




Connect


Rodolfo Falcone

Amministratore Delegato, Red Hat Italia

Red Hat Summit: Connect 2025



Red Hat Summit: Connect 2025



Posteitaliane




Andamento mercato digitale ('25E - '28E)

■ Mercato Digitale senza PNRR


■ Impatto PNRR

Variazioni in Mln € e Variazioni %




Fonte: NetConsulting cube, Giugno 2025

Andamento nei mercati verticali ('23-'24)



Fonte: NetConsulting cube, Giugno 2025

Andamento dei Digital Enabler ('24-'28E)



Ci aspetta un viaggio “adrenalinico”...

- Corsa all'AI
- Cybersecurity
- Multi-Cloud Ibrido
- IT Green
- Virtualizzazione vs Cloud
- Sovranità digitale
- Normative (DORA, ecc.)
- ...




Stiamo entrando nell'era
dell'**intelligenza
di massa**:
un miliardo di persone
ha accesso
all'Intelligenza Artificiale



Ma siamo pronti a gestirne l'impatto?


- Come **valorizzare** un miliardo di persone che usano l'AI, senza perdere il **controllo** nel caos?
- Come ricostruire la **fiducia** in un mondo dove chiunque può creare qualsiasi cosa?
- Come **proteggere** il valore dell'esperienza umana mentre **democratizziamo** l'accesso alla conoscenza?



Risultati ricerca



- **Lavorare senza AI:**
i **team superano** i singoli individui
- **Individui con AI:**
eguagliano le performance dei team senza AI,
dimostrando che l'AI può replicare efficacemente
alcuni benefici della collaborazione umana
- **Team che usano l'AI:**
lavorano più **velocemente** risparmiando tra il 12 e 16%
del tempo




Risultati ricerca




Le persone che **utilizzavano l'AI**:

- Hanno riportato **livelli significativamente più alti** di **emozioni positive** (entusiasmo e voglia di fare) rispetto a coloro che lavoravano senza AI
- Hanno riportato livelli **più bassi** di **emozioni negative** come ansia e frustrazione.



Lavorare con l'AI rende più "felici"



L'evoluzione dell'AI indica una nuova direzione

- Ci stiamo spostando dall'essere **collaboratori** che plasmano il processo a "**supplicanti**" che ricevono l'output
- Dobbiamo probabilmente diventare **intenditori di output** piuttosto che di processo





Connect

Udo Urbantschitsch


Vice President, GEO Technology Sales, Red Hat




Connect


Innovate freely, operate efficiently:
Build on a unified AI and application platform

Business is driven by major technology waves





**How do we navigate the transition
from the cloud era to the AI era?**



What have we learned from the cloud era ?



Sovereign cloud



Public cloud



Private cloud

Location

Security

Operations

Deployments

Automation

Consistency

200 Cloud
services

200 CNCF
projects

Innovation

FinOps

Platform
engineering

GitOps

DevOps

Collaboration

Open hybrid cloud

Collaboration

DevOps

Platform
engineering

GitOps

FinOps

Innovation

Workloads

Workloads

Workloads

Consistency

Modern application platform

Integrated containers and virtual machines

Consistent, secure infrastructure platform

Automation platform

Location



Physical



Virtual



Private



Sovereign



Public



Edge

Building success on Red Hat platforms

3-year ROI

540%

 **Red Hat**
Enterprise Linux^[1]

668%

 **Red Hat**
Ansible Automation
Platform^[2]

468%

 **Red Hat**
OpenShift^[3]

[1] "The Business Value of Red Hat Solutions Versus Non-Paid Open Source Alternatives," IDC, March 2023.

[2] "The Business Value of Red Hat Ansible Automation Platform," IDC, March 2024.

[3] "The Total Economic Impact" Of Red Hat OpenShift Cloud Services," Forrester Consulting, Feb. 2024.

Building success across cloud providers



Microsoft Azure

Google



ORACLE

A global imperative

Digital Sovereignty



79%

Of organisations looking operational sovereignty require data only to be accessed by EU residents from within the EU¹

.....

80%


Of EU organisations say that transparency and auditability are the top open source benefits for sovereign AI²

[1] *Source: [IDC: The Strategic Importance of Digital Sovereignty, Sept. 2025](#)

[2] *[The State of Sovereign AI: Exploring the Role of Open Source Projects and Global Collaboration in Global AI Strategy](#) (Linux Foundation Research, August 2025)

Red Hat's approach to digital sovereignty

Cloud Sovereignty Framework



Premium support, right at home



Sovereign Support

Engineers from and operating within your region who have passed a security background check




Sovereign Data

Customer support data is stored and accessed exclusively within the region and bound by its laws



Premium Skills

High touch, security-focused support engineers that are proficient in relevant compliance frameworks



Is there still more opportunity to optimize your cloud?



