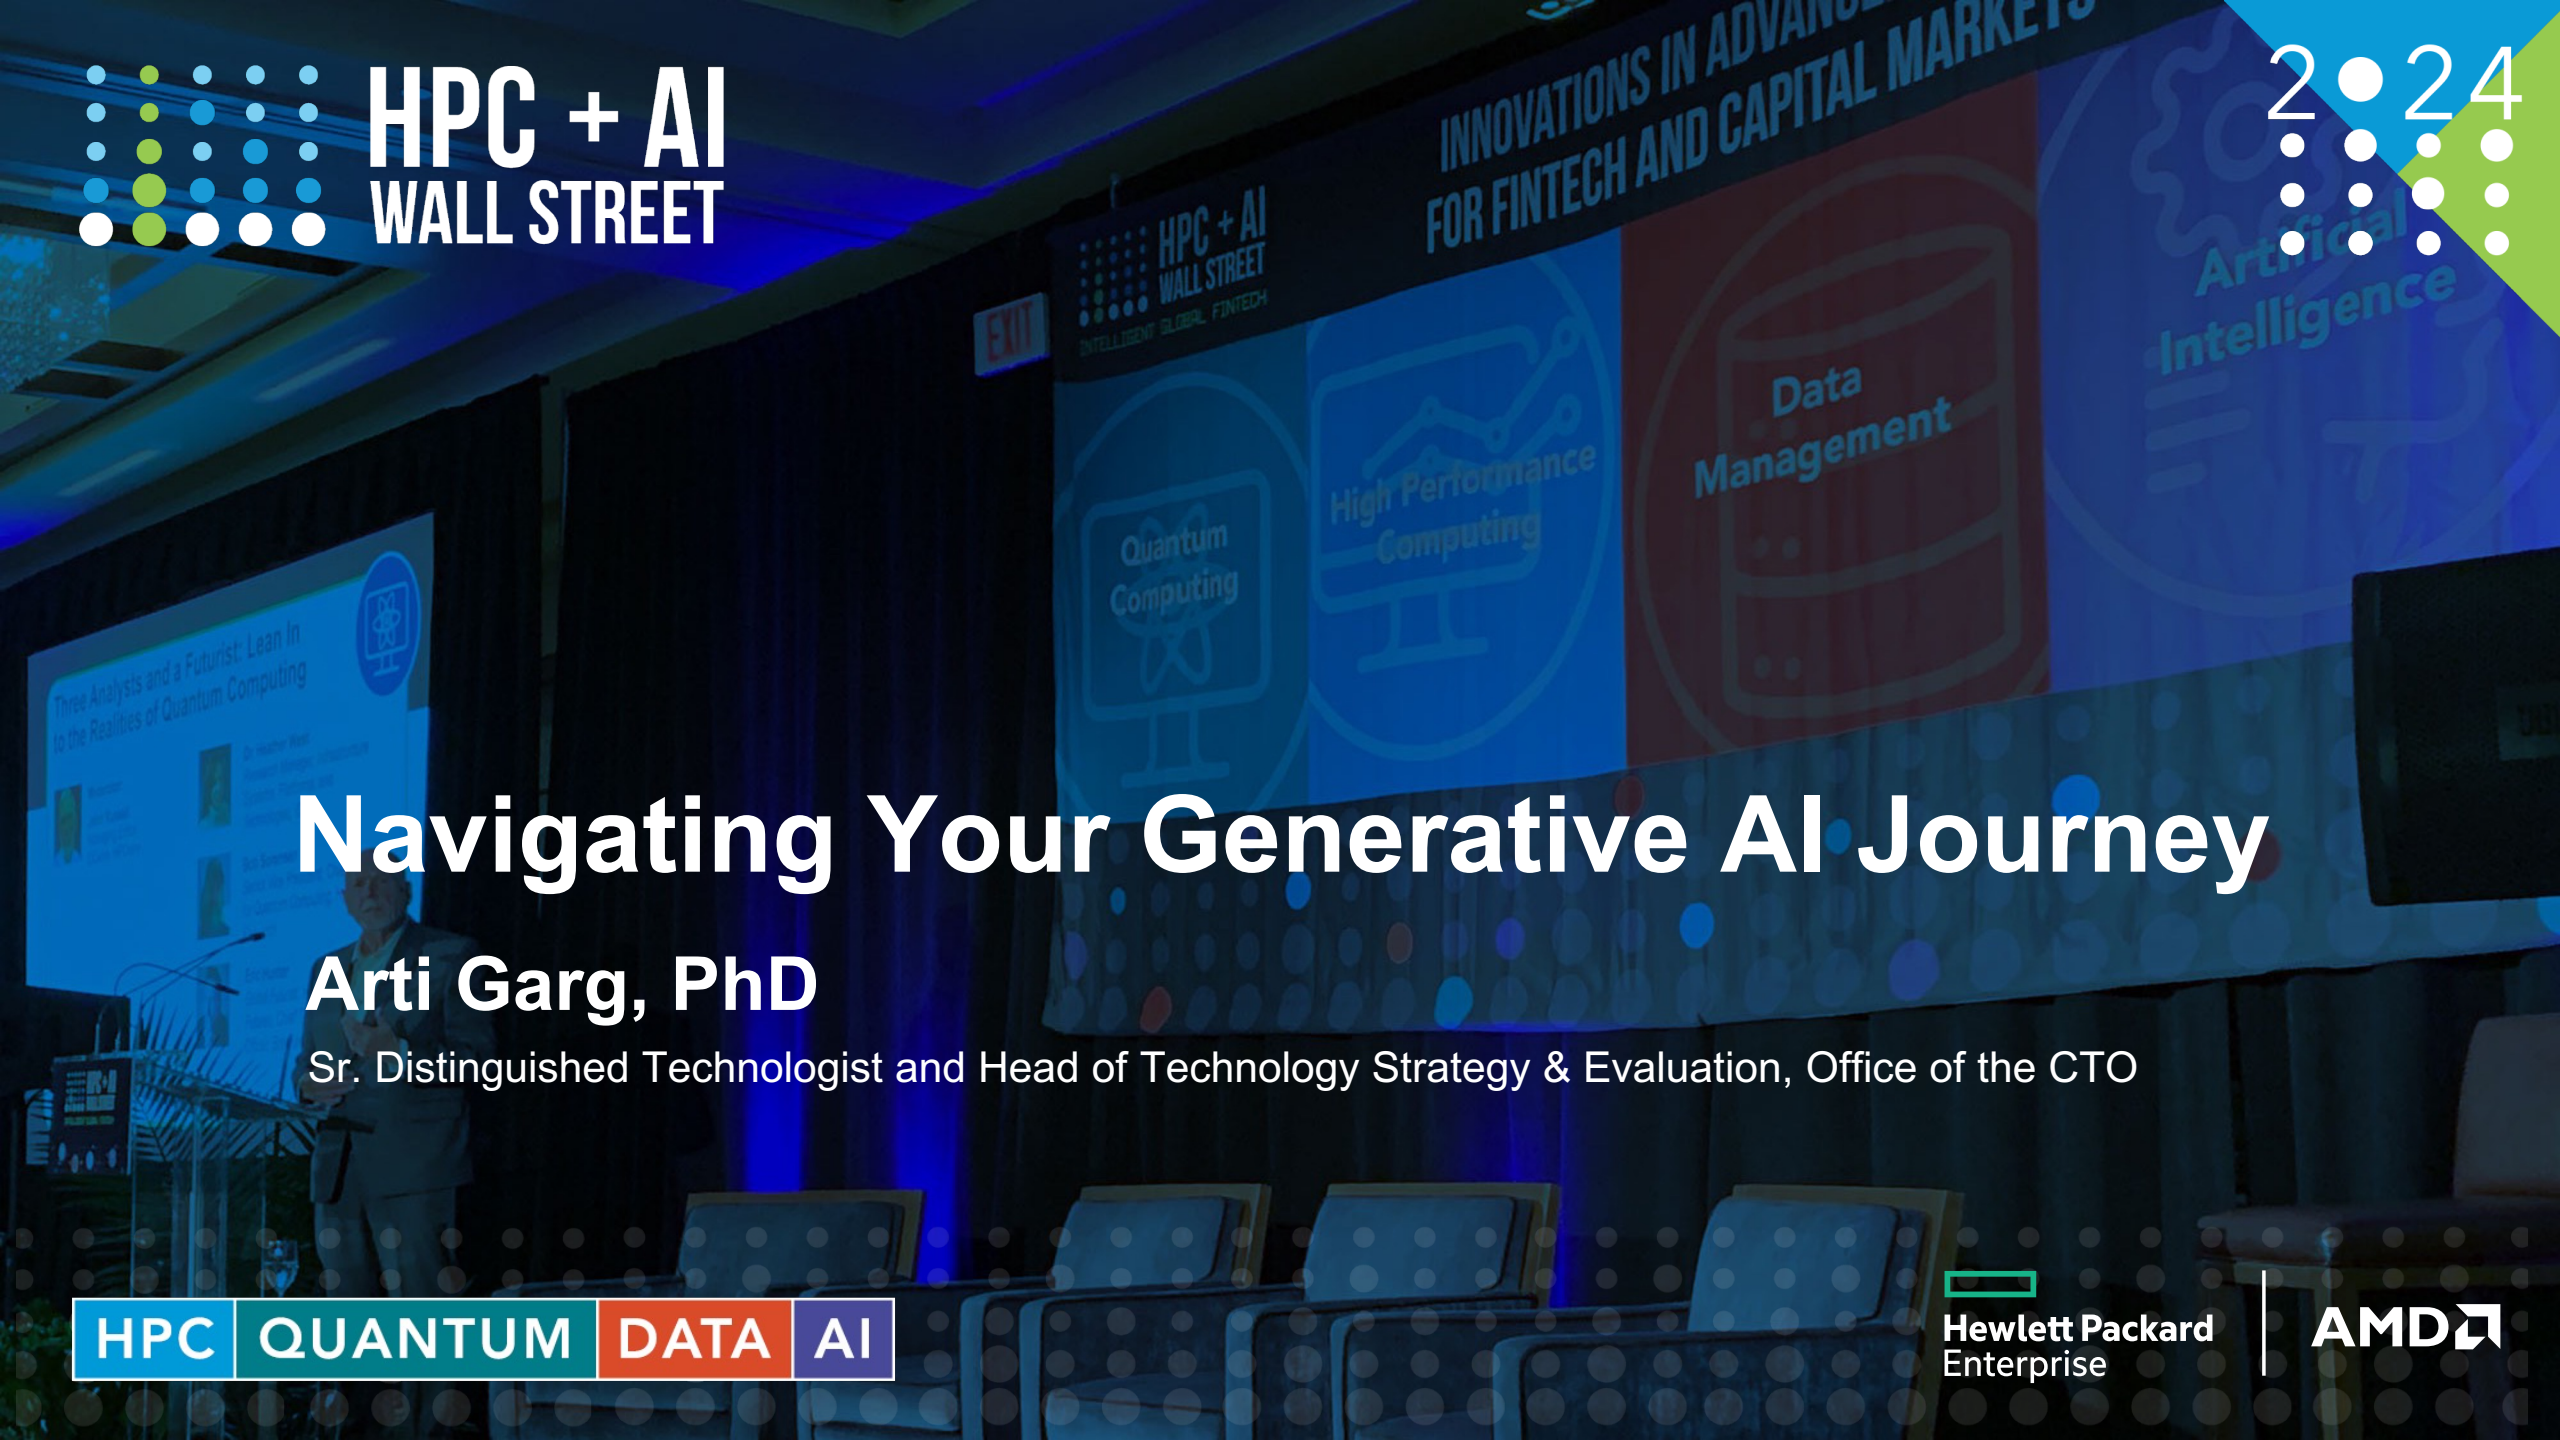

Navigating Your Generative AI Journey

Arti Garg, PhD

Sr. Distinguished Technologist and Head of Technology Strategy & Evaluation, Office of the CTO

