No
Microsoft Windows Server 2012 R2 VM						No
Microsoft Windows Server 2016 VM						No

Deeper partnerships on OpenShift Virtualization

On-Prem HW + Storage

Products for OpenShift
Virt using CSI (container
storage interface)



Backup/DR

Products for OpenShift



Networking

Products for OpenShift Virt
using CNI (container
networking interface)




Cloud Services

Current public cloud
providers offering OpenShift
virtualization



Additional Information



[Listings](#) of current partner products that are certified or completed statement of support.

Visit this [source page](#) to see the current 'in progress integrations' and to submit requests for additional partner product integrations.



* This is not an exhaustive list of ISV partners, with new partners being added all the time.

Technical Overview

Powered by KubeVirt

- ▶ Open Source, written in Go
- ▶ Initiated in 2016 by Red Hat
- ▶ Contributions by other companies
e.g (v)GPU support by Nvidia
- ▶ CNCF sandbox project since 2019
- ▶ CNCF incubating project since 2022
- ▶ Provides an API for running KVM based virtual machines in Kubernetes
- ▶ Goal: run those VMs alongside with containerized workloads



Red Hat Contributions to KubeVirt


Red Hat actively contributes to the KubeVirt project and is currently ranked at the first place as the most active company with the following stats:

- ▶ **14.32k** contributions in the last quarter
- ▶ **836** pull requests in the last quarter

	min	max	avg	current	total
Red Hat Inc.	3.22 K	11.10 K	7.16 K	3.22 K	14.32 K
International Business Machines Corporation	167.00	703.00	435.00	167.00	870.00
Gitpod GmbH	69.00	395.00	232.00	69.00	464.00
NVIDIA Corporation	90.00	94.00	92.00	94.00	184.00
Hashnode	37.00	128.00	82.50	37.00	165.00
Zyda	17.00	129.00	73.00	17.00	146.00
Google LLC	0	81.00	40.50	0	81.00
ARM	13.00	37.00	25.00	13.00	50.00
Ænix	16.00	25.00	20.50	25.00	41.00
SUSE LLC	0	25.00	12.50	0	25.00
All GeekHaven IIIT Allahabad	0	17.00	8.50	0	17.00
Independent	5.00	11.00	8.00	11.00	16.00
The Linux Foundation	0	13.00	6.50	0	13.00
CNCF	4.00	7.00	5.50	7.00	11.00
AssetCues	0	9.00	4.50	0	9.00
Mirantis Inc.	0	8.00	4.00	8.00	8.00
NetApp Inc	2.00	4.00	3.00	4.00	6.00
Kasten	0	6.00	3.00	0	6.00
devguard GmbH	0	5.00	2.50	0	5.00
Jd.Com	0	5.00	2.50	0	5.00
Cloudbase	0	3.00	1.50	0	3.00
Kuzzle	0	3.00	1.50	0	3.00


OpenShift Virtualization uses KVM

- ▶ OpenShift Virtualization uses **KVM**, the Linux kernel hypervisor and a core component of Red Hat Enterprise Linux kernel with 10+ years of production use.
- ▶ **QEMU** uses KVM to execute virtual machines
- ▶ libvirt provides a management abstraction layer
- ▶ Available on Bare Metal and AWS
- ▶ Windows Server Virtualization Validation Program (**SVVP**) certification



Containerizing KVM

Trusted, mature KVM wrapped in modern management and automation



Dedicated API

```
vm.yaml
1  apiVersion: kubevirt.io/v1alpha3
2  kind: VirtualMachine
3  metadata:
4    name: testvm
5  spec:
6    running: false
7    template:
8      metadata:
9        labels:
10         team: Tiger
11      spec:
12        domain:
13          devices:
14            disks:
15              - disk:
16                 bus: virtio
17                 name: rootfs
18            interfaces:
19              - name: default
20        resources:
21          requests:
22            memory: 1GB
```

Declarative

Like anything in Kubernetes, the KubeVirt API is declarative, and follows Kubernetes API conventions.