Connect

Andrea Leoncini

Manager Technology Sales, Red Hat Italia

Enterprise
Linux 10

Ansible
Automation
Platform

OpenShift

Red Hat
AI

Efficiency



Modernization



Simplify
deployments

Expand
automation

Cloud Native
AppDev

Enter
The new era



Connect

**Paolo
Zampieri**

RHEL Sales Specialist

**Marcello
Panci**

Ansible Sales Specialist

**Rinaldo
Bergamini**

Italy OpenShift Platform
Leader

**Natale
Vinto**

Dev Experience &
Evangelism Director

Efficiency and Security



Red Hat Enterprise
Linux 10

RHEL image mode

RHEL Lightspeed - AI assistant


RHEL Cloud-optimized

Image mode – Contain drift and accelerate delivery



With image mode for Red Hat Enterprise Linux:

- ▶ **Speed time to market** → Using DevOps and CI/CD tools and practices
- ▶ **Streamline operations** → With updates and rollbacks like your smartphone
- ▶ **Enhance security** → With compliant immutable systems



RHEL Lightspeed – Addressing the Linux skills gap



With Red Hat Enterprise Linux Lightspeed:

- ▶ **Red Hat's expertise** → With decades of Linux experience at hand
- ▶ **Proactive guidance** → Provides relevant information and guidance
- ▶ **Level up skills** → AI command line assistant for junior and senior users

```
paolo@rhel10:~  
paolo@rhel10:~$ c "aiutami a capire perchè il servizio ssh non parte correttamente"
```

Cloud optimized RHEL – Deploy workloads faster



Accelerate cloud adoption with RHEL:

- ▶ **Integrated, pre-configured and tuned** → For AWS, Azure and CGP
- ▶ **Observability and reporting** → RHEL in same cloud provider dashboards
- ▶ **Security from boot to runtime** → Secure Boot, Confidential Computing
- ▶ **Post-quantum cryptography** → Quantum-resistant algorithms ready



Efficiency



Ansible Automation
Platform

55%

reduction in lead times*

80%

reduction in error
response times*

65%

reduction in application
performance errors**

90%

improvement in server
patching time*

HashiCorp Terraform and Ansible integration

Ansible Lightspeed intelligent assistant


*MAPFRE **Dai-ichi Life

Modernizing IT operations




Standardize automation
operations

Ansible Automation Platform



Event-driven automation
& AIOps


Event-Driven Ansible +
Ansible Lightspeed




Policy enforcement of all
automation

Ansible Policy as Code

Standardize Automation Operation



AIOps & Policy Enforcement




86%
of organizations are
developing cloud-native
applications¹

76%
of organizations say
learning our architecture requires
high cognitive load, limiting
productivity¹

156%
YoY growth of
malicious packages²

An open and complete platform


from code to production - any workloads




[1] "State of Platform Engineering in the Age of AI," A Red Hat report, conducted by Illuminas, Oct. 2024.

[2] "State of the Software Supply Chain," Sonatype, Oct. 2024.

OpenShift Virtualization adoption





**How do we navigate the transition
from the cloud era to the AI era?**





TURKISH AIRLINES

"Thanks to Red Hat and OpenShift AI, our business teams—our citizen data scientists—can build their own models quickly and independently."

Emre Yavuz

Head of Data and AI,
Turkish Technology - Technology
Company of Turkish Airlines

Accelerating the developer experience—everywhere

\$10M

saved annually in fuel costs

Dynamic pricing

Fraud prevention

Customer service

Operations optimization

Logistics optimization


60

AI models in
production

40

AI models in
development





How can we prepare for AI success ?

AI is a strategic enabler across industries

Predictive AI runs businesses today; generative AI brings innovation to the enterprise

Revenue generation

- ▶ Chatbots
- ▶ Campaign and content marketing
- ▶ Developer assistants
- ▶ Guided selling
- ▶ Product innovation

Cost optimization

- ▶ Automated AI support
- ▶ Knowledgebase search & summarization
- ▶ Doc summarization
- ▶ AI-optimized logistics
- ▶ Augmented product R&D

Risk management

- ▶ Sentiment analysis
- ▶ Predict employee attrition
- ▶ Contract risk assessment
- ▶ Fraud detection
- ▶ AI-assisted security operations

AI demand is high, but there are risks



A striking number of executives believe that **AI is going to revolutionize their business.**

55%

of leading digital organizations picked **AI as a top new investment priority** ¹

58%

of CIOs will spend more than **5% of their CY25 IT budget on gen AI** ²

30%


Gen AI projects will be abandoned after proof of concept by end of 2025 ³

[1] "AI Transformation Is Driving Digital Business," IDC, June 2024.


[2] "How Will AI Investment Requirements Impact CIO Influence and IT Skills Shortage?," IDC, March 2024.

[3] "Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept By End of 2025" Gartner press release, July 2024.

Open Hybrid Cloud enabled Cost Reduction, Choice, and Flexibility




Open Hybrid Cloud + AI enables Cost Reduction, Choice, and Flexibility




Red Hat and Red Hat AI


Open, Agentic, Enterprise-Ready



The uncertainty of the AI transition



Red Hat and AI - Enabling AI Success



Red Hat Developer Lightspeed for Migration Toolkit for Applications (MTA)



Accurate AI-generated code suggestions

Uses MTA's migration data to generate accurate code solutions that are tailored for your code




Apply code fixes with a click from IDE

Automatically apply code fixes to one or all occurrences of a migration issue within an app



Refactor similar applications efficiently

Once an application is migrated, code-change suggestions for similar apps become more accurate.



Modernize apps to cloud-native at scale with intelligent app refactoring

Open Hybrid Cloud + AI enables successful Technology Era transition



Red Hat AI

**Accelerate the development and delivery
of AI solutions** across hybrid cloud environments



Red Hat Open Hybrid Cloud

Red Hat
Summit

Connect

Grazie



linkedin.com/company/red-hat



facebook.com/redhatinc



youtube.com/user/RedHatVideos



twitter.com/RedHat





Connect

Giampiero Cannavò

MED & Italy Ecosystem Director, Red Hat





Mission



Costruire, connettere e catalizzare un ecosistema aperto di partner, affinché insieme possiamo offrire ai clienti soluzioni rilevanti basate sulle tecnologie Red Hat.



