

September 2024

# Red Hat Modernization

Modernizing & Advancing the Data Center

**Partner Enablement Package**

How our valued partners can build modernization solutions and advance the data center with AI, based on latest Intel® technology

intel®



Red Hat

# Contents



## Introduction

- Value Proposition
- AI Ecosystem Support
- Why Intel + Red Hat



## CentOS EOL

- CentOS Linux EOL
- Deploy RHEL on Intel
- Why Choose 5<sup>th</sup> Gen Intel® Xeon® Processors for Server Refresh



## OpenShift Virtualization

- Server Virtualization Opportunity
- OpenShift Virtualization Overview & Business Value
- Server Consolidation Benefits
- Live Migration Stability



## OpenShift AI

- AI Market Opportunity
- Transformative AI Joint Offerings
- OpenShift AI Overview
- Intel + Red Hat AI
- Red Hat® Validated Patterns
- Intel® Gaudi® AI Accelerator



## Call to Action



## Resources



# Benefits & Value Proposition



## For SIs / SPs

- Connect with your customers, show them how to save costs, deliver improved performance and modernize to a cloud native offering
- Server & Software upgrades unleash new AI & Security capabilities
- Identify multiple revenue streams by articulating the value of modernizing the data center

## For End Customers

- Achieve lower TCO through server consolidation, resulting in energy savings
- Reduce the number of servers -> lower SW costs
- Cloudify the data center with OpenShift
  - Single platform to manage VMs and Containers
  - Flexibility to run On-Prem and in the Cloud

Why Faster Refresh Cycles And Modern Infrastructure Management Are Critical To Business Success

# AI Ecosystem Support

"The AI landscape is evolving at breakneck speed, and organizations need solutions that are both agile and powerful to keep pace. Red Hat OpenShift AI, combined with the breadth and performance of Intel's AI portfolio, offers a compelling solution for businesses to accelerate their AI journey and unlock new possibilities."

*Chris Wright, chief technology officer and senior vice president, Global Engineering, Red Hat*



"Cisco is actively working to advance generative AI technology to bring new and exciting opportunities to customers, such as faster innovation, enhanced decision-making and mitigated risks. We strongly support the collaboration between Red Hat and Intel as we continue to work closely with them to offer new and innovative solutions for our joint customers."

*Daniel McGinnis, vice president, Product Management, Cisco*



"Dell Technologies helps customers take full advantage of artificial intelligence by enabling the deployment of AI models and applications, especially generative AI, across enterprise environments. Dell collaborates with Red Hat and Intel to make it easy for customers to deploy and operate Red Hat OpenShift AI on Dell APEX Cloud Platform and Dell PowerEdge with Intel accelerators."

*Greg Findlen, senior vice president, Product Management, Dell Technologies*



"Together, Red Hat and Intel are committed to helping organizations develop and unlock the full potential of scalable and more secure AI systems based on open ecosystems and standards. Intel's AI technologies help minimize the complexities associated with AI workloads and unlock improved performance, empowering businesses to innovate faster and achieve their AI goals."

*Justin Hotard, executive vice president and general manager, Data Center and AI Group, Intel*



"Lenovo is committed to shaping a future where AI enhances every aspect of computing environments. This commitment is embodied in our collaboration with Red Hat and Intel, aiming to harness the full potential of AI through our Lenovo Edge-to-Cloud Solutions. By using Red Hat OpenShift AI, powered by the latest 5th Gen Intel Xeon Scalable Processors with Intel AMX, and soon-to-include Intel Gaudi 3 AI Accelerators, we are optimizing and speeding up AI deployments. This synergy allows us to deliver sophisticated AI solutions that operate seamlessly across the hybrid cloud, meeting the evolving needs of our diverse clientele."

*Kamran Amini, vice president and general manager, Server, Storage & Software Defined Solutions, Lenovo*



"At Supermicro, we envision a world where AI is seamlessly integrated into every facet of our computing solutions. To bring this vision to life, Supermicro is working closely with industry leaders like Red Hat and Intel. By leveraging Red Hat OpenShift AI and the powerful capabilities of 5th Gen Intel Xeon Scalable Processors with Intel AMX and Intel's Gaudi 2 and Gaudi 3 AI Accelerators—we are enhancing and accelerating AI operations across hybrid cloud environments, ensuring our customers have access to the most advanced AI tools on Supermicro servers."

*Ray Pang, vice president, Technology & Business Enablement, Supermicro*





# Intel and Red Hat: Leaders in Open Technology

---

Together, Red Hat and Intel provide agile, cloud-ready network architectures based on high-performance, industry-standard platforms and open, software-defined infrastructure.

Partner Showcase:  
Intel and Red Hat - Leaders in Open Technology



25 years in  
**partnership**

**Top contributor**  
to Linux kernel  
and open source



**Solutions**  
optimized  
through joint  
roadmaps

# Why Intel + Red Hat

## The Latest Technology

Intel collaborates closely with Red Hat to implement the newest technologies and software with a focus on optimization. Customers gain leading architecture, workload-optimized containers, and integrated solutions.

## Advanced Security

Intel® Software Guard Extensions (Intel® SGX) and Intel® Trust Domain Extensions (Intel® TDX) provides advanced data protection.

(\*\* Powered by Intel® Device Plugins Operator for Red Hat OpenShift Container Platform)

[Check out the Intel® Device Plugins Operator here-->](#)

## Advanced AI Tools

Developers at industry-leading independent software vendors (ISVs) and enterprises use Intel® tools and framework optimizations to build their AI platforms, systems, and applications. Developers can learn, test, and deploy AI tools such as Intel® AI Analytics Toolkit (AI Kit), Intel® Distribution of OpenVINO™ toolkit, Intel® Tiber™ AI Studio, and Habana, using the Red Hat OpenShift AI sandbox.

[Check out the Red Hat OpenShift AI Sandbox here-->](#)

## Featured Solutions

### **Streamline Enterprise AI Adoption and Deployment**

Combining Red- Hat OpenShift AI and Intel Enterprise AI provides a unified, consistent, validated, and integrated hardware and software solution for enterprise AI, right out of the box.


### **GenAI Inferencing on UCS X-Series with 5th Gen Intel® Xeon® Scalable Processors on Red Hat OpenShift AI**



# Intel® Xeon® Scalable processors for your data center

## Your business can count on:

- Shared **architectural platform** for 4th Gen Intel® Xeon® processors and 5th Gen Intel® Xeon® processors
- Designed to **accelerate performance** across the fastest-growing workloads
- Most **built-in accelerators** of any CPU on the market to help maximize performance and power efficiency for emerging workloads
- Advanced **security technologies** to help protect data
- Scale with the **most choice**, no matter the deployment path
- Intel's most **sustainable** data center processors



# Upgrade Opportunity: CentOS Linux

**Support ended on June 30<sup>th</sup> 2024** Current CentOS Linux users now need to find an alternative Linux offering to maintain support

Migration to **Red Hat Enterprise Linux (RHEL)** gives you an opportunity to gain more value for your organization



## Security resources

Industry-leading security response and data



## Partner ecosystem

4,900 ISVs, 5,500 IHVs



## Life cycle support and flexibility

Reliable cadence, 10+ years major, 2+ year minors



## Product leadership

Lead the Enterprise Linux ecosystem and roadmap



## Proactive analytics

Vulnerability reporting, cost management, etc.