Domain-specific

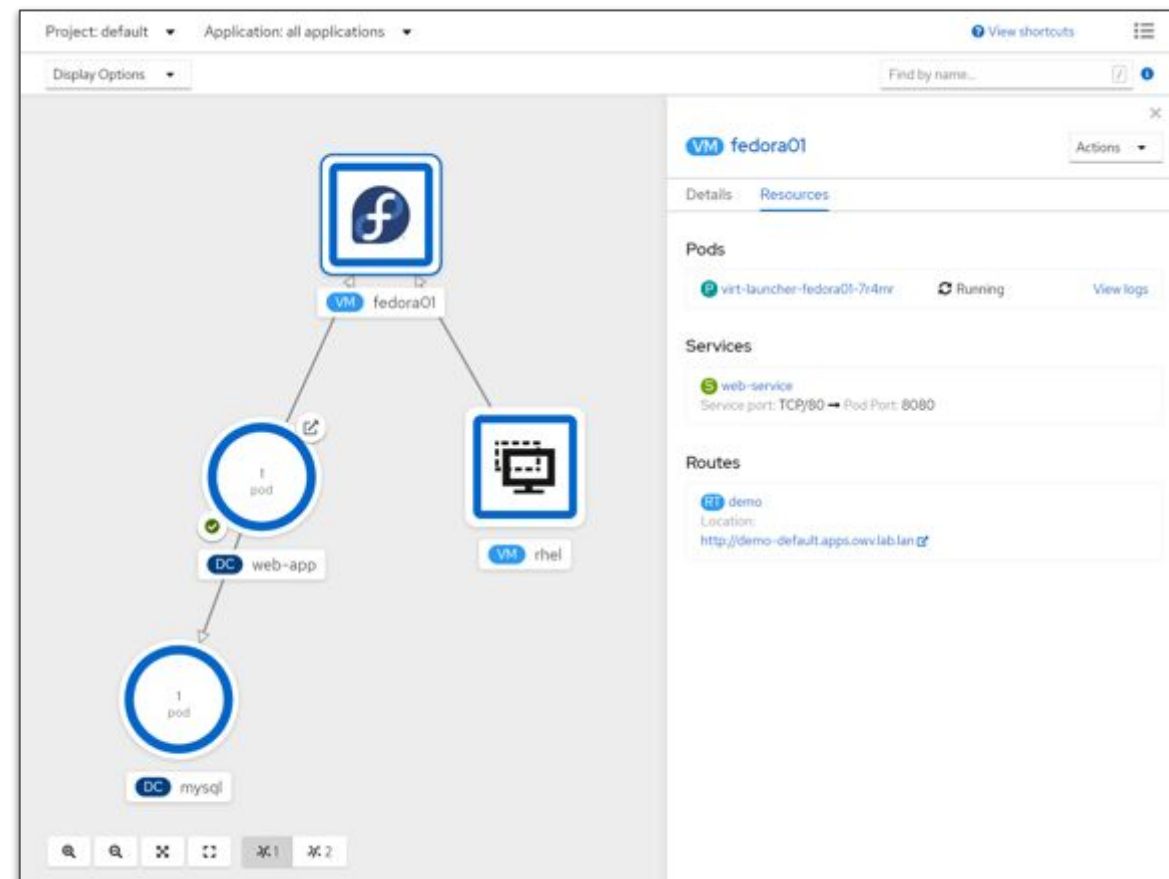
VMs are inherently differently defined than containers. Reusing the pod API is not explicit enough for all the necessary details—and due to differences.

Divide and conquer

Due to the dedicated API, it is straightforward to add virtualization-specific functionality

Using VMs and containers together

- Virtual Machines connected to pod networks are accessible using standard Kubernetes methods:
 - Service
 - Route
 - Pipelines
 - etc.
- Network policies apply to VM pods the same as application pods
- VM-to-pod, and vice-versa, communication happens over SDN or ingress depending on network connectivity




laC for Cloud Native Virtualization

Red Hat tools for GitOps & IAC



OpenShift GitOps based on [Argo CD](#) (Included in OpenShift Container Platform and OpenShift Platform Plus)



OpenShift Pipelines based on [Tekton](#) (Included in OpenShift Container Platform and OpenShift Platform Plus)



Ansible Automation Platform, a unified solution for strategic automation that combines the security, features, integrations, and flexibility needed to scale automation across domains.

Ansible Automation Platform capabilities



Applications

- ▶ DevOps
- ▶ CI/CD
- ▶ GitOps



Network

- ▶ Configuration management
- ▶ Infrastructure awareness
- ▶ Network validation



Cloud

- ▶ Orchestration
- ▶ Operationalisation
- ▶ Governance



Security

- ▶ Investigation enrichment
- ▶ Threat hunting
- ▶ Incident response



Infrastructure

- ▶ Deployment
- ▶ Provisioning
- ▶ Management



Edge

- ▶ Extend security
- ▶ Scalability
- ▶ Interoperability

Next Gen approach to VM provisioning

A process that can be optimized down to a few minutes

Virtual Machine

- ▶ CPU: 4 vCPU, 1 core
- ▶ Memory: 16GB
- ▶ Disk: 30 GB
- ▶ OS: RHEL

Additional filesystems

- ▶ data: 500GB, disk
- ▶ logs: 100GB, partition

Application platform


- ▶ JBoss 7.4 Update 11

Firewall rules

- ▶ Ingress: SSH, HTTPS
- ▶ Egress: *.redhat.com

DNS & LB

- ▶ api.service.org
- ▶ Healthcheck: HTTPS port



Demo Time



argo



Provisioning VMs with OpenShift GitOps

OVERVIEW: The GitOps way uses Git repositories as a single source of truth to deliver infrastructure as code.

During this demo ArgoCD is employed to keep the desired and the live state of clusters in sync at all times.

LEARN: How to manage the lifecycle of VMs using a purely declarative, GitOps approach.




Deploy VMs and configure external entities with Ansible Automation Platform

OVERVIEW: Ansible Automation Platform provides a complete framework to fully automate the provisioning tasks, from the creation of the virtual machine, up to software configuration.

During this demo Ansible Workflow Jobs are employed to deploy the VM and apply all the necessary configurations, including service desk management.

LEARN: How to fully manage the lifecycle of VMs using Ansible Automation Platform.



Unattended Windows VM creation with Openshift Pipelines

OVERVIEW: Red Hat OpenShift Pipelines offers an efficient solution to manage the release lifecycle of virtual machine images.