Connect

Penny Philpot

Vice President Ecosystem EMEA, Red Hat



The logo for Red Hat Summit, featuring the words "Red Hat" in a small font above the word "Summit" in a larger, bold font, all contained within a red speech bubble shape.

Red Hat
Summit

Connect

Nicola Reina

Industry Technical Specialist, Intel Corporation

intel[®]



Red Hat





Red Hat

intel®

Open-source software: Intel is committed

Intel® has a long history with Linux®, actively participating in open-source development and collaboration with the Linux community, to ensure hardware is well-supported and delivers optimal performance on Linux-based systems. Intel contributes to more than 100 different open-source projects, from the Linux kernel to cloud orchestration and plugins for Kubernetes.




Intel is one of the key contributors to Linux
Linux Kernel Corporate Contributor since 2007¹



Red Hat

intel®

Evolution of AI Applications in Enterprise Use Cases






Red Hat

intel®

The AI Hierarchy: Mapping ML, DL, and GenAI with Intel

Discover how Intel® processors fuel AI workloads across inference, training, and next-generation GenAI applications.





Red Hat

intel®



intel
xeon

Intel® Xeon® 6 Processor

1.9x

higher performance per watt at a typical
40% server utilization vs. prior
generation

Designed for
Efficiency

2.5x

higher HPC performance
vs. prior generation

Significant
Performance Leaps

5.5x

higher AI Inferencing performance
vs. AMD EPYC

Unmatched
Performance

See [9G2, 9H9, 9A3] at [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel Xeon 6. Results may vary





Red Hat



intel®

Apply for a **free** Gaudi 3 Proof of Concept in **30 seconds**

Choose your GenAI or Virtualization PoC:

- ▶ Building Inference, RAG, AgenticAI, Model-as-a-Service, and other AI Use Cases with Intel Gaudi and Xeon
- ▶ Optimize fine tuning with intel Gaudi

Why work with Intel + Red Hat?:

- ▶ Benefit from access to free highly qualified experts from Red Hat and Intel and free access to the latest hardware to build your AI use case / application.

If selected, a Intel / Red Hat representative will contact you via email.



The Intel logo is centered on a dark blue background. It features a small, light blue square positioned above the first vertical stroke of the letter 'i'. The word 'intel' is rendered in a white, lowercase, sans-serif typeface.



Connect

Sergio Gianotti

Head of Italy Public Sector, AWS



Red Hat



Red Hat e AWS

Espongono le tue opportunità di innovazione

Dal 2008, Red Hat e AWS collaborano per fornire piattaforme e soluzioni moderne, coerenti e scalabili per la trasformazione delle aziende e della PA



Consistenza e
disponibilità
Cross-infrastructure



Esperienza unificata in
ambienti Ibridi



Infrastruttura scalabile
e affidabile



Supporto
Enterprise grade



AWS European Sovereign Cloud




Data sovereignty

Resiliency

Operational autonomy






Siamo entrati in una nuova Era...

Generative AI



Grazie



Red Hat
Summit

Connect


Stefano Stinchi

Direttore Divisione Enterprise, Microsoft Italia



Microsoft's European Digital Commitments

We will help build
a broad AI and cloud
ecosystem across Europe.

The background of the slide is a composite image. It features a night-time aerial photograph of a city, likely London, with the River Thames and the London Bridge visible. Overlaid on this image is a complex, glowing blue network of lines and dots, resembling a digital or data network. The lines connect various points across the frame, creating a sense of global connectivity and digital infrastructure.

Red Hat and Microsoft Partnership



Why choose Red Hat and Microsoft?

Combining solutions from two market leaders

Gain consistency from the datacenter to the cloud for maximum flexibility



Trusted global partners



Integrated solutions



Collaborative support



Engineering
collaboration



Advanced security



Azure Red Hat OpenShift

A jointly supported, turnkey application platform native to Azure.



- Comprehensive application platform
- Hybrid cloud flexibility
- Fully managed Day 2 operations
- Integrate with Azure services
- Consistent updates and patches
- Single invoice and draw down on MACC

65%

Shortened dev cycle¹

50%

Improved operational efficiency¹

¹ ["The Total Economic Impact™ of Red Hat OpenShift Cloud Services by Forrester, "Mar. 2024."](#)

Grazie



Connect

Fabio Grassini

Regional Enterprise Sales Director, Red Hat Italia



Red Hat Summit: Connect 2025



Posteitaliane





Connect

Libera il potenziale del 5G

Ivano Guardini | Head of Network Cloud, TIM

Simone Bruni | Enterprise Account Manager, Red Hat Italia



Red Hat

TIM



Infrastruttura



Rete 5G



Cloud

I valori che ci uniscono e ci caratterizzano sono la **passione** e il **coraggio**, che da sempre ci aiutano a cogliere le sfide del mercato, l'**inclusione**, perché crea valore per l'intera società, e infine l'**integrità**, per meritare e mantenere la **fiducia** dei nostri interlocutori.