HPC | QUANTUM | DATA | AI

Hewlett Packard Enterprise | AMD

HPC + AI WALL STREET

At HPE, we believe in changing the way people live and work.

The World Before November 2022…

HPC + AI WALL STREET

Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, the World's Largest and Most Powerful Generative Language Model

Oct 11, 2021 +9 Like Discuss (1)

By Paresh Kharya and Ali Alvi

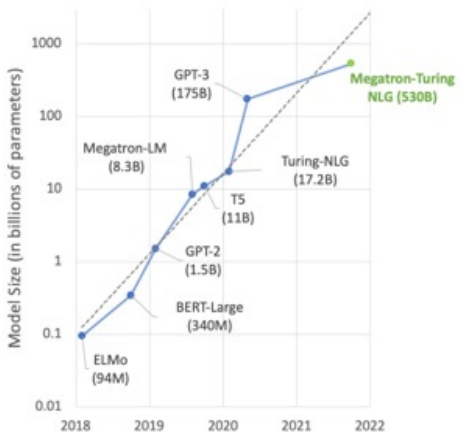

Top 9 Use Cases of Computer Vision in Manufacturing

GPT-3 powers the next generation of apps

Over 300 applications are delivering GPT-3–powered search, conversation, text completion, and other advanced AI features through our API.

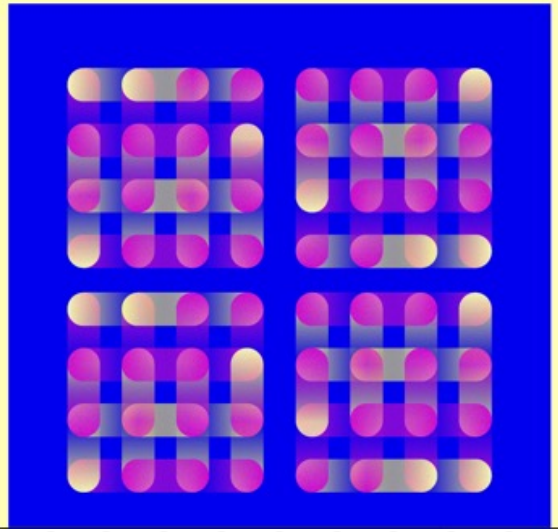

…and Then This Happened…

HPC + AI WALL STREET

- OpenAI's chatGPT
- OpenAI's GPT-4
- Google's Bard
- Anthropic's Claude 2
- Google's Gemini
- AWS's Q
- …many more

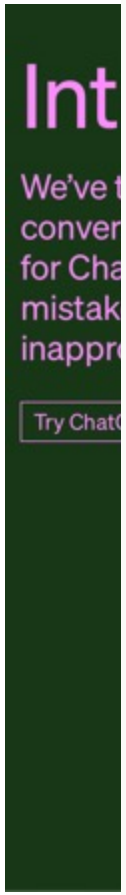

The AI Arms Race

Siddharth Sharma · Follow
9 min read · Feb 9

2

> *"It's a new day in search. The race starts today… We're going to move fast."*
> *— Satya Nadella, Microsoft CEO*

Altman and Lecun (Forbes)

In recent years, the field of Artificial Intelligence has seen a rapid rise in the development of large language models. These models, based on deep