LEARN: How to create and customize custom boot sources using a dedicated Tekton pipeline.

Red Hat
Summit

Connect

Thank you



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



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



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



twitter.com/RedHat



Connect

3A per automation in Openshift: Ansible, ArgoCD, ACM



Red Hat

Yuri Francalacci

Architect
Red Hat

Amedeo Salvati

Principal Consultant
Red Hat

3 use cases and technology

- ▶ *Openshift deployment*
Ansible
- ▶ *Upgrade OpenShift EUS-to-EUS*
Ansible + ACM
- ▶ *Openshift cluster and apps configuration*
Ansible + ACM + ArgoCD + ...

Automation in OpenShift

Why?

Some benefits can be obtained with automation in OpenShift:

- ▶ No repetitive tasks
- ▶ Reduce human error
- ▶ Give back time to user
- ▶ Repeatable at scale

Red Hat OpenShift Automated Deployment and Configuration with Ansible

- ▶ *What?*
The Project Goals
- ▶ *Why?*
The Reasons
- ▶ *How?*
The Implementation & The Usage
- ▶ *Well?*
The Impact

The Project Goals

Our Customer's Needs

Provide ready-to-use OpenShift clusters to their clients, that were:

- ▶ fully configured with Auth, trusted CAs, Operators, Storage, ...
- ▶ installed on-premise on their VMware farm, *agnostic mode*
- ▶ easily scalable both horizontally and vertically
- ▶ deployable in short time

The Reasons

For Using Ansible

Why not use other tools or products, instead?

- ▶ ACM: A centralized management of all clusters was not needed, nor possible
- ▶ HCP: Final client needed to have full control of their Control Plane Nodes
- ▶ Assisted Installer: Doesn't cover all post-installation configuration aspects

The Implementation

Four Phases. Four Ansible Roles



check

- Verifies Inventory sanity
- Warns of mistakes
- *Fail early* paradigm



prepare

- Downloads *CLI* tools
- Creates *installconfig.yml*
- Adds *manifest* files
- Builds *.ign* files



deploy

- Uploads *OVF* template
- Creates & starts VMs
- Deletes *Bootstrap*
- Waits for completion



configure

- Deals with:
 - Pull secrets
 - Operators & Custom Resources
 - Nodes' Label & Taints
 - Ingress Controller
 - ODF Storage installation
 - Trusted CAs and Proxy
 - Identity Provider
 - Monitoring & Logging
 - Image Registry & Sources
 - SSL Certs for Ingress & API
 - Any additional Kubernetes object in YAML format



Usage

The Customer experience

Designs the OCP cluster and nodes sizing.

Requests IP address plan and firewall rules.

Configures Inventory as per design.

Can use previous Inventories as *blueprints*.

Runs the playbooks

`check.yml`, `prepare.yml`,
`deploy.yml`, `configure.yml`

(or just `main.yml`).

The Impact

After Automation

This level of Ansible Automation means:

- ▶ No previous OpenShift installation experience is needed
- ▶ Process is reproducible (reinstallation, node redeployment)
- ▶ Reduced *time to market* (from *days* to *a few hours*)
- ▶ Human errors brought to minimum, thanks to copy+paste
- ▶ Increase in our Customer's interest in Ansible