Sviluppo della Rete 5G TIM

Una rete 5G cloud-native progettata per abilitare il business di domani



- ▶ Massima **Qualità**
- ▶ Più **veloce**, più **reattiva**, più **intelligente**
- ▶ Abilita applicazioni **mission-critical** e **real-time** services
- ▶ Per le imprese: più **efficienza**, più **innovazione**, più **valore**



Challenge



Time To Market



Flessibilità



Efficienza Operativa



Ottimizzazione Costi



Benefici della soluzione



Time To Market

Adozione paradigma **Agile** per il deployment

Riduzione tempi **onboarding** delle CNF



Flessibilità

Soluzione **Open Source** (riduzione del rischio di lock-in)

Ampio **ecosistema** di NEP/Partner certificati



Efficienza Operativa

Semplificazione architettura

Automazione tramite GitOps

Modalità di deployment e gestione

standardizzate (replicabilità)



Ottimizzazione Costi



Riutilizzo server e network fabric OpenStack

Efficienza nell'utilizzo dell'infrastruttura

Costruire su **best practice e competenze** acquisite (internalizzazione)



Architettura Logica CaaS



Red Hat OpenShift (OCP)

Cluster distribuiti geograficamente per ospitare le funzioni di rete 5G cloud-native

Advanced Cluster Management (ACM)


Gestione centralizzata dei cluster OpenShift

Zero Touch Provisioning per il deploy automatico e lo scale-out dei cluster


OpenShift tramite paradigma GitOps



Prossimi passi



■ Completamento 5G Dual Mode Core e abilitazione 5G Standalone



■ Ulteriore incremento del livello di automazione: cambi release e cambi di configurazione trasparenti per le CNF



■ Cambio di passo nell'efficienza operativa: AIOps



■ Paradigma FinOps per controllo ed ottimizzazione consumi



Red Hat
Summit

Connect

Grazie



[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

EXSCALATE: Cloud-Native Success Story

From molecules to candidates: come Dompé ha industrializzato l'intera piattaforma Exscalate su OpenShift AI

Anna Fava | Senior Software Engineer, Dompé

Giuseppe Romano | Partner Sales Executive, Red Hat Italia



Dompé



Red Hat








Overcoming the challenges of drug discovery

Drug discovery is complex. Progressing from biological targets to finding candidates for clinical trials is a costly, time consuming and risky process. This hinders innovation as scientists and researchers cannot easily and quickly test their hypotheses for potential drugs.

As a result, effective candidates often don't even make it past the preclinical phase.



Exscalate embraces complexity

Exscalate leverages advanced technologies such as **AI**, a **vast chemical library** together with the power of **exascale supercomputing** to effectively tackle the complexities of the drug discovery process.



Drug Discovery Process

Purpose: Identify promising compounds that interact well with the target.
Method: High-throughput screening (HTS) of chemical compound libraries, followed by biological activity analysis and refinement of chemical structures to improve efficacy and reduce toxicity.

Purpose: Confirm that lead compounds interact with the target in biological models.
Method: In vitro (cells and tissues) and in vivo (animal models) tests to verify the biological activity, specificity, and safety of the compounds.

Purpose: Evaluate the safety and efficacy of the drug in humans.

Identification of Therapeutic Target



Purpose: Identify a specific protein or gene associated with the disease to target with a drug.
Method: Biological and genetic studies to understand the disease mechanism.

Screening and Optimization



Purpose: Produce new chemical compounds that can act as lead compounds.
Method: Chemical synthesis of new compounds based on molecular docking results. Optimization of chemical and physical properties to improve interaction with the target. Testing these compounds to evaluate their efficacy and safety, followed by further structural modifications to optimize their properties.

Synthesis of Lead Compounds



Lead Validation



Preclinical Development



Purpose: Prepare the drug for human clinical trials.


Clinical Development



Red Hat

EXscalate – workflow

An end-to-end platform for docking, scoring, molecular dynamics and binding free energy calculations



Red Hat

EXscalate Pillars



Tangible Novel Druglike Chemical Space

Trillions of “easy to synthesize” compounds for virtual screening

Chemical cartridge for indexing and searching of large chemical libraries



Comprehensive Therapeutic Target Database

Is a repository of three-dimensional protein structures that have been carefully annotated with functional and topological details.

A dynamic 3D environment for virtual screening and MD simulations



Ligand Generator structure-based drug design

Drug screening software, is the central tool of the EXSCALATE platform.

LiGen is the result of a long-standing collaboration between Dompé, Cineca and the Politecnico di Milano.

Integration with IT4I Lexis platform



Polypharmacology profiling

AI Pose selection
To streamline scoring, unrealistic poses are filtered


ML Scoring Function
The scoring function evaluates poses

BFE Scoring Function
The degree of stabilization of the system




Platform Overview

Scientific Workflow



OpenShift Cluster



End-to-end AI-driven drug discovery platform

Cloud-native, scalable workflows on GPU infrastructure

Automation accelerates scientific decision-making

GPU-accelerated docking engine



AI pose selection and ML-based scoring

Physics-based Binding Free Energy analysis

Full reproducibility & scientific traceability

CI/CD Automation (Tekton)

Scientific Workflows
Commit → Build → Scan → Deploy → Run




Parallel execution of ligand expansions and docking

No manual intervention required



Trusted Software Supply Chain





Gitea repositories for versioned source code

Quay registries (public/private) with image signing

Automated scans and policy-based promotion

Data & Secrets Management




Secure DB access via OpenShift Operators

Automated secrets rotation and RBAC


Unified APIs for scientific data services

Impact & KPIs



~85%

faster
deployments



90+

workloads now
fully cloud-native



Zero

downtime
during updates

Next-Gen Roadmap

AI-powered selection for next-step candidates (Reinforcement Learning)