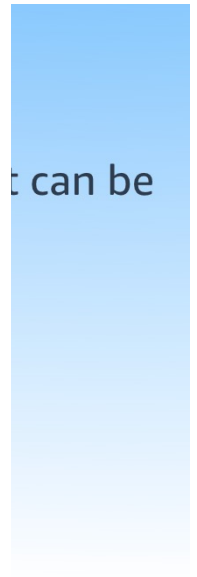

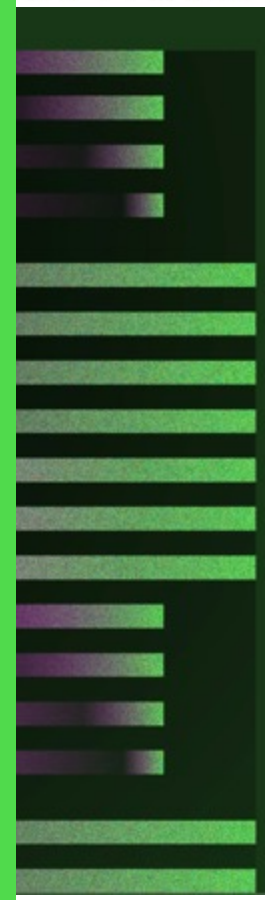

…and Everything Changed

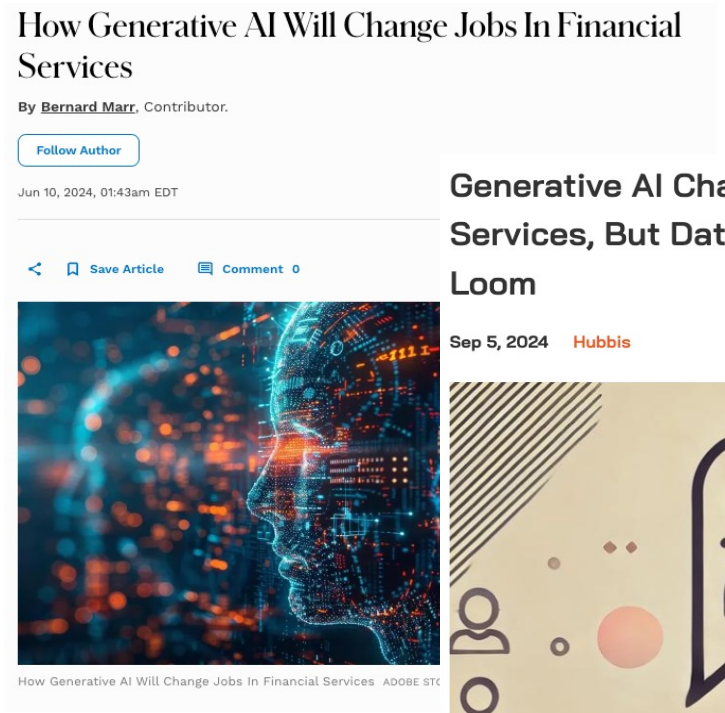

How Generative AI Will Change Jobs In Financial Services

By Bernard Marr, Contributor.