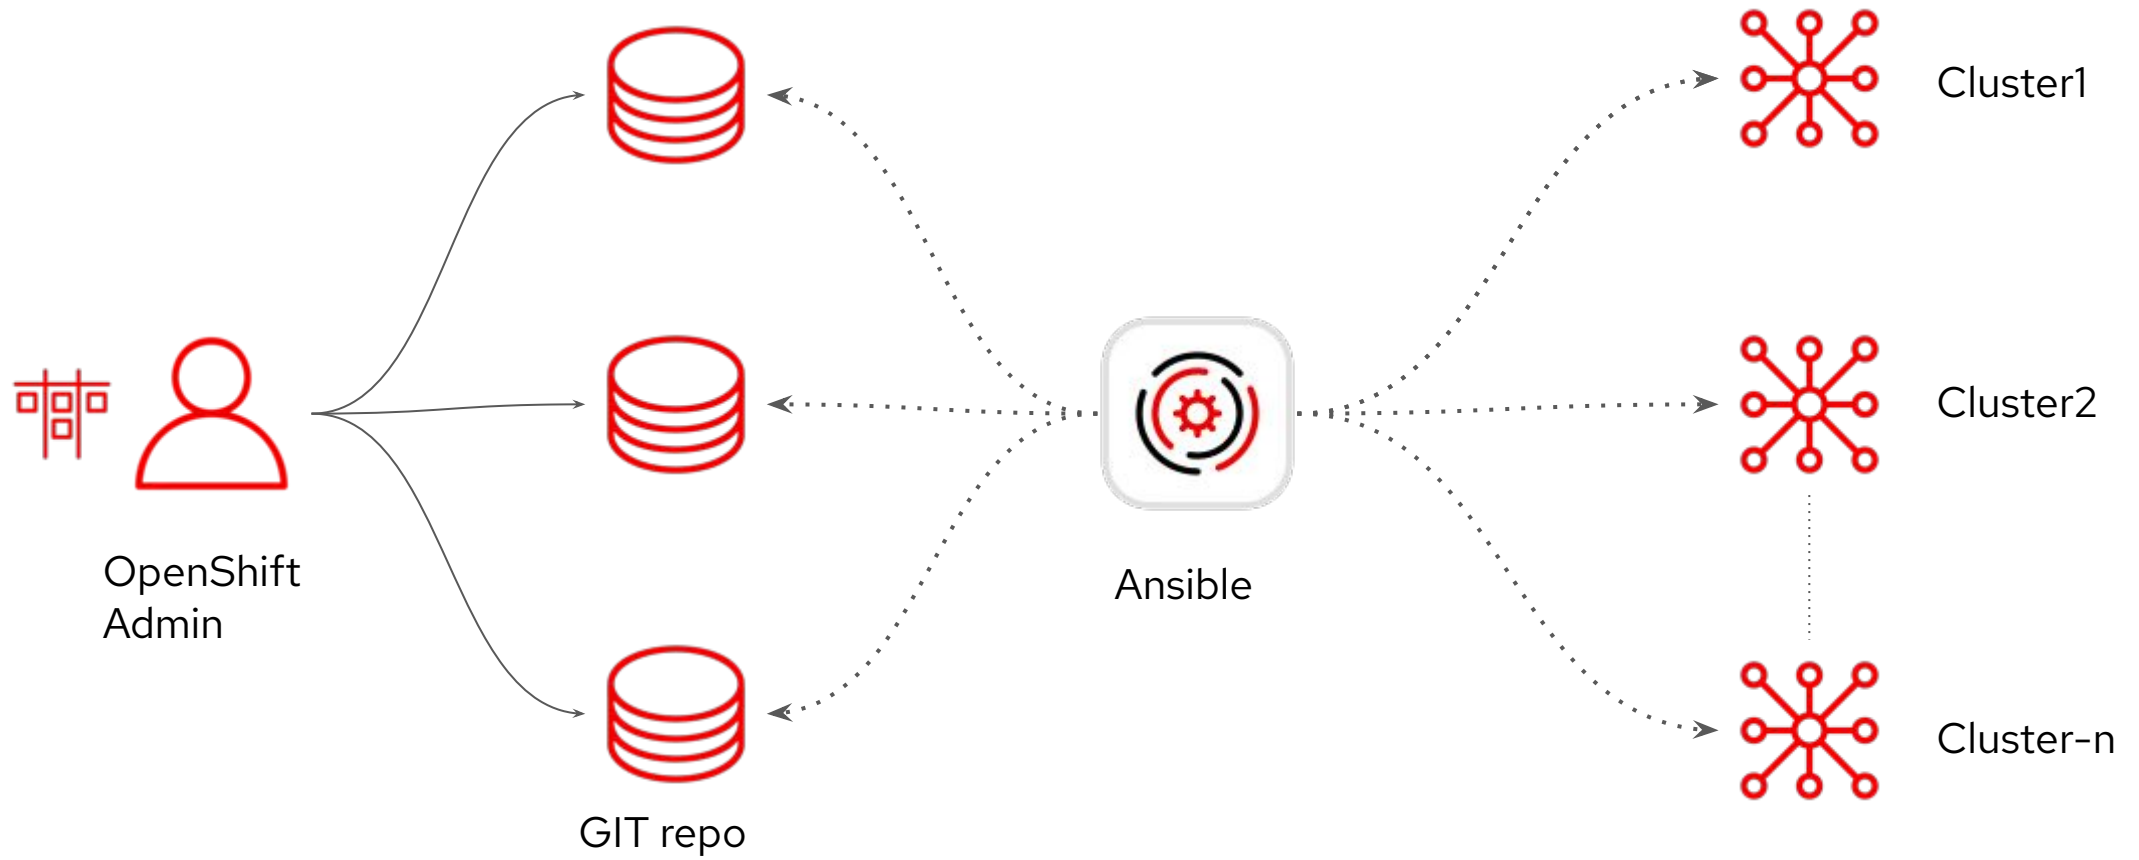
Red Hat OpenShift Automation for Cluster Upgrade EUS-to-EUS

The Project Goals

Our Customer's Needs

- ▶ Reduce extensive and time-consuming procedures
- ▶ Reduce manual activities
- ▶ Scalability of the upgrade: requirement to run it on over 25 clusters
- ▶ Make the best use of procedural and declarative automation depending on the situation