Federated data governance for collaborative discovery

Compliance-by-design pipelines for broader adoption



Red Hat
Summit

Connect

Grazie



EXSCALATE®





Connect

Accelerare la trasformazione cloud-native della sanità pubblica con Red Hat OpenShift

Licia Mignardi | Data Center & Cloud Director, LEPIDA

Alessandro Spigaroli | CEO Extra Red & Executive Director ENGINEERING



extrared



Red Hat



Red Hat
Summit

Connect

Grazie



linkedin.com/company/red-hat



facebook.com/redhatinc



youtube.com/user/RedHatVideos



twitter.com/RedHat





Connect

Efficienza e innovazione: la strategia vincente di **Anima Sgr** con l'**A.I.**

Massimo Airoidi | Chief Information Officer, Anima Sgr
Gaetano La Rosa | CEO EXTRAORDY - WeAreProject



Red Hat



Chi è **ANIMA** 

Punto di riferimento del Risparmio Gestito in Italia



Patrimonio complessivo in gestione di oltre 200 Miliardi di Euro



Oltre 1.000.000 di Clienti: gruppi finanziari, imprese e privati



Dalla sperimentazione al valore in produzione

Contesto

**Team dedicato
Data Scientist e di
sviluppo ML**

**Strumenti
e Tecnologie**

Dettagli

- **Provisioning**: tempi lunghi per nuovi ambienti
- Dipendenze & conflitti: **"works on my machine"**
- **Pipeline manuali** (ingestione, training, deploy)
- Tracciabilità limitata

- Il team **Data Science** investiva troppo tempo su **infrastruttura e dati**, poco su modelli.
- Complessità operativa alta, time-to-insight lento, scarsa ripetibilità e governance non by-design.

La Sfida: dal potenziale della A.I. al valore industriale


Principali difficoltà

Gestione sistemi infrastrutturali complessi.

Perdita di tempo prezioso nella parte di execution della pipeline dei task operativi che permettono al modello di andare in run

Focus sui Modelli e soluzioni

Perdita di focus da quello che sanno fare meglio:
Continuous Improvement dei modelli e soluzioni di AI



La svolta con adozione di Platform e Metodologie

Cosa abbiamo implementato

- Red Hat **OpenShift AI**.
- **Best Practice** MLOps end-to-end
- Ambienti "su misura" standardizzati
- Automazione infrastrutturale

Focus sull'AI, non sull'Infrastruttura

- Ottimizzazione modelli e Big Data
- Velocità, semplicità e meno errori.
- Supporto esperto a livello globale di **Red Hat** e gli **Architetti certificati di EXTRAORDY**.



I vantaggi e l'impatto operativo (KPI)

Cosa cambia
per i team di
Anima Sgr

- **Before → After**
 - Provisioning ambienti: settimane → < 24 ore.
 - Rilasci modelli: manuale → pipeline automatizzate con quality gate.
 - Tracciabilità: ad-hoc → audit trail di dati/modelli/versioni.
 - Focus **Data Scientist**: ops diffuso → + tempo su feature engineering e ottimizz.
- **Quick wins**
 - Riduzione errori di deploy/rollback.
 - Coerenza ambienti dev/test/prod ("niente più works-on-my-machine").
 - Velocità esperimenti: più iterazioni a parità di tempo e budget.
- **Metriche target (12 mesi)**
 - **Time-to-Market** modelli: -30 -50%.
 - Produttività **Data Scientist** (tempo su core): +25-40%.
 - Incidenti da rilascio: -50%.
 - Modelli in produzione/anno: ×2.

Prossimi obiettivi per Anima Sgr

Integrazione OpenShift

aumentare la velocità con cui trasformiamo un'idea in una soluzione concreta grazie ad **OpenShift**

- Adozione del paradigma del **Model Serving - API standard**
- Introduzione di nuovi **acceleratori** Hardware Nvidia



Time to Market

Riduzione del **Time To Market** e favorendo una **collaborazione** sempre più stretta tra i reparti, in ottica **DevOps** su tutto lo stack con **EXTRAORDY**

- Task amministrativi: **Architetture**
- Container Image: **Innovation**
- Sviluppo Modelli: **Data Scientist**

Red Hat
Summit

Connect

Grazie



linkedin.com/company/red-hat



facebook.com/redhatinc



youtube.com/user/RedHatVideos



twitter.com/RedHat



The shape of Cloud to come

Posteitaliane

The Fool(s) on the Hill

Posteitaliane

#contacts



Vincenzo Marchese

CTO



<https://www.linkedin.com/in/vmarchese/>



vincenzo2.marchese@posteitaliane.it



Luca Sabolo

Head of IT Infrastructure



<https://www.linkedin.com/in/luca-sabolo>




luca.sabolo@posteitaliane.it


Come Together
Italy's Platform Company

Posteitaliane

Everybody got something to hide except me and my company



Everybody got something to hide except me and my company



Highway to Hybrid


Any Cloud, Same Experience

Posteitaliane

Security & Continuity in the Multi-Hybrid Cloud

#hybridcloud

Containerization makes it possible to move almost any workload anywhere.