## Support and expertise

Industry-leading team, TAMs



More than

# 90%

of companies in the Fortune 500 rely on Red Hat technologies<sup>1</sup>

## How to Migrate

[Watch Video](#)



Streamline your migration with an automated conversion tool

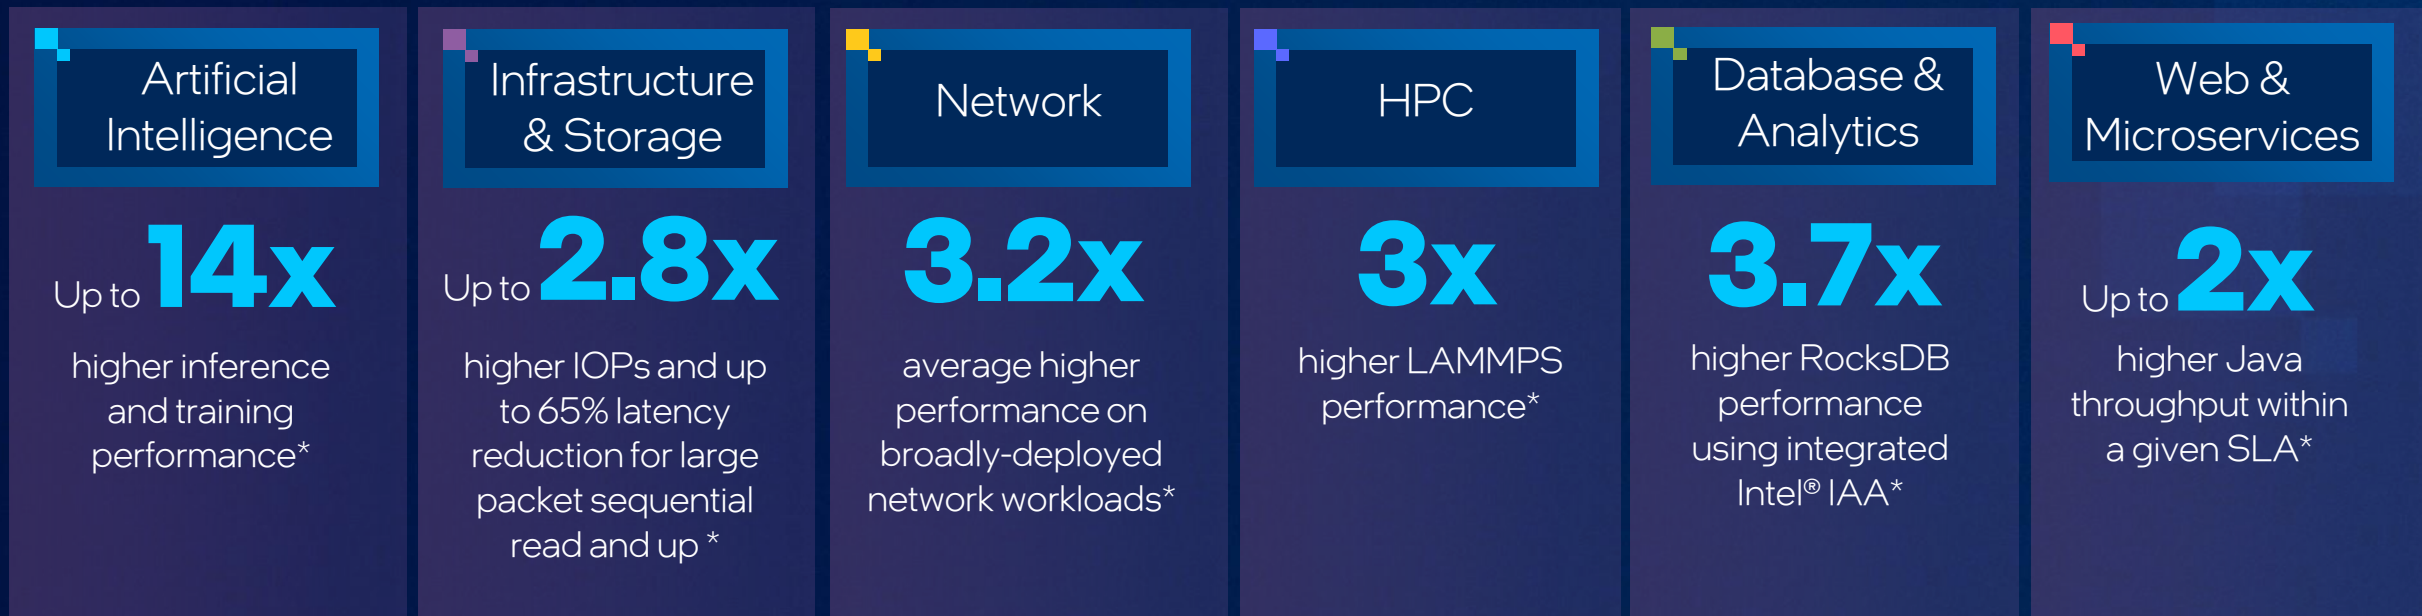
[Convert2RHEL](#)

<sup>1</sup><https://www.redhat.com/en/about/company>



# Deploy RHEL on Intel® Xeon®

Upgrade to 5th Gen Intel® Xeon® Scalable processors for significant performance gains



5th Gen Intel® Xeon® Scalable processors vs. 3rd Gen Intel® Xeon® processors

**Migrate to RHEL on Intel® Xeon® Scalable processors to access significant performance gains (as above) + reduced power & TCO**

READ MORE >

[RHEL on 5th Gen Intel® Xeon®](#)

Access the [Intel® Xeon® Processor Advisor Suite](#) to calculate the best route to lower TCO and your ROI

\*See [A17, N16, N11, H10, D1, W3] at intel.com/processor claims: 5th Gen Intel® Xeon® Scalable processors. Results may vary.

# Why Choose 5th Gen Intel® Xeon® processors for Server Refresh?

intel.  
xeon®

## Lower Total Cost of Ownership (TCO)

Intel's portfolio of hardware, software, systems, and tools can help advance your data center's overall efficiency, creating energy savings and reducing your carbon footprint, without sacrificing performance, while giving you the TCO and flexibility you need.

Up to  
**77%**  
reduction  
in TCO<sup>1</sup>

## Increase Efficiency

Intel® Accelerator Engines boost CPU utilization, reduce electricity consumption resulting in lower impact on the environment.

**10x**  
better efficiency  
(perf/watt)  
with built-in  
accelerators<sup>3</sup>

## Gain Better Performance

By delivering more performance per core with built-in accelerators, 5th Gen Intel® Xeon® processors help you meet requirements for even the most demanding workloads.

**84%**  
Performance  
gain<sup>2</sup>

## Optimize the Data Center

Intel® Xeon® processors deliver the low-latency, high-bandwidth capabilities required by modern and AI-infused workloads. Replacing aging infrastructure with these speedy and energy efficient processors will help you keep pace with rapidly evolving market needs.

Up to  
**16:1**  
server  
consolidation<sup>1</sup>

## Stay Secure with Confidential Computing

With Intel, you can choose from the most deployed confidential computing options in data centers on the market today—now including application or VM-level isolation.

**300+**

Organizations have engaged with Intel to develop and deploy Confidential Computing services<sup>4</sup>

**LEARN MORE**

What's the right transition  
for your customer?