Starting Point



Create and Evolve

Create procedural and declarative automations to address specific problem

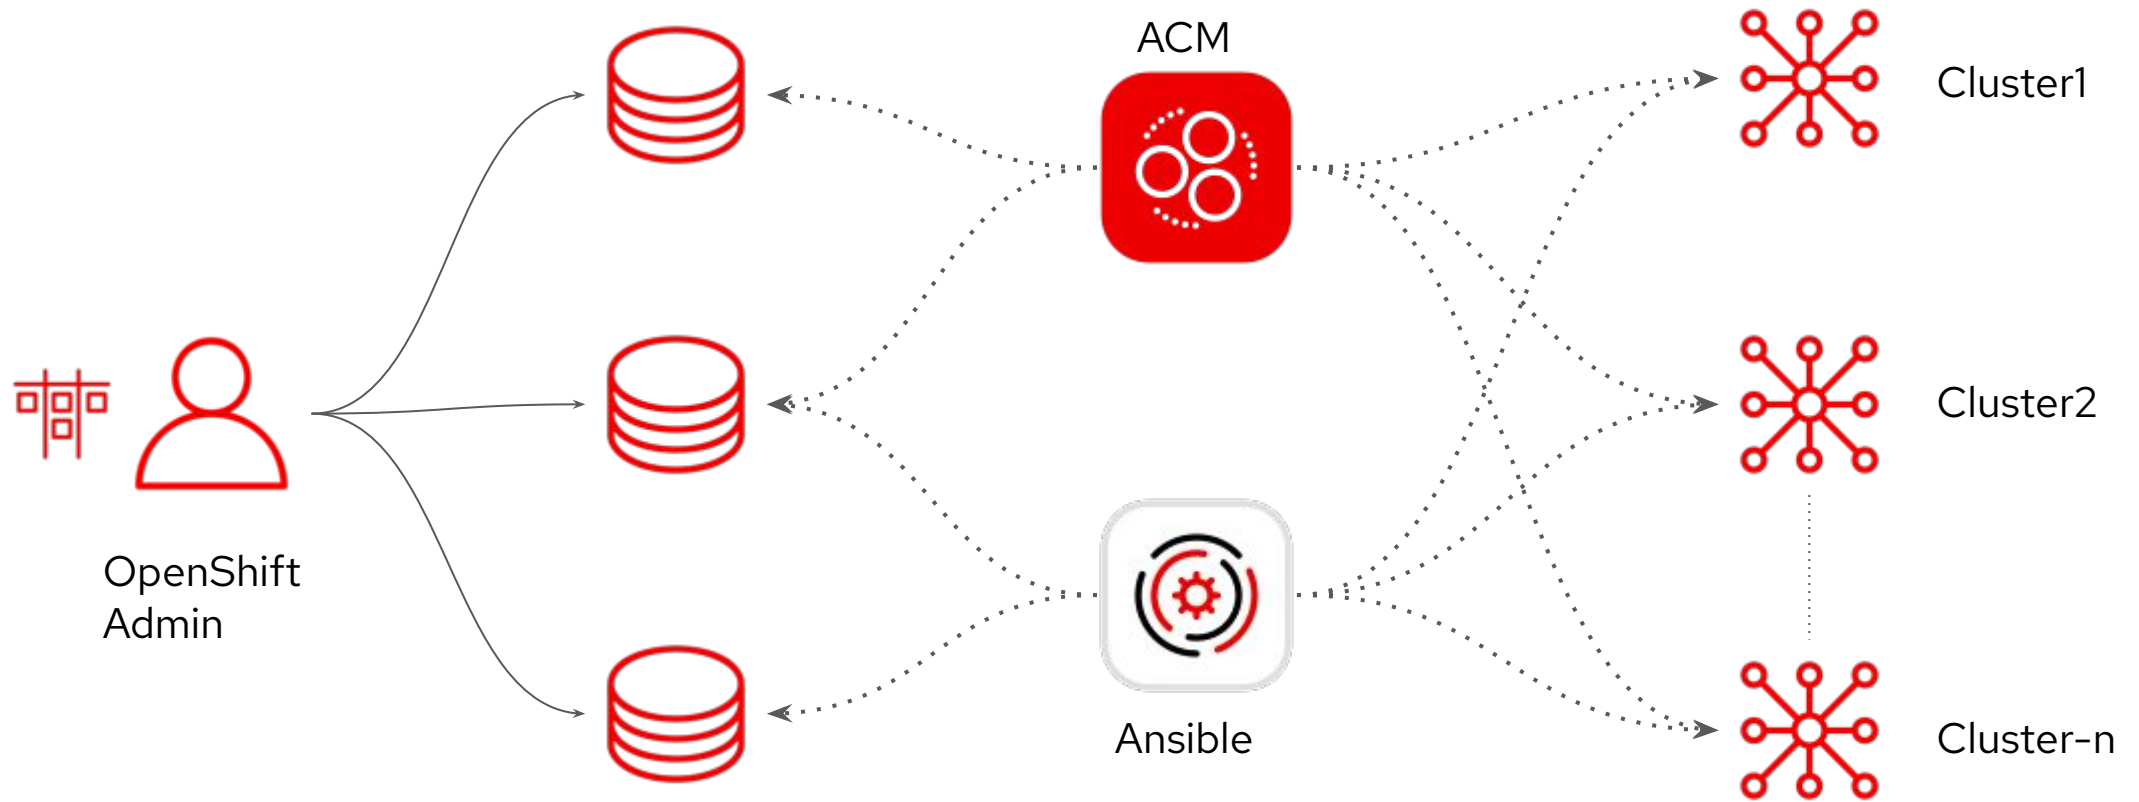


Evolve

Tune phase

Expand

Last Status



The Implementation

Actors and steps



Playbooks

- Git based file as single pane of glass for apps version
- Update current cluster configuration git repos
- Manual run



Architects

- From manual to automated run
- Define procedural and declarative automations
- Definition and control via ACM policies