Fly me to the cloud

Security & Continuity in the Multi-Hybrid Cloud

Posteitaliane

#architecture



Two owned data centers in Italy for critical workloads and sensitive data.

Two hyperscale cloud providers (3 EU regions for Azure, 2 for AWS) that ensure scalability and access to best-of-breed technologies.


Kubernetes to abstract the infrastructure and guarantee workload resilience and portability.

Fly me to the cloud

Security & Continuity in the Multi-Hybrid Cloud

Posteitaliane

#architecture




A next-generation data platform that extracts data in real time from legacy vertical systems and exposes them as Data Products (Data Mesh)

Fly me to the cloud

Security & Continuity in the Multi-Hybrid Cloud

The new FEU (Front End Unico) as an orchestrator to enable access to sales funnels.

A security layer based on SPID and open standards to secure access from digital channels.



Posteitaliane


#architecture

Extending the infrastructure abstraction to the edge by leveraging SD-WAN.


All of this following the principles of automation and Infrastructure as Code within a DevOps culture.

Whole Lotta Cloud

The Gen3 OpenShift Journey



Posteitaliane




Sultans of Scale

Gen3 Cloud


The Poste Italiane Container Platform is fully automated and serves as the **infrastructure standard** for the development and delivery of hybrid workloads.

It provides a **single pane of glass** for centralized management of delivery, configurations, and costs—ensuring efficiency, consistency, and full operational control.



Posteitaliane

#hybridcloud
#kubernetes



Walk on the Edge Side

Gen3 Edge


Posteitaliane

#hybridcloud #edge

The **Poste Italiane Edge Container Platform** empowers business operations across the territory.


Aligned with hybrid cloud principles of **automation, control, and system decoupling** from infrastructure—enabled through **Zero Touch Provisioning (ZTP)** and **GitOps**—it ensures flexibility and resilience.

This platform simplifies software distribution and unlocks the adoption of new technologies, opening the door to **innovative business scenarios and digital transformation** across the network.



Every Breath you Trace

Observability 2.0



Posteitaliane


We will, we will watch you

Posteitaliane

#observability

Poste Italiane's Observability Platform unifies **metrics**, **traces**, and **logs** with three layers: **integration** (ingest), **collection** (scalable tiered retention), and **UI** (visualize–analyze–act).

Blending **managed services** with **custom integrations**, it delivers a governed, flexible single pane of glass across hybrid and legacy environments.



Sympathy for the standards

Built on open standards—OpenTelemetry (plus OpenMetrics and OpenTracing)—and proven tooling: OTel Collector, Prometheus, Jaeger, Elasticsearch/OpenSearch, Kibana, and Grafana.

For scale, we rely on AWS managed services: Amazon Managed Service for Prometheus, Amazon Managed Grafana, Amazon OpenSearch Service, and CloudWatch—fully managed observability for large-scale workloads.



The sound of metrics


Posteitaliane

#observability

80 Mln of Active Timeseries (60 sec samples)

90.000 span/sec

95 TB/month of logs, over 76B record/month



Source: Internal Data – October 2025

I still haven't found the metric I'm looking for


Ok, but how do we monitor:

Assets



PK-Watch


A set of scrapers to collect info on assets (hosts, hardware, services, ...) and sends the info to the company's asset management tools



Podnoise exporter

A golang agent that identifies "noisy" pods, i.e. pods that produce a large number of log lines, the number of exposed metrics, CPU and Memory unused allocation, IOPS and unmounted persistent volumes


Generative AI Pipelines




- Human Feedback
- Context Precision
- Context Recall
- Response Relevancy
- Faithfulness
- ...

Bridge over troubled Virtual Machines

Gen 3 Virt



Posteitaliane



Exit Music (For a VM)


Posteitaliane

#hybridcloud
#virtualization

To progressively phase out VMware, we have extended the Gen3 approach to virtual machine orchestration by leveraging **OpenShift Virtualization**.

Gen3 Virt brings this vision to completion, enabling seamless workload delivery on Kubernetes while inheriting all the core strengths of Gen3—automation, compliance, security, and a unified management experience (single pane of glass).





Kind of Blue(print)

The Poste Italiane Kubernetes blueprint (based on the Red Hat distribution) is built on:

- Zero Touch Provisioning (ZTP) installations
- Bare metal or cloud IaaS nodes (single, three, or multiple)
- Infrastructure and application-level **GitOps**
- **Extended cluster components** that simplify integration into the Poste Italiane operational ecosystem



This blueprint ensures a robust, standardized, and easily scalable foundation for Kubernetes adoption across the organization.

Default Layout

- 2 x INTEL Xeon Gold 5420+ 2Ghz, 28C
- 1024GB di Ram
- 2 x SSD 1.6TB SAS R1
- 2 x Scheda LAN DualPort 10 GbE Base-T su bus PCI-E
- 1 x DualPort 10 GbE Base-T integrata
- 1 x Scheda LAN DualPort 1 GbE Base-T embedded
- 2 x HBA FC 16/32 DualPort con transceivers e supporto FC-SP2
- 2 x Power supply

Posteitaliane

#hybridcloud
#virtualization



The Times They Are AI-Changin'

Should we talk about that?