<sup>1,2,3</sup> See [T7, G1, T13] at [intel.com/processorclaims:5th Gen Intel® Xeon® Scalable processors](https://www.intel.com/processorclaims:5th%20Gen%20Intel%20Xeon%20Scalable). Results may vary

<sup>4</sup> Source: [Four Facts: Intel at the Foundation of Confidential Computing](#)



# Server Virtualization Opportunity

Global server virtualization software market size was

**\$7.88 billion** in 2022<sup>1</sup>

Market is projected to reach

**\$12.32 billion** by 2031<sup>1</sup>

## Challenges of traditional Virtualization environments

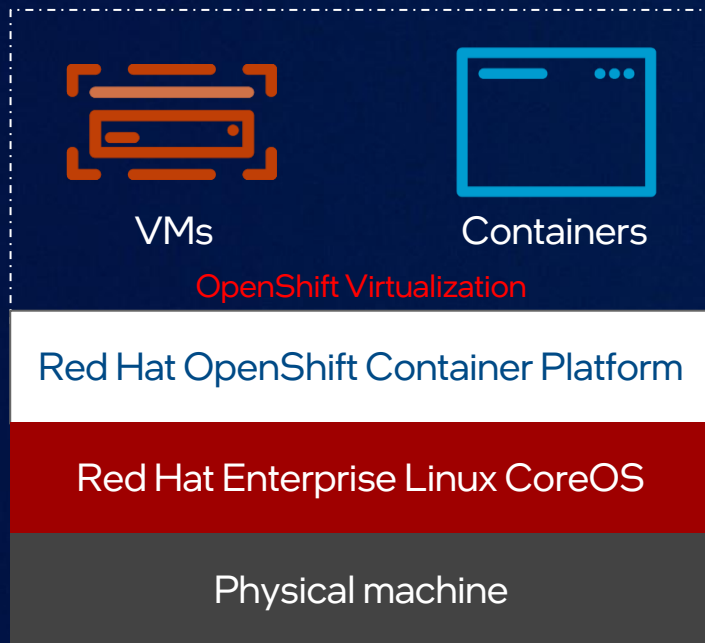
-  Increasing cost
-  Slow Evolution
-  Supporting Growth
-  Risk
-  Developer Productivity

## Benefits of Virtualized Environments

-  Cost reduction for operating infrastructure
-  Innovate at speed
-  Scalability
-  Security focused
-  Integrated development tools

# OpenShift Virtualization

Red Hat OpenShift Virtualization supports virtual machines alongside containers and serverless workloads, enabling a single platform and common set of enterprise tools to support rapidly evolving infrastructure.



## Features:

- Run VMs in OpenShift
- Performance, stability, scalability, and reliability of KVM, the Linux kernel-based hypervisor
- RHEL guest entitlements are included
- Supports Microsoft Windows guests - Microsoft Server Virtualization Validation Program (SVVP)
- Manageability and ecosystem of OpenShift
- Unified platform for running VMs and Containers
- Included feature of the OpenShift application platform



# OpenShift Virtualization Business Value

## Business Value:

- Trusted by FSI, healthcare, and government for enterprise software
- Hybrid cloud ready — Built for DevOps and DevSecOps practices for greater agility and flexibility
- Migration stability with similar CPU architecture
- Support for multiple Confidential Computing options for virtual machines and containers



Benefits to upgrading infrastructure

From 1 <sup>st</sup> to 5 <sup>th</sup> Gen Intel <sup>®</sup> Xeon <sup>®</sup> processor	Reduce servers	Reduce energy and CO2	Reduce TCO	Recover costs (months)
8168 → 8562Y+	60%	44%	41%	19
6138 → 8558	74%	39%	59%	10
5118 → 5520+	74%	55%	66%	5
4110 → 5520+	82%	63%	76%	3

\*EMEA energy costs greater than US. See backup for configurations. Results may vary.

# Save Power and Money on New Server Purchases

Deploy fewer 5th Gen Intel® Xeon® processor-based servers to meet performance and TCO goals

Comparisons of replacing **50** servers based on 3rd Gen Intel® Xeon® with new 5th Gen Intel® Xeon® processors

	Web (NGINX TLS)	Data Services (RocksDB)	Artificial Intelligence (NLP w/ Bert-Large)	Artificial Intelligence (Recommender w/ DLRM)
Number of 5th Gen Intel® Xeon® processor-based servers	29 servers	14 servers	13 servers	10 servers
Lower fleet energy	870 MWh	1,482 MWh	1,644 MWh	1,705 MWh
Reduced CO2 emissions*	368,757 kg	628,372 kg	696,768 kg	722,607 kg
TCO savings*	\$442K	\$1,198K	\$1,273K	\$1,425K

\* Estimated over 4 years.

\*See [T10, T9, T11, T12] at [intel.com/processor](https://www.intel.com/processor) claims: 5th Gen Intel® Xeon® Scalable processors. Results may vary. 14





# Replace Aging Servers to Save Energy and Costs

Significantly reduce data center infrastructure space, power and costs

Comparisons of replacing **50** servers based on 1st Gen Intel® Xeon® with new 5th Gen Intel® Xeon® processors

	Web (NGINX TLS)	Data Services (RocksDB)	Artificial Intelligence (NLP w/ Bert-Large)	Artificial Intelligence (Recommender w/ DLRM)
Number of 5th Gen Intel® Xeon® processor-based servers	7 servers	9 servers	3 servers	4 servers
Lower fleet energy	843.5 MWh	1191.5 MWh	1697.1 MWh	1259.2 MWh
Reduced CO2 emissions*	357,586 kg	505,136 kg	719,458 kg	533,822 kg
TCO savings*	\$254K	\$192K	\$541K	\$449K

Access the [Intel® Xeon® Processor Advisor Suite](#) to calculate the best route to lower TCO and ROI

# Live Migration Stability

Bring cloud-native functionality to virtual machines with Red Hat in a single open platform

## Why Migrate?

Empower operations teams to innovate

Take advantage of a cloud-like experience, everywhere

Benefit from a trusted enterprise open source platform

Experience a complete application platform that integrates with existing infrastructure, tools, and services

Intel recommends migration from **Xeon<sup>®</sup>** to **Xeon<sup>®</sup>** whenever possible:

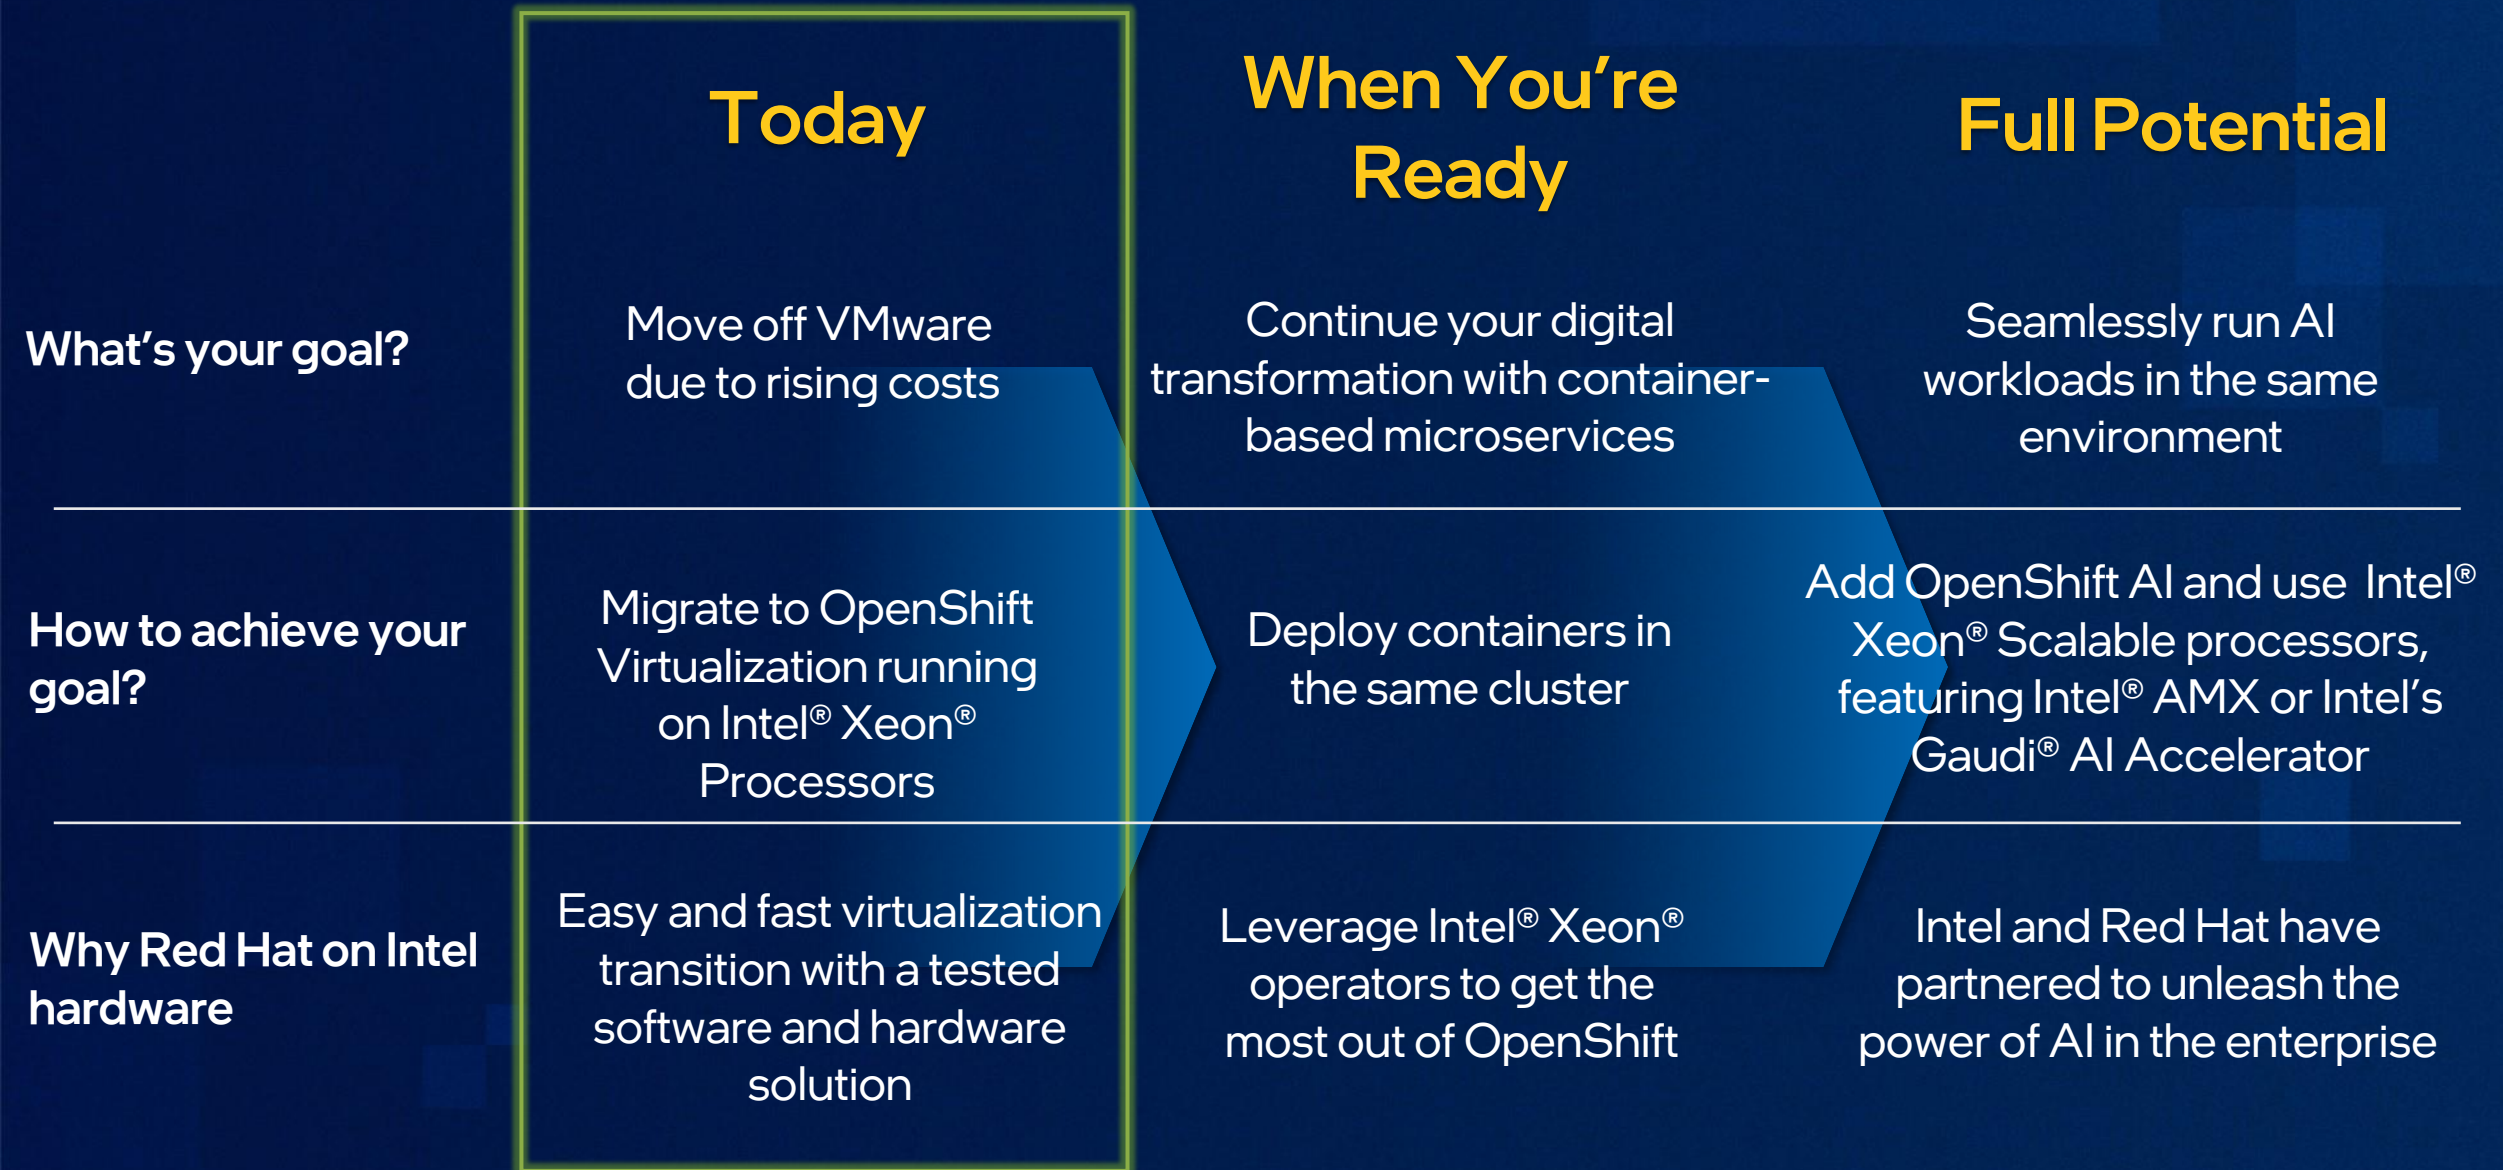


- Live migration requires a similar CPU architecture to reduce risk in the migration process
- Red Hat Migration Tool Recommendations:
  - *"In all cases, the CPU features of the destination node must be a superset of the CPU features on the source node."*  
Source: Red Hat Migration [Webpage](#)

Get Started with the [Migration Toolkit for Virtualization](#)



# Why Partner with Intel for Digital Transformation



# Summary

1. OpenShift Virtualization provides a mature, stable and scalable solution that's easy to get started
2. Full automation including; declarative, self-healing and secure
3. Integration with a broad ecosystem of partners
4. Part of an application platform that allows added features and capabilities

[ACCESS NOW >](#)

[Red Hat OpenShift Virtualization + Intel — Better Together Pitch Deck](#)

## Next Steps

Engage customers to understand if they are happy with their existing Virtualization offering or if familiar with OCP Virtualization

Recommend OpenShift Virtualization running on Intel® 4th or 5th Gen Xeon® processors & Connect with corresponding Red Hat/Intel AE

**Begin Migration Journey**



# Every Company Will Be an AI Company



Virtually every company will become an “**AI company,**” AI will be used to create new services, improve internal business processes, and/or improve employee productivity