OpenShift Admins

- Perform EUS-to-EUS upgrade with a click
- Activity that can be performed at any time window
- Release in Git Flow way



Developers

- Be able to adopt the new features offered by the new versions in a short time
- Double version jump with a single application restart



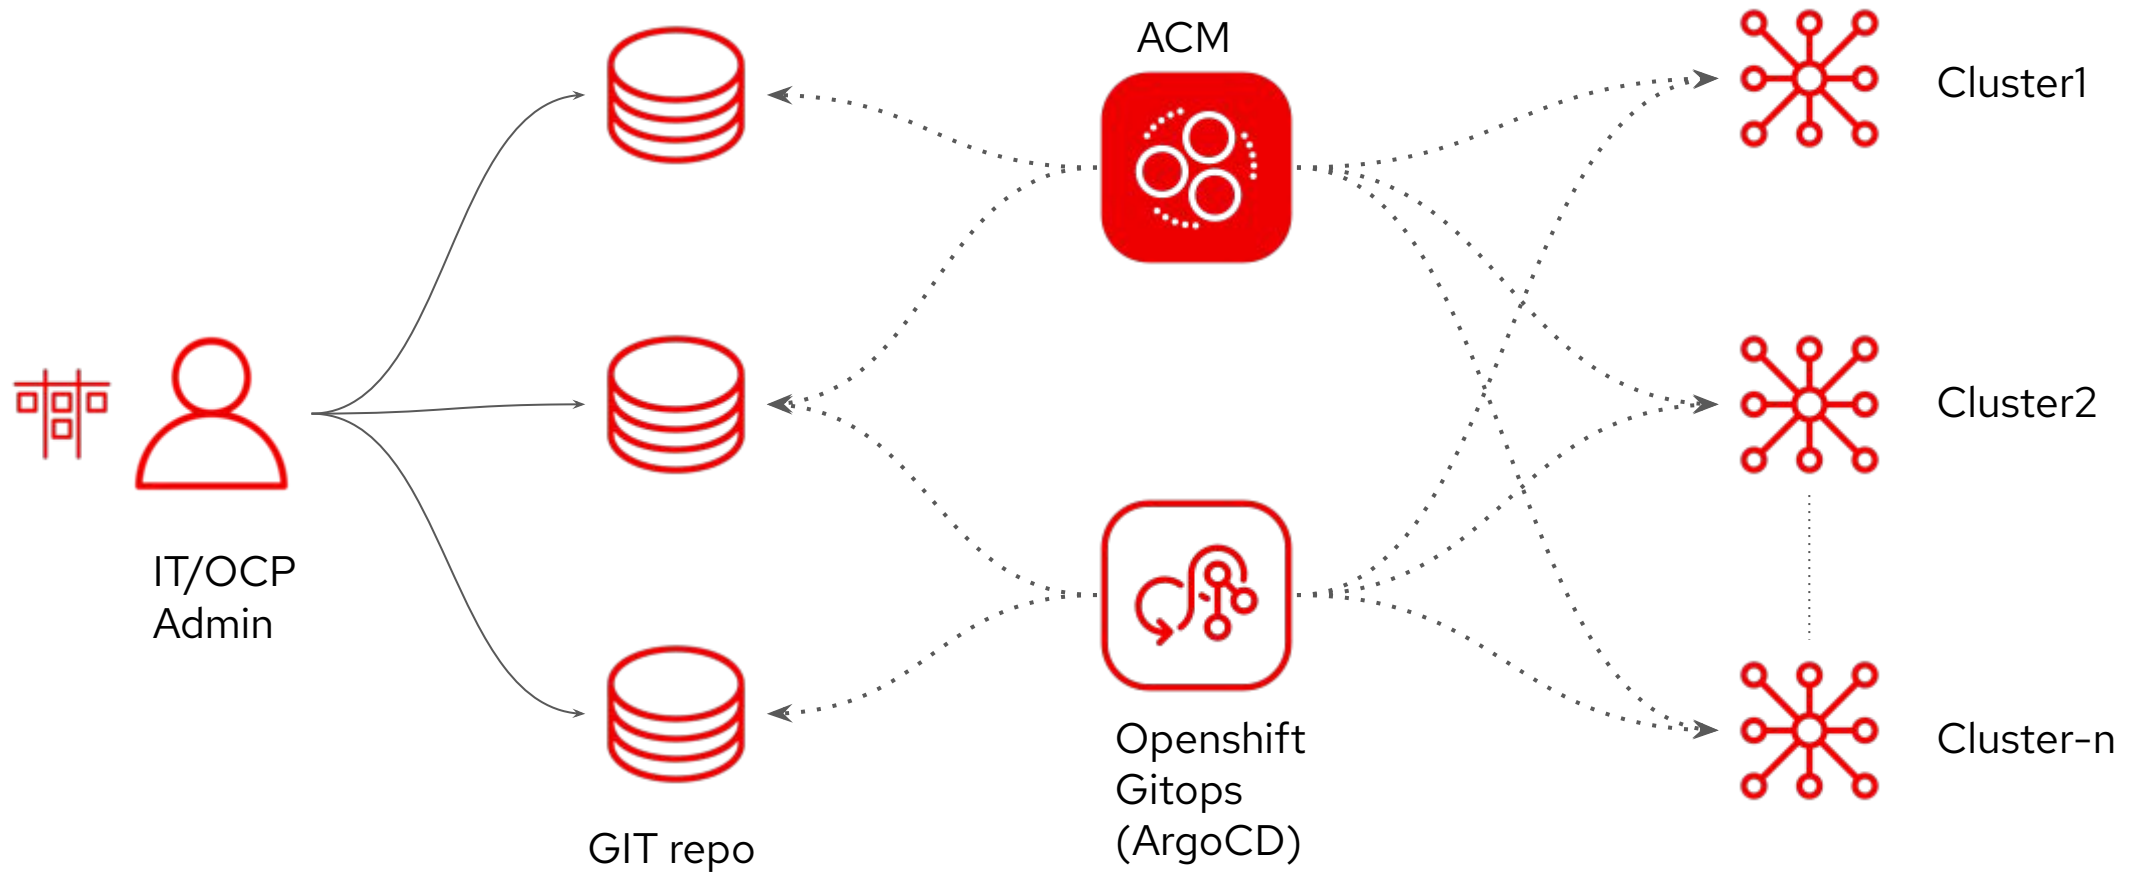
Red Hat OpenShift Automation for Cluster and Application Configuration