Paranoid Androids

Getting better

Select AI technical performance benchmarks vs. human performance

Source: AI Index, 2025 | Chart: 2025 AI Index report





Figure 2.1.33'

...and smaller

Smallest AI models scoring above 60% on MMLU, 2022–24

Source: Abdin et al., 2024 | Chart: 2025 AI Index report



Paranoid Androids


Posteitaliane

#ai

...and cheaper

Inference price across select benchmarks, 2022–24


Source: Epoch AI, 2025; Artificial Analysis, 2025 | Chart: 2025 AI Index report



...and more profitable (?)

Cost decrease and revenue increase from analytical AI use by function, 2024

Source: McKinsey & Company Survey, 2024 | Chart: 2025 AI Index report



Source: Artificial Intelligence Index Report 2025 – Stanford University


With a Little Help from My Agents

Posteitaliane

#ai #strategy

Our AI Strategy


Flexibility



Awareness *(training, webinars,...)*

Governance *(ethical committee)*


Knowledge




Here, There and Everywhere (Inference)

The multi hybrid-cloud model allows us to deliver any type of workload on-premises, on hyperscalers, and at the edge.

The possibility of having nodes with **GPUs** therefore enables us to bring inference wherever it is needed for service delivery.






Stairway to AI

To deliver LLM (Large Language Model) workloads on-premises, in the case of sensitive workloads, we adopt the same **containerisation** strategy that we have implemented for traditional workloads.

Gen3 AI, also based on the Gen3 paradigm, is a solution powered by **NVIDIA** and **vLLM**.

The default layout of Gen3 AI nodes is based on the Nvidia HGX B200 architecture.



Posteitaliane

#hybridcloud #ai




Eight Days a Week

24×7 Ops in Figures

Ok Computer: *the numbers*

Posteitaliane

#results



2200 projects (namespace) in dev stage
2700 systems (namespace) in production

~64306 repositories

297 projects in Continuous Testing

678 unit test, 312 e2e test



OPENSIFT

50 clusters/~1000 nodes
~11k containers/~68k POD (500 VM)
~8000 cores in production (21k total)

32502 build pipelines

24667 deploy pipelines in production



5097 chart in production (111660 releases)
4927 charts in dev (42651 releases)

~550 full IaC projects

~130 Terraform modules

~ 70 Ansible modules

24 Packer OS images



Source: Internal Data – October 2025

Greatest Hits

Field-Proven Gen3 Projects



Hello Storyteller

Posteitaliane

#ai #softwareengineering

Storyteller leverages **generative AI agents** to guide users through interactive interviews, turning project ideas into structured **requirements**. AI automatically **extracts** and **organizes** these requirements into an editable list and a **PRD** document.

Project Idea Capture

Guided AI Interviews

Automated Requirements Extraction

Completeness and suggestions
Dashboard



We Can Work It Here: POLIS

A FEU dedicated on-prem deployment **hardened** for POLIS

Multi Factor Authentication
protection with **FIDO2**
cryptographic tokens

~2100 postal offices enabled
~110.000 passport requests accepted
~42 *Questure* served

UFFICIO: 65115 PDL 02 DATA CONTABILE: 09/03/2023 CASSA: 4 N.OPERAZIONE: 1 TESTO711

POLIS > HOMEPAGE ESCI

Identifica il cittadino

Richiedi un documento di identità valido e la tessera sanitaria o il codice fiscale

Codice fiscale

Inserisci

PROSEGUI

P O L I S
STESSI DIRITTI, STESSI SERVIZI,
STESSA OPPORTUNITÀ.

Stato Periferico: 10/10/2023 13:54

A team supreme

Posteitaliane


#team

Heading to a single platform for containers and VMs



Centralized governance, with strong technical and operational ownership by Poste italiane

OCP Gen2 → Gen3



Posteitaliane

Solution Design, Planning & Monitoring,
BluePrint Definition, Run Management


SOURCESENSE

Tech. Support

OTHER SI

Project Governance

VMWARE → GENVIRT



Posteitaliane

Solution Design, Planning & Monitoring,
BluePrint Definition, Run Management,
Network delivery

SOURCESENSE


Delivery Tech. Support

Red Hat

Design & delivery Tech. Support

OTHER SI

Planning, Governance, Delivery Assurance



Good riddance (time of gen2 life)

Posteitaliane

#hybridcloud #virtualization




1.800 Namespace



700 Cloud



1.100 On Prem



48 Clusters

14 Gen2

34 Gen3

Migration Progress

Migrated Namespaces: 884
49% complete




Includes modernization:

- Global routing improvements
- Deprecated API remediation
- Security & route hardening
- Container registry modernization



The Gen2 → Gen3 migration is halfway through, with 884 namespace already upgraded and modernized. The program is on track for full completion by June 2026, delivering a more secure, compliant, and cloud-native OCP platform



Symphony in three movements



Applications in the **Finance** domain (or supporting it) fully compliant with BankIT regulations





Applications in the **Insurance** domain (or supporting it), subject to IVASS regulation

● ● ● Everything **else** (Postal Offices, Logistic, Corporate, ...)



The Long and Winding Roadmap






Giant Steps (in virtualization)

Posteitaliane

#hybridcloud #virtualization



25 Gen 3 Virt Farms



650 nodes



~21.000 Virtual Machines

Source: Internal Data – October 2025

Bonus Tracks

Posteitaliane

Side A

Garantire la Conformità della PCI DSS 4.0 in Ambienti OpenShift Multi-Cluster



Riccardo Tagliaferri



Amedeo Salvati



Francesco Cuccagna

Side B

Percorso di modernizzazione verso una virtualizzazione cloud native basata su Openshift Virtualization



Stefano Tommaso Lombardi



Paolo Conti



Matteo Santucci



Valentino Uberti



Posteitaliane

#security #virtualization

What a wonderful world of
Open Source!