## AI is now central to business strategy

**79%** of corporate strategists say AI is critical to success in 2024<sup>1</sup>

**37%** CAGR, 2023-2330 growth is underway in the global AI market<sup>2</sup>

GenAI is supercharging AI's potential

**80%** of enterprises will use GenAI by 2026<sup>4</sup>

AI architects face too many choices

**46%** of experts say infrastructure is the biggest challenge in productizing LLMs<sup>3</sup>

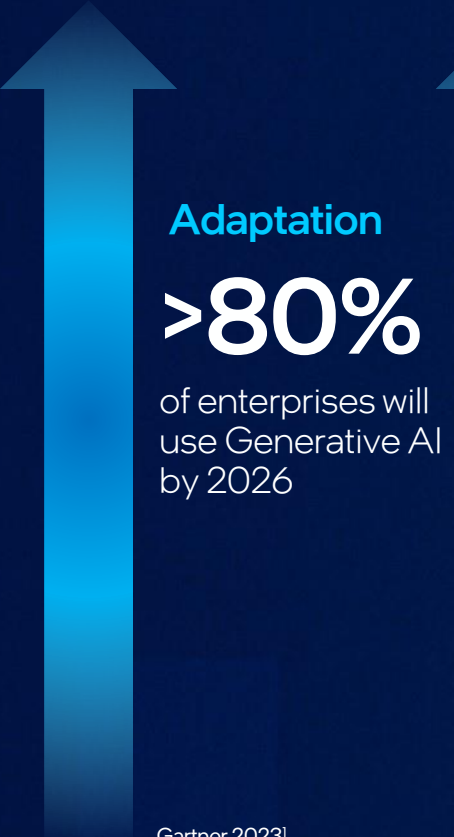
<sup>1</sup> Skim AI, “10 Enterprise AI Statistics to Know in 2024.” <https://skimai.com/10-enterprise-ai-stats-to-know-in-2024/>.

<sup>2</sup> Grandview Research, “Artificial Intelligence Market Size and Share Report, 2030.” <https://www.grandviewresearch.com/industry-analysis/artificial-intelligence-ai-market>.

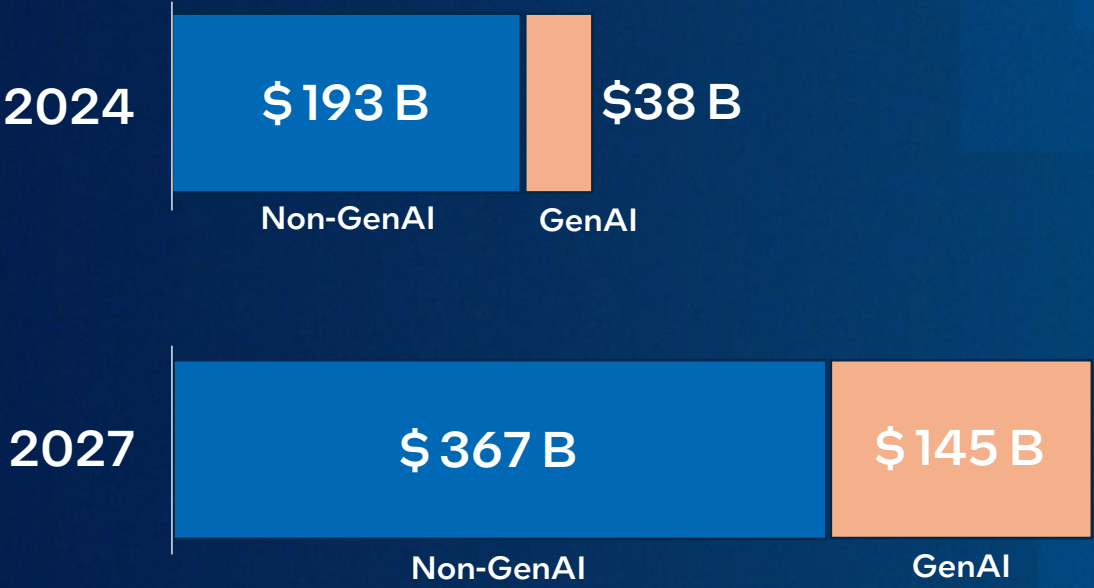
<sup>3</sup> mlinsider. “The state of Generative AI and Machine Learning at the end of 2023.” [https://cnvrg.io/wp-content/uploads/2023/11/ML-Insider-Survey\\_2023\\_WEB.pdf](https://cnvrg.io/wp-content/uploads/2023/11/ML-Insider-Survey_2023_WEB.pdf).

<sup>4</sup> Gartner, October 11, 2023. “Gartner says More than 80% of Enterprises Will Have Used Generative AI APIs or Deployed Generative AI-Enabled Applications by 2026.” <https://www.gartner.com/en/newsroom/press-releases/2023-10-11-gartner-says-more-than-80-percent-of-enterprises-will-have-used-generative-ai-apis-or-deployed-generative-ai-enabled-applications-by-2026>.

# AI Inflection Point



## Expected Spend on Artificial Intelligence



IDC AI and Generative AI Spending Guide Forecast<sup>2</sup>

1 <https://www.gartner.com/en/documents/4558899>  
2 <https://www.idc.com/fdsfile/download/0cb3a164-e692-412b-a7ce-514e5f3bcee3>



# Transformative AI Joint Offerings



Enterprise AI  
from **intel**®


- Red Hat and Intel offer **validated and integrated** hardware and software elements for enterprise AI
- Together, they **reduce the complexity** of selecting and integrating solution components
- Customers attain **faster time to value**, with lower cost and less risk

# Red Hat® OpenShift® AI




## Red Hat OpenShift AI

Provide data scientists an  
open ML platform



### Built on top of OpenShift

A secure, supported and proven  
enterprise platform on hybrid cloud



### Designed for machine learning

An easy-to-manage platform, with hardware  
acceleration and an ecosystem of ML tools



### Empowered data science

A self-service infrastructure with foundational  
elements from open source



# OpenShift is a Trusted Enterprise Kubernetes

## Red Hat OpenShift

Red Hat OpenShift offers a comprehensive platform powered by Kubernetes to efficiently build, update, and scale applications. Enhance productivity and speed up the application deployment process by using a full suite of services, all adaptable to your preferred infrastructure.

## Red Hat OpenShift AI


Red Hat OpenShift AI serves as a versatile, scalable MLOps platform equipped with tools for constructing, deploying, and managing applications powered by AI. OpenShift AI was developed using open-source technologies: it offers reliable, operationally consistent features for teams to experiment, serve models, and launch innovative applications. OpenShift AI supports the complete lifecycle of AI/ML experiments and models, both on-premises and in the public cloud.

[Learn more](#)

[Integrate and simplify production AI with Intel and Red Hat](#)

The industry is moving towards a Hybrid Multi-Cloud operating model. Red Hat OpenShift is the most widely deployed multi-cloud container platform at **47%**

**MSS**



# Red Hat Supercharges AI Workloads with Intel® Technologies from the Datacenter to the Edge



- Red Hat OpenShift AI and Intel's AI technologies enable organizations to accelerate generative AI and machine learning innovation on a trusted, integrated platform
- Red Hat OpenShift AI now supports Intel® Gaudi® AI accelerators
- Get Started and accelerate AI development with [Intel® Tiber™ Developer Cloud](#)

[PRESS RELEASE](#)

Red Hat and Intel's open-source approach to AI innovation empowers organizations to deploy anywhere on their platform of choice, accelerating time to market and providing ubiquitous AI building blocks that are more cost-effective at scale.