The Project Goals

Our Customer's Needs

- ▶ Verify application version on each environment easily
- ▶ Reduce manual activities needed to apply a "gitops" change
- ▶ Remove serialization and overload on IT team
- ▶ Reduce Mean-time-to-change / Time To Market

Starting Point



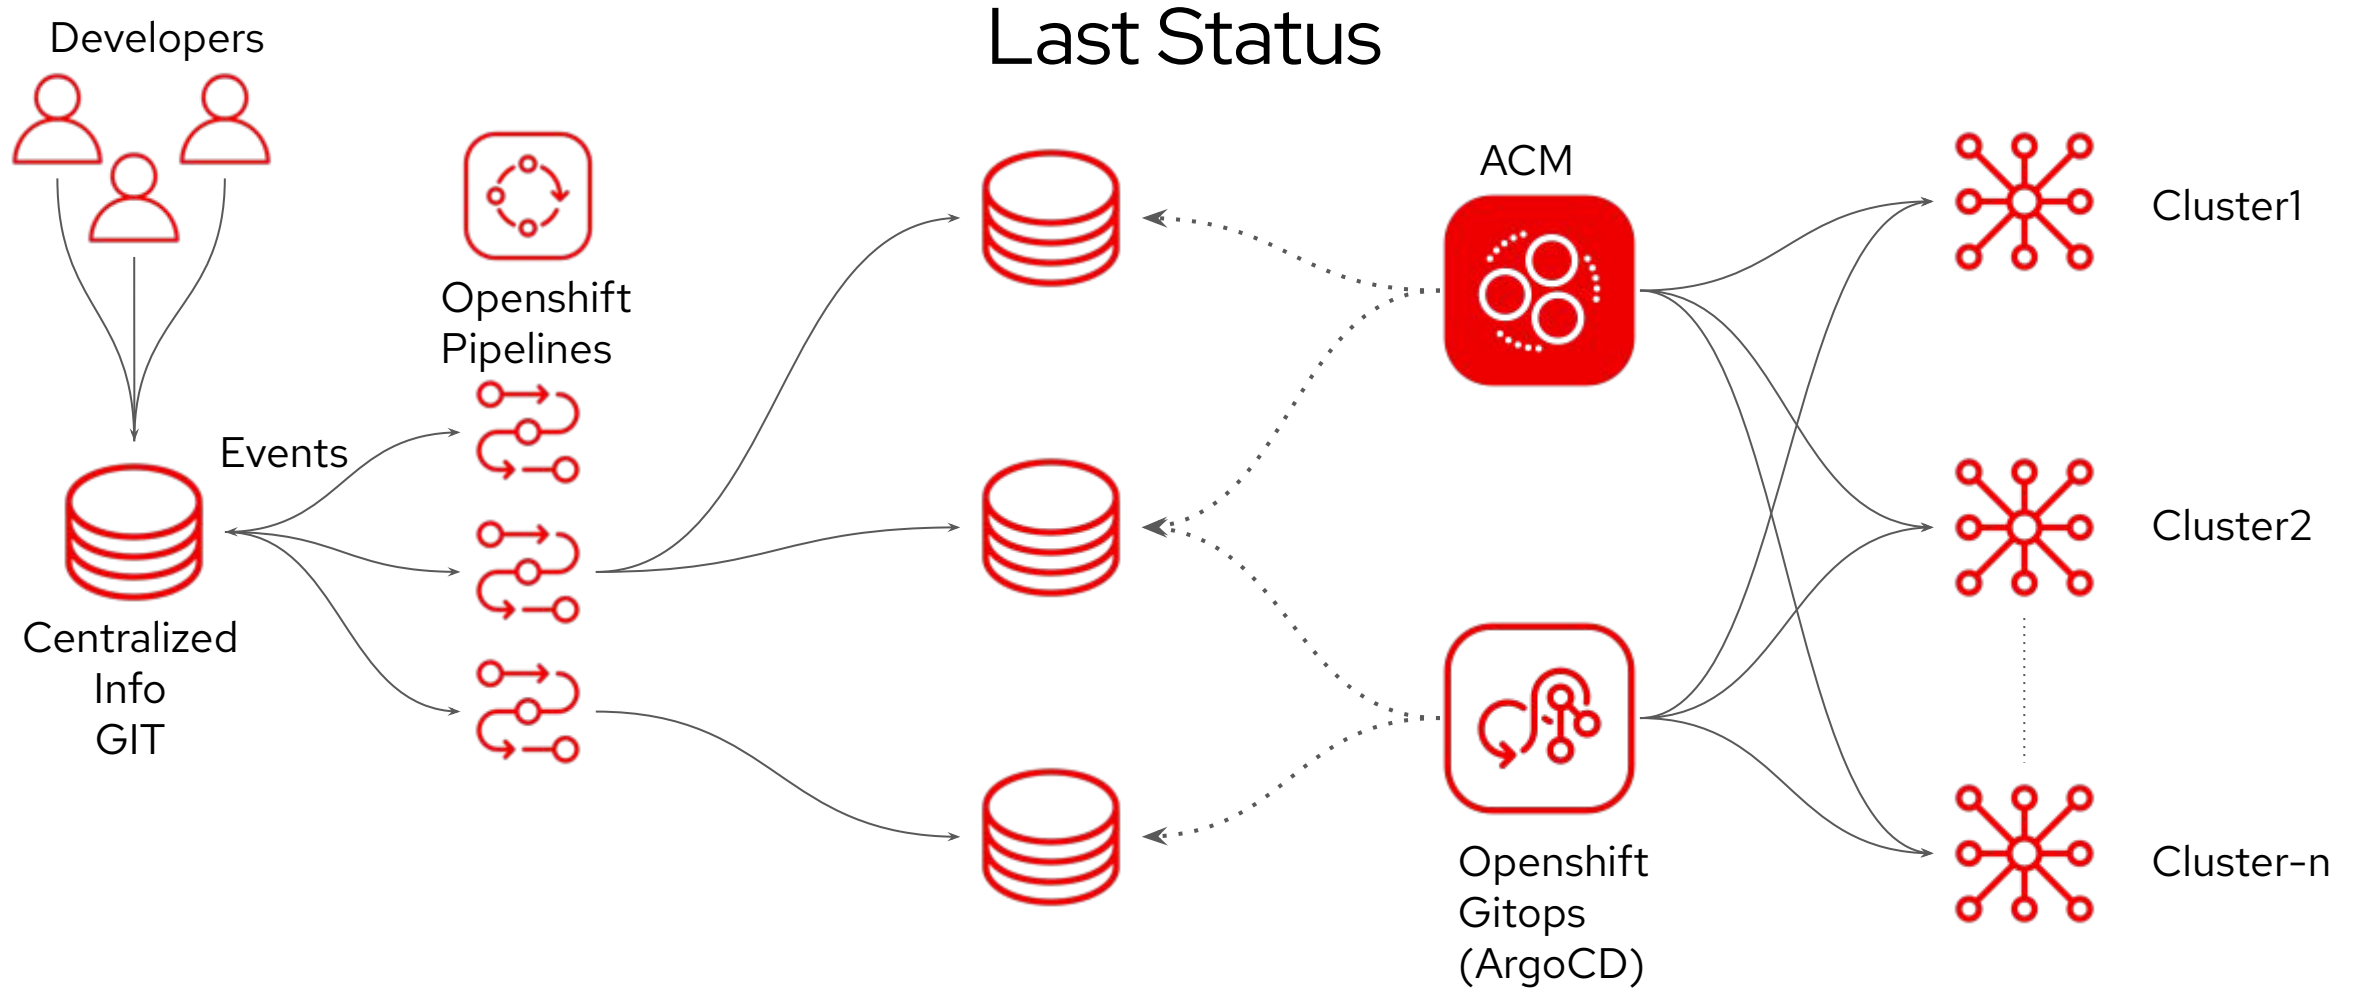
Project phases

Create multiple small automations to address specific problem

Create GIT based automation

Run automation on request

Self-service model



The Implementation

Actors and steps



Playbooks

- Git based file as single pane of glass for apps version
- Update current cluster configuration git repos
- Manual run



Pipelines

- Git based
- From manual to automated run
- Automatically invoked via git webhooks
- Definition and control via ACM policies



Developers

- Store source of truth in developers' git server
- Create pipeline to convert from dev to infra git file structure
- Release in Git Flow way



Self service

- Security and validation pipelines
- Pull-request based deployments
- Automatically deploy to test env w/ zero touch
- Deploy to prod via pull-request approval



The Impact

After Automation

- ▶ Developers can deploy apps in self-service mode
- ▶ IT team can focus on platform management
- ▶ Prevent human error with automatic validation

Continuous evolution

New requirements and ideas

Feedback loop

Integration with customer
ticketing system

Automatic document
generation

Summary

Ansible

maximum
flexibility

ACM

centralised
governance &
compliance

ArgoCD

application
deployment &
management
flexibility

Red Hat
Summit

Connect

Thank you



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



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



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



twitter.com/RedHat