Thanks for tuning in

Posteitaliane

CHI SIAMO

Poste Italiane costituisce la più grande **platform company omnicanale** di servizi in Italia, attiva nei settori **logistica, corrispondenza e pacchi**, nei **servizi finanziari e assicurativi**, nei **sistemi di pagamento**, nella **telefonia** e nel **mercato dell'energia**.

Con oltre 160 anni di storia, il **Gruppo** rappresenta una realtà unica per dimensioni, riconoscibilità, capillarità e fiducia da parte della clientela e fornisce un importante contributo al tessuto economico, sociale e produttivo del **Paese**.



Il nostro **Purpose**

Crescere responsabilmente grazie al decisivo contributo delle proprie persone per il successo sostenibile, l'innovazione, la digitalizzazione e la coesione sociale del Paese.

La nostra **Strategia**

Il Piano Strategico 2024-2028 «The Connecting Platform» conferma Poste come la più grande piattaforma Phygital del Paese

LAVORARE CON NOI È TUTTA UN'ALTRA STORIA.

Abbiamo scritto la storia del Paese,
con noi puoi scrivere la tua.
Scopri le nostre opportunità lavorative
sul sito posteitaliane.it




PERCHE' SCEGLIERE POSTE

Lavorare in Poste Italiane significa far parte di un'azienda **dinamica** ed **innovativa**, capace di valorizzare la propria tradizione ma sempre con uno **sguardo al futuro**.

Attraverso le **opportunità di inserimento** nelle più stimolanti Funzioni aziendali, Poste Italiane costituisce il luogo ideale in cui ciascuno può **scrivere la propria storia professionale** contribuendo a scrivere la **storia del Paese**.

I NOSTRI RICONOSCIMENTI



Promuoviamo una
Cultura Inclusiva



Valorizziamo
le Persone



Welfare & Benessere



Opportunità di
lavoro in tutta Italia

L'ESPERIENZA DI SELEZIONE IN POSTE ITALIANE

Posteitaliane

LA CANDIDATURA

Posteitaliane

AZIENDA INVESTITORI GOVERNANCE MEDIA CONTRATTI TRASPARENTI SOSTENIBILITÀ CARRIERE

Opportunità Professionali

In Poste Italiane valorizziamo l'esperienza professionale delle nostre persone rendendole protagoniste del loro percorso di carriera. Mettendo a frutto le proprie competenze e abilità distintive, ciascuno può partecipare alla trasformazione della nostra organizzazione e traghettare il purpose dell'azienda: crescere responsabilmente attraverso il contributo decisivo delle proprie persone, per il successo sostenibile, l'innovazione, la digitalizzazione e la coesione sociale del paese.

Metti la tua esperienza al lavoro!

Scegli la funzione che fa per te e usa la tua esperienza per guidare il successo del Gruppo mentre sviluppi le competenze del futuro.



Posizioni aperte

Scopri la posizione più adatta a te sul nostro portale di recruiting e invia la tua candidatura.



Candidatura spontanea

Ti interessa lavorare con noi? Inviaci il tuo curriculum con una candidatura spontanea. Ti ricontatteremo nel caso in cui il tuo profilo dovesse risultare in linea con le nostre esigenze.

Vai alla sezione **CARRIERE**
del sito www.posteitaliane.it

PROCESSO DI SELEZIONE



Gli iter sono costruiti per farti vivere la migliore esperienza e prevedono incontri con **recruiter e manager** e percorsi **innovativi e digitali**, differenziati per seniority, che consentono di mettere in campo le proprie **attitudini e talenti** attraverso **dinamiche immersive** e coinvolgenti come ad es. challenge, tech talks, virtual recruiting day ecc."



Il nostro processo di inserimento prevede l'affiancamento a un **tutor** per creare un'esperienza personalizzata e si avvale di una **piattaforma digitale** attraverso la quale manager e nuova risorsa condividono i piani di azione, i documenti e pianificano incontri con focus su accoglienza, affiancamento, formazione e feedback.

RESTIAMO IN CONTATTO, SEGUICI SUI NOSTRI CANALI SOCIAL

Posteitaliane



Posteitaliane





Connect

Danilo Maggi

Marketing Manager, Red Hat Italia



Sessioni pomeridiane

Orario	Sessione	Sala
14.00 - 17.00	Red Hat	Auditorium 1° piano
	Hybrid Cloud & Security & AI/ML	Sala Piano Terra
	Hybrid Cloud & Data & AI/ML	Sala Piano Terra
	IT Automation & Security & AI/ML	Sala Piano Terra
	Application & Virtualization Modernization & IT Automation / Observability	Sala Piano Terra
	Application/Virtualization Modernization & AI/ML	Sala Piano Terra
	Application/Virtualization Modernization & AI/ML	Sala Piano Terra
	RIOS - Sviluppare la fiducia con l'Open: compliance, governance, AI	Sala Piano Terra
	Silent Theatre Microsoft	Booth Microsoft Piano Terra
	Silent Theatre Red Hat Partner	Area espositiva Piano Terra
	Sessione Difesa (accesso riservato)	Sala Piano Terra

Red Hat
Summit

Connect

Grazie



[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