# Intel AI Software

## Intel AI Tools

### Data Analytics

Distributed data ingestion and preprocessing

Intel® Distribution of Modin

### Machine Learning

Training and prediction

Intel® Extension for Scikit-learn

Intel® Optimization for XGBoost

### Deep Learning

Modeling, training, transfer learning, fine-tuning, inference

PyTorch Optimizations from Intel

TensorFlow Optimizations from Intel

### Optimization & Inference

Perform model compression and deploy to the cloud, edge, or PC

Intel® Neural Compressor

OpenVINO

Intel® AI Reference Models

Pretrained, Intel-optimized models

Intel® Tiber™ Studio

Full-stack machine learning operating system

Intel® Gaudi® Software

Leverage dedicated AI accelerators



Open Platform for Enterprise AI



Intel® Tiber™ Edge Platform



Intel® Tiber™ Developer Cloud



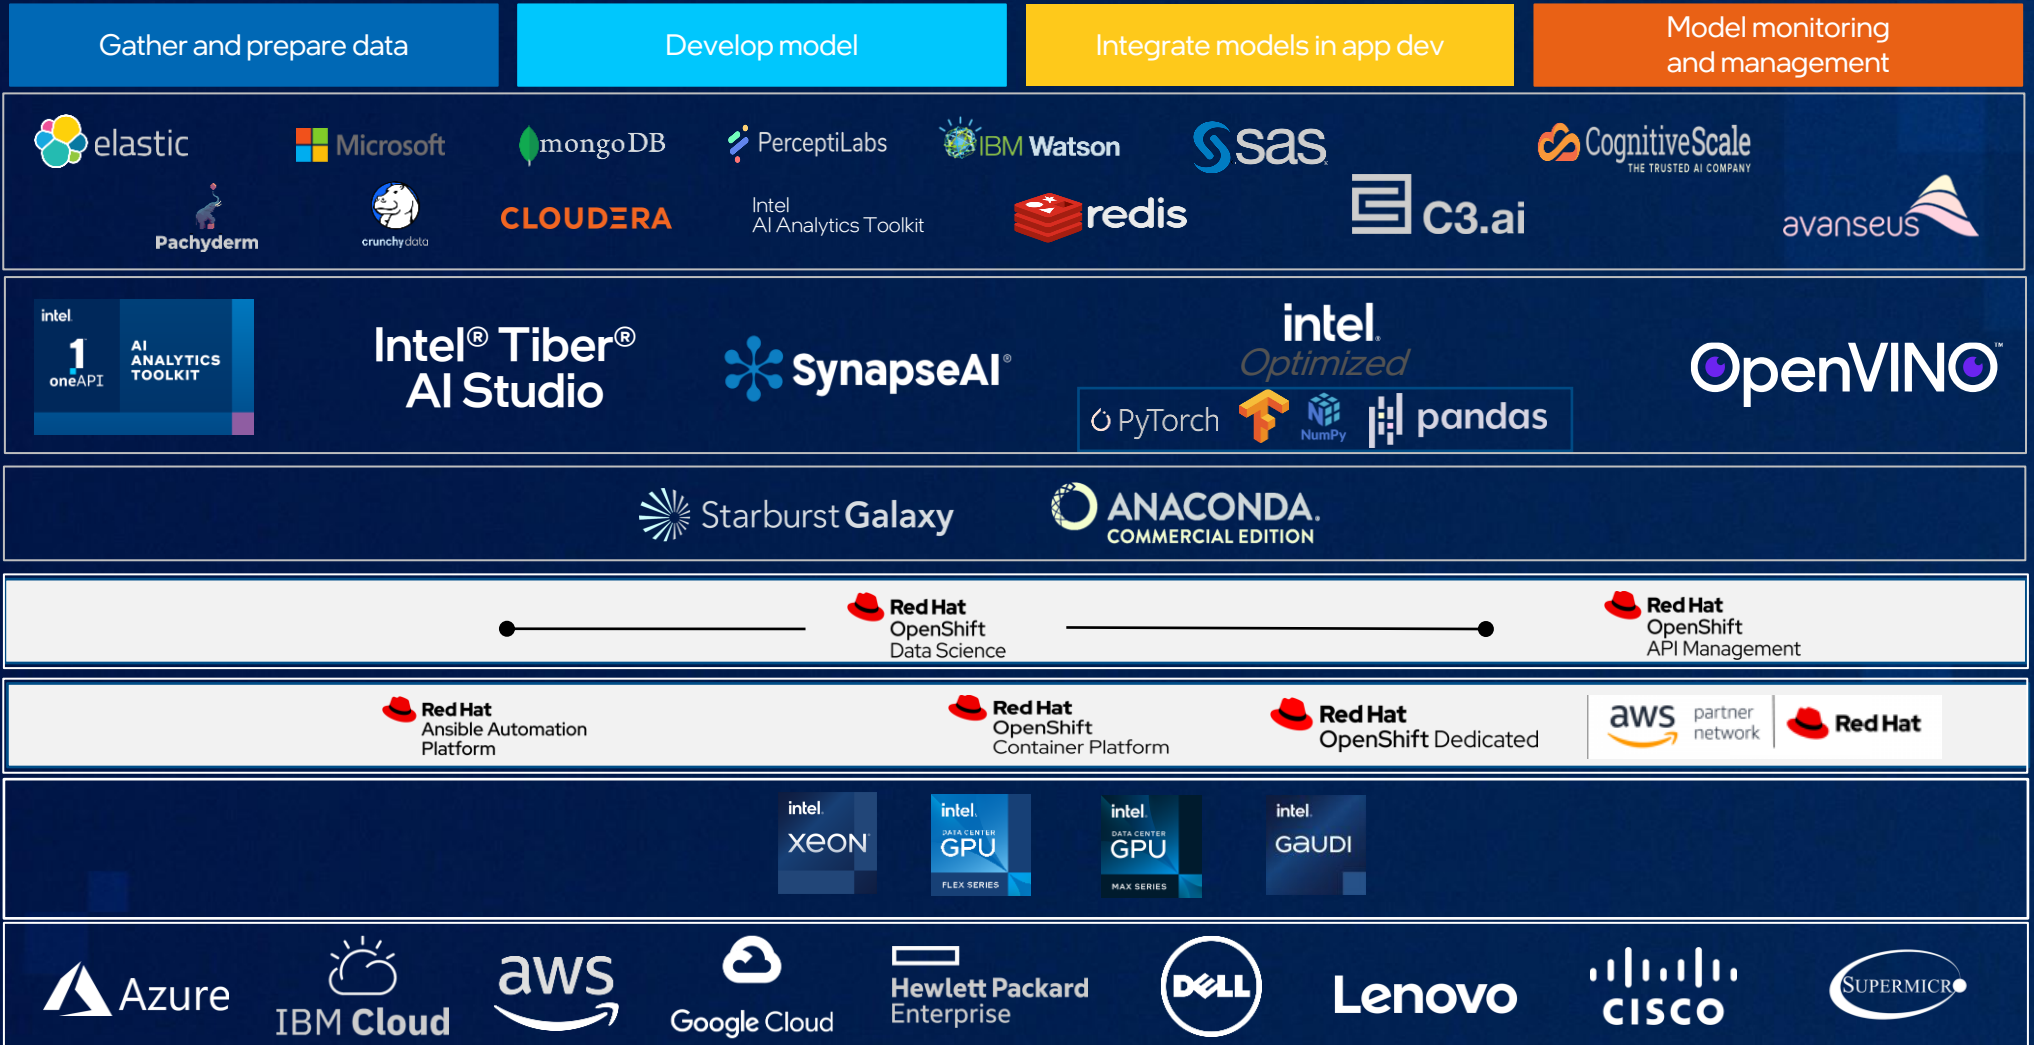
oneAPI Open, Cross-Architecture Programming Model

# Red Hat® + Intel® AI

## Integrating + Simplifying the AI Ecosystem

Boosting Business with AI

Red Hat\* and Intel® AI and ML: The Perfect Combination for Data Scientists






# Accelerate Red Hat® OpenShift® AI Workflows: Using Intel's Newest Processor Features with Red Hat® Validated Patterns

## Solution Benefits

Extended Red Hat® validated patterns from Intel make it easy to enable AI and security features :

- **Effortless deployment** of comprehensive, fully operational AI workflows, eliminating the need for extensive manual configurations
- **Accelerated AI model training and inference** with Intel® Advanced Matrix Extensions
- **Enhanced security** using Intel® Software Guard Extensions, which supports confidential computing
- **Offloading of encryption and compression** to a dedicated accelerator to improve application performance, using Intel® Quick Assist Technology

## PROOF POINT



\*configuration details in Solution Reference Architecture linked across, results may vary

READ MORE

[Solution Reference Architecture](#)

### Solution Architecture Highlights

Figure 1 illustrates the solution stack for Intel's extension of Red Hat validated patterns. Through node feature discovery, the patterns can determine if a particular feature, such as Intel AMX or Intel SGX, is available. OpenShift AI, Kubernetes management software, and other pattern components all integrate with OpenShift.

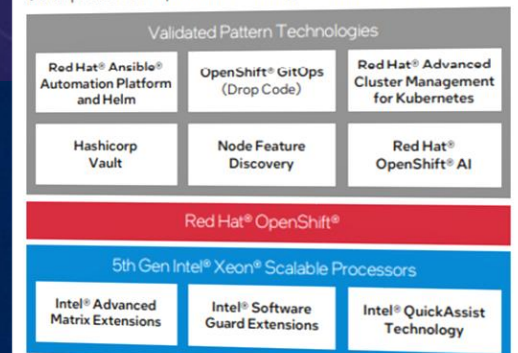


Figure 1. Extending Red Hat® validated patterns to include Intel® technologies can accelerate AI workloads and enhance security.

# Improve Time-to-Value with Validated Patterns

With **Red Hat Validated Patterns**, organizations can find tested and supported solutions for many AI use cases, from customized chatbots to fraud detection to medical diagnosis.

Benefits include:

- **Automatic deployment** of full application stack
- **Wide range of customization**, including Intel® accelerators like:
  - Intel® Advanced Matrix Extensions (Intel® AMX)
  - Intel® Quick Assist Technology (Intel® QAT)
  - Intel® Software Guard Extensions (Intel® SGX)
  - Intel® Trust Domain Extensions (Intel® TDX)
- **Variety of use cases** including Medical Diagnosis and Fraud Detection
- List of validated patterns with maturity and support information is publicly available on <https://validatedpatterns.io/patterns/>

**LEARN MORE >**

[Integrate and Simplify Production AI with Intel and Red Hat](#)



# Intel Enterprise AI for Red Hat® OpenShift® AI and Red Hat® Enterprise Linux

SOLUTION BRIEF

[Streamline AI Adoption and Deployment Using Intel Enterprise AI with Red Hat® OpenShift® AI](#)

Customer and Ecosystem Managed Software

Intel® Tiber™  
Edge Platform

PyTorch

OpenVINO™

Intel® Tiber™  
AI Studio

 **Red Hat**  
OpenShift AI

 **Red Hat**  
Ansible Automation  
Platform

 **Red Hat**  
OpenShift  
Container Platform

 **Red Hat**  
Enterprise Linux

intel  
CORE  
ULTRA

intel.  
ARC  
GRAPHICS

intel  
XEON

intel  
GAUDI

intel.  
IPU

Intel brings together a rich set of scalable open source and commercial software, ready to go out of the box, which can be integrated with Red Hat OpenShift AI

Intel® Tiber™ Edge  
Platform

OpenVINO™

Intel® Tiber™  
AI Studio

oneAPI

Intel® Tiber™  
Developer Cloud

# Accelerating Generative AI and Large Language Models with Intel® Gaudi® AI Accelerator

intel.  
GAUDI

Intel® Gaudi® 2 delivers leading performance and optimal cost savings for AI training

[Press Release](#)

The Gaudi® 2 deep learning accelerator performs competitively on deep learning training and inference, with up to **2.4x faster performance than Nvidia A100**<sup>1</sup>

[Newsroom](#)  
[Tech Article](#)

Gaudi® 2 delivers compelling performance vs. Nvidia's H100<sup>2,3</sup> for GPT-3 and GPT-J

[Newsroom](#)  
[ML Commons Announcement](#)

COMING SOON:

## Intel® Gaudi® 3 AI accelerator

[Bringing Choice to GenAI with Performance, Scalability and Efficiency](#)

[Intel® Gaudi® 3](#) will deliver a significant leap in AI training and inference for global enterprises looking to deploy GenAI at scale

[READ MORE](#)

[Intel® Gaudi® 2 Remains Only Benchmarked Alternative to NV H100 for GenAI Performance](#)

[Intel Breaks Down Proprietary Walls to Bring Choice to Enterprise GenAI Market](#)

<sup>1</sup>Performance varies by use, configuration, and other factors; workloads and configuration details available at: [intel.com/performanceindex](https://intel.com/performanceindex) Results may vary.

<sup>2,3</sup>Performance varies by use, configuration, and other factors; workloads and configuration details available at: <https://mlcommons.org/2023/09/mlperf-results-highlight-growing-importance-of-generative-ai-and-storage/> Results may vary.



# Next Steps



[Sign up for Intel® Tiber™ Developer Cloud](#)



[Try a 60-day trial of Red Hat OpenShift AI](#)



[Download pre-packaged cloud-native Intel® AI software](#)

## MORE INFORMATION

- [Red Hat® and Intel® AI and Machine Learning: The Perfect Combination for Data Scientists](#)
- [Essential Tools for Jumpstarting AI Development Projects](#)
- [How to Use Intel®-Optimized AI Software in the Cloud](#)
- [Speed Up Machine Learning Training on CPUs with AI Tools](#)
- [Deploying Intel Enterprise AI with Red Hat OpenShift AI Made Simple](#)
- [Red Hat OpenShift AI + Intel — Better Together Pitch Deck](#)

# Developer Resources

- Create a Red Hat OpenShift AI Account
  - [Developer Sandbox](#)
- Intel Technology Enabling for OpenShift
  - [GitHub](#)
- [Developer Resources from Intel and Red Hat](#)
- [Optimize AI with Powerful Drop-In Acceleration from Intel & Red Hat\\*](#)
- [How to Get the Most Out of Red Hat OpenShift\\* AI with Intel® AI Tools](#)
- [Streamline AI Adoption and Deployment Using Intel Enterprise AI with Red Hat® OpenShift® AI](#)

AI application evaluation on Intel® hardware through Red Hat\* OpenShift\* AI platform using the Intel® Developer Cloud for the Edge

- [devcloud.intel.com/edge](https://devcloud.intel.com/edge)
- Container Playground Entry point
  - <https://www.intel.com/devcloud-containers>
- Container Playground Developer Guide
  - <https://www.intel.com/content/www/us/en/develop/documentation/devcloud-containers/top.html>



# Notices and Disclaimers

- Performance varies by use, configuration and other factors. Learn more on the [Performance Index site](#).
- Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.
- Your costs and results may vary.
- Intel technologies may require enabled hardware, software or service activation.
- © Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

The Intel logo is centered on a dark blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small, light blue square is positioned above the letter "i". To the right of the word "intel" is a registered trademark symbol (®).

intel®



# What's the right transition for your customer?

Guide your customer to the best refresh option	Current install base	Good	Better	Best
	Intel Xeon Silver 4000 Series	4514Y 16C/150W	4516+ 24C/185W	5520+ 28C/205W
	Intel Xeon Gold 5000 Series	5515Y 8C/185W	5520+ 28C/205W	6530 32C/270W
	Intel Xeon Gold 6000 Series	6526Y 16C/195W	6548Y 32C/250W	8558 48C/330W
	Intel Xeon Platinum 8000 Series	8562Y+ 32C/300W	8568Y+ 48C/350W	8592+ 64C/350W

intel.  
**XEON**

## Exceptional Performance & Efficiency

5th Gen Intel® Xeon® processors deliver impressive performance-per-watt gains across all workloads, plus outsized performance and lower TCO

- Optimize AI, HPC, network, data analytics & storage workloads with Intel® Accelerator Engines
- Enhanced platform capabilities
  - 3x increase in shared last-level cache
  - PCIe 5 – Double I/O bandwidth
  - DDR5 – Increase memory bandwidth
  - CXL® 1.1 – Next Gen I/O for low latency and performance
  - Optimized Power Mode enables energy savings with minimal impact of performance\*
  - Advanced security technologies to help protect data with Intel Software Guard Extensions (Intel SGX) and Intel® Trust Domain Extensions (Intel TDX)
- To learn more about SKU transitions and refresh opportunities:
  - <https://xeonprocessoradvisor.intel.com>

Refresh your infrastructure today, to be ready for tomorrow's demands

Refresh from  
1<sup>st</sup> Gen Intel® Xeon® to  
5<sup>th</sup> Gen Intel Xeon

**AI**  
(BertLarge)

up to **42x**  
higher performance<sup>1</sup>

up to **20x**  
higher performance/  
watt<sup>1</sup>

Refresh from  
2<sup>nd</sup> Gen Intel® Xeon® to  
5<sup>th</sup> Gen Intel Xeon

**AI**  
(DLRM)

up to **10x**  
higher performance<sup>2</sup>

up to **6x**  
higher performance/  
watt<sup>2</sup>

Intel vs. AMD

**AI**  
(Recommender - DLRM)

up to **2.8x**  
higher batched inference  
performance<sup>3</sup>


up to **2.6x**  
higher performance/  
watt<sup>3</sup>



**TCO Advisor Tool**

<sup>1,2,3</sup> See [A37, A38, A208] at [intel.com/processorclaims](https://www.intel.com/processorclaims): 5<sup>th</sup> Gen Intel Xeon Scalable processors. Results may vary  
\*Enabled in platform BIOS, visit [intel.com/processorclaims](https://www.intel.com/processorclaims): 5<sup>th</sup> Gen Intel Scalable processors for more information

# What's the right transition for your customer?

Guide your customer to the best refresh option	Current install base	Good	Better	Best
	Intel Xeon Silver 4000 Series	4514Y 16C/150W	4516+ 24C/185W	5520+ 28C/205W
	Intel Xeon Gold 5000 Series	5515Y 8C/185W	5520+ 28C/205W	6530 32C/270W
	Intel Xeon Gold 6000 Series	6526Y 16C/195W	6548Y 32C/250W	8558 48C/330W
	Intel Xeon Platinum 8000 Series	8562Y+ 32C/300W	8568Y+ 48C/350W	8592+ 64C/350W

intel.  
**XEON**

## Exceptional Performance & Efficiency

5th Gen Intel® Xeon® processors deliver impressive performance-per-watt gains across all workloads, plus outsized performance and lower TCO

- Optimize AI, HPC, network, data analytics & storage workloads with Intel® Accelerator Engines
- Enhanced platform capabilities
  - 3x increase in shared last-level cache
  - PCIe 5 – Double I/O bandwidth
  - DDR5 – Increase memory bandwidth
  - CXL® 1.1 – Next Gen I/O for low latency and performance
- Optimized Power Mode\* enables energy savings with minimal impact of performance
- Advanced security technologies to help protect data with Intel Software Guard Extensions (Intel SGX) and Intel® Trust Domain Extensions (Intel TDX)
- To learn more about SKU transitions and refresh opportunities:
  - <https://xeonprocessoradvisor.intel.com>

## Intel® Xeon®, the processor designed for AI

### Total Cost of Ownership Savings

3<sup>rd</sup> Gen Intel® Xeon® to 5<sup>th</sup> Gen Intel Xeon

**AI**

(Recommender -DLRM)

up to **5:1**  
server consolidation<sup>1</sup>

up to **72%**  
TCO savings<sup>1</sup>

### Performance & Efficiency Gains

3<sup>rd</sup> Gen Intel® Xeon® to 5<sup>th</sup> Gen Intel Xeon

**AI**

(Real Time Inference)

up to **14x**  
higher performance<sup>2</sup>

up to **9.5x**  
higher performance/watt<sup>2</sup>

### Intel vs. AMD

**AI**

(Recommender - DLRM)

up to **2.8x**  
higher batched inference performance<sup>3</sup>

up to **2.6x**  
higher performance/watt<sup>3</sup>



**TCO Advisor Tool**

<sup>1,2,3</sup> See [T12, A16, A208] at [intel.com/processorclaims](https://www.intel.com/processorclaims): 5<sup>th</sup> Gen Intel Xeon Scalable processors. Results may vary

\*Enabled in platform BIOS



# Benefits to upgrading infrastructure configurations

## 3:1 - 5th Gen Xeon Server Consolidation:

Up to 3:1 consolidation and 61% TCO savings with 5th Gen Intel Xeon processors: Calculations as of June 2024 based on the Intel® Node TCO & Power Calculator using default cost, power and TCO assumptions over a 5-year TCO horizon comparing replacing 50 older servers with Intel Xeon 4214 processors with new servers using new Intel Xeon 5520+ processors. Results may vary.

[4214: <https://www.spec.org/cpu2017/results/res2020q2/cpu2017-20200427-22184.html>  
5520+: <https://www.spec.org/cpu2017/results/res2024q2/cpu2017-20240325-42544.html> ]

## 1st Gen Intel Xeon to 5th Gen Intel Xeon:

Calculations as of June, 2024 based on the Intel® Node TCO & Power Calculator using default cost, power and TCO assumptions over a 5-year TCO horizon comparing replacing 50 older servers with Intel Xeon 4110 processors with new servers using new 4th Gen Intel Xeon processors. Results may vary. Performance measurements based on published SPECrate®2017\_int\_base on spec.org as of March 28, 2023

8168: <https://www.spec.org/cpu2017/results/res2019q3/cpu2017-20190827-17261.html>  
8562Y+: <https://www.spec.org/cpu2017/results/res2024q1/cpu2017-20240311-42071.html>  
6138: <https://www.spec.org/cpu2017/results/res2021q4/cpu2017-20211122-30177.html>  
8558: <https://www.spec.org/cpu2017/results/res2024q2/cpu2017-20240325-42453.html>  
5118: <https://www.spec.org/cpu2017/results/res2019q3/cpu2017-20190805-16469.html>  
5520+: <https://www.spec.org/cpu2017/results/res2024q1/cpu2017-20240311-41976.html>  
4110: <https://www.spec.org/cpu2017/results/res2020q4/cpu2017-20201015-24218.html>

## 2nd Gen Intel Xeon to 5th Gen Intel Xeon:

Calculations as of June 2024 based on the Intel® Node TCO & Power Calculator using default cost, power and TCO assumptions over a multi-year TCO horizon comparing replacing 50 older servers with Intel Xeon 4110 processors with new servers using new 4th Gen Intel Xeon processors. Results may vary. Performance measurements based on published SPECrate®2017\_int\_base on spec.org as of March 28, 2023

8260: <https://www.spec.org/cpu2017/results/res2021q3/cpu2017-20210815-28685.html>  
8558: <https://www.spec.org/cpu2017/results/res2024q2/cpu2017-20240325-42453.html>  
6230: <https://www.spec.org/cpu2017/results/res2021q1/cpu2017-20210301-25037.html>  
5218: <https://www.spec.org/cpu2017/results/res2020q3/cpu2017-20200731-23587.html>  
6538Y+: <https://www.spec.org/cpu2017/results/res2024q1/cpu2017-20240115-40663.html>  
4214: <https://www.spec.org/cpu2017/results/res2021q3/cpu2017-20210824-28887.html>  
5520+: <https://www.spec.org/cpu2017/results/res2024q1/cpu2017-20240311-41976.html>