Follow Author

Jun 10, 2024, 01:43am EDT

Save Article Comment 0

How Generative AI Will Change Jobs In Financial Services ADOBE STO

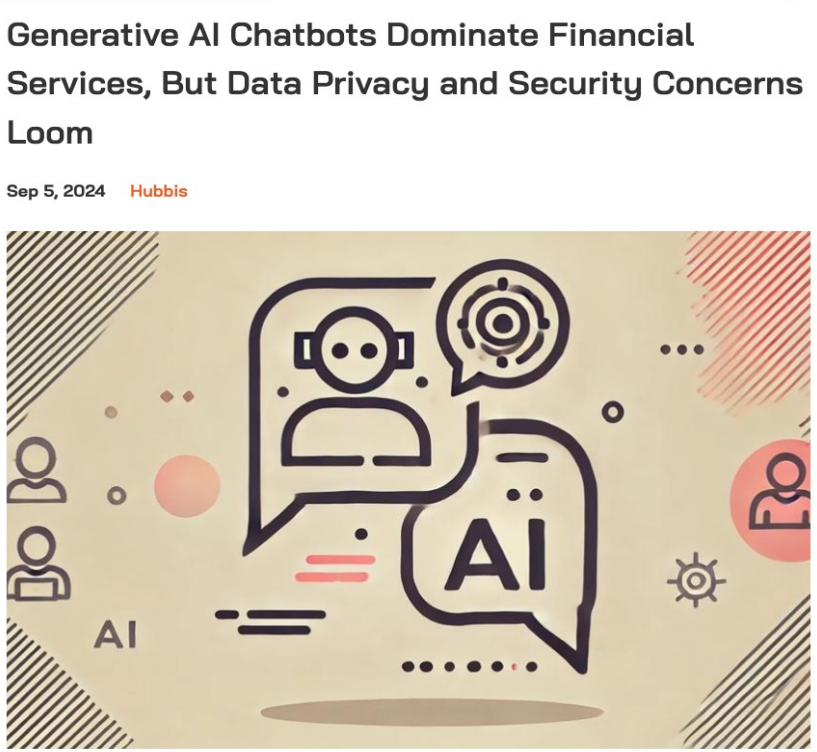

Generative AI Chatbots Dominate Financial Services, But Data Privacy and Security Concerns Loom

Sep 5, 2024 Hubbis

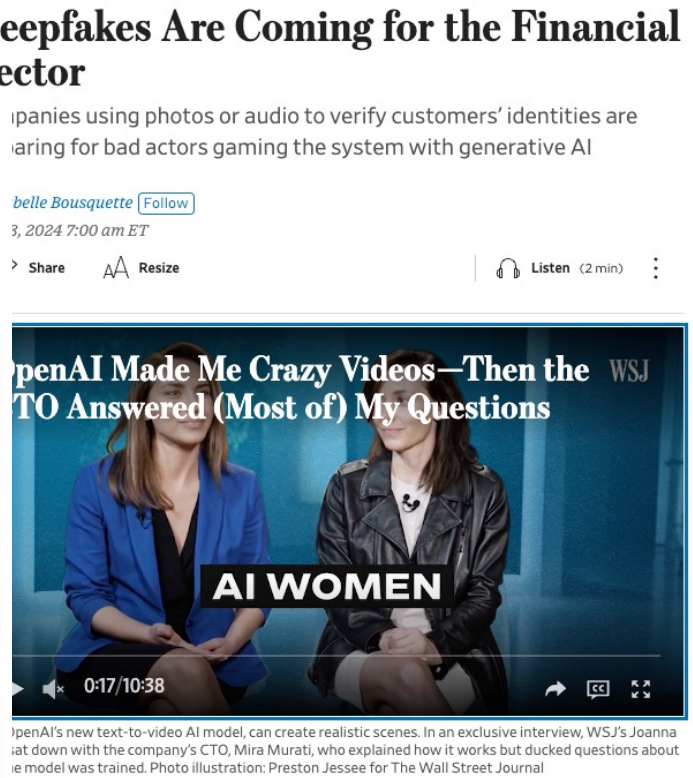

Deepfakes Are Coming for the Financial Sector

…panies using photos or audio to verify customers' identities are …aring for bad actors gaming the system with generative AI

…belle Bousquette Follow

…3, 2024 7:00 am ET

Share Resize Listen (2 min)

…penAI's new text-to-video AI model, can create realistic scenes. In an exclusive interview, WSJ's Joanna …at down with the company's CTO, Mira Murati, who explained how it works but ducked questions about …e model was trained. Photo illustration: Preston Jessee for The Wall Street Journal

Where Do You Go From Here?

1. What is needed to **get started ?**
2. Which **one** (company, model, open source)?
3. Is it **enterprise ready**?
4. What about **governance**?
5. Which **use cases**?

What Is Required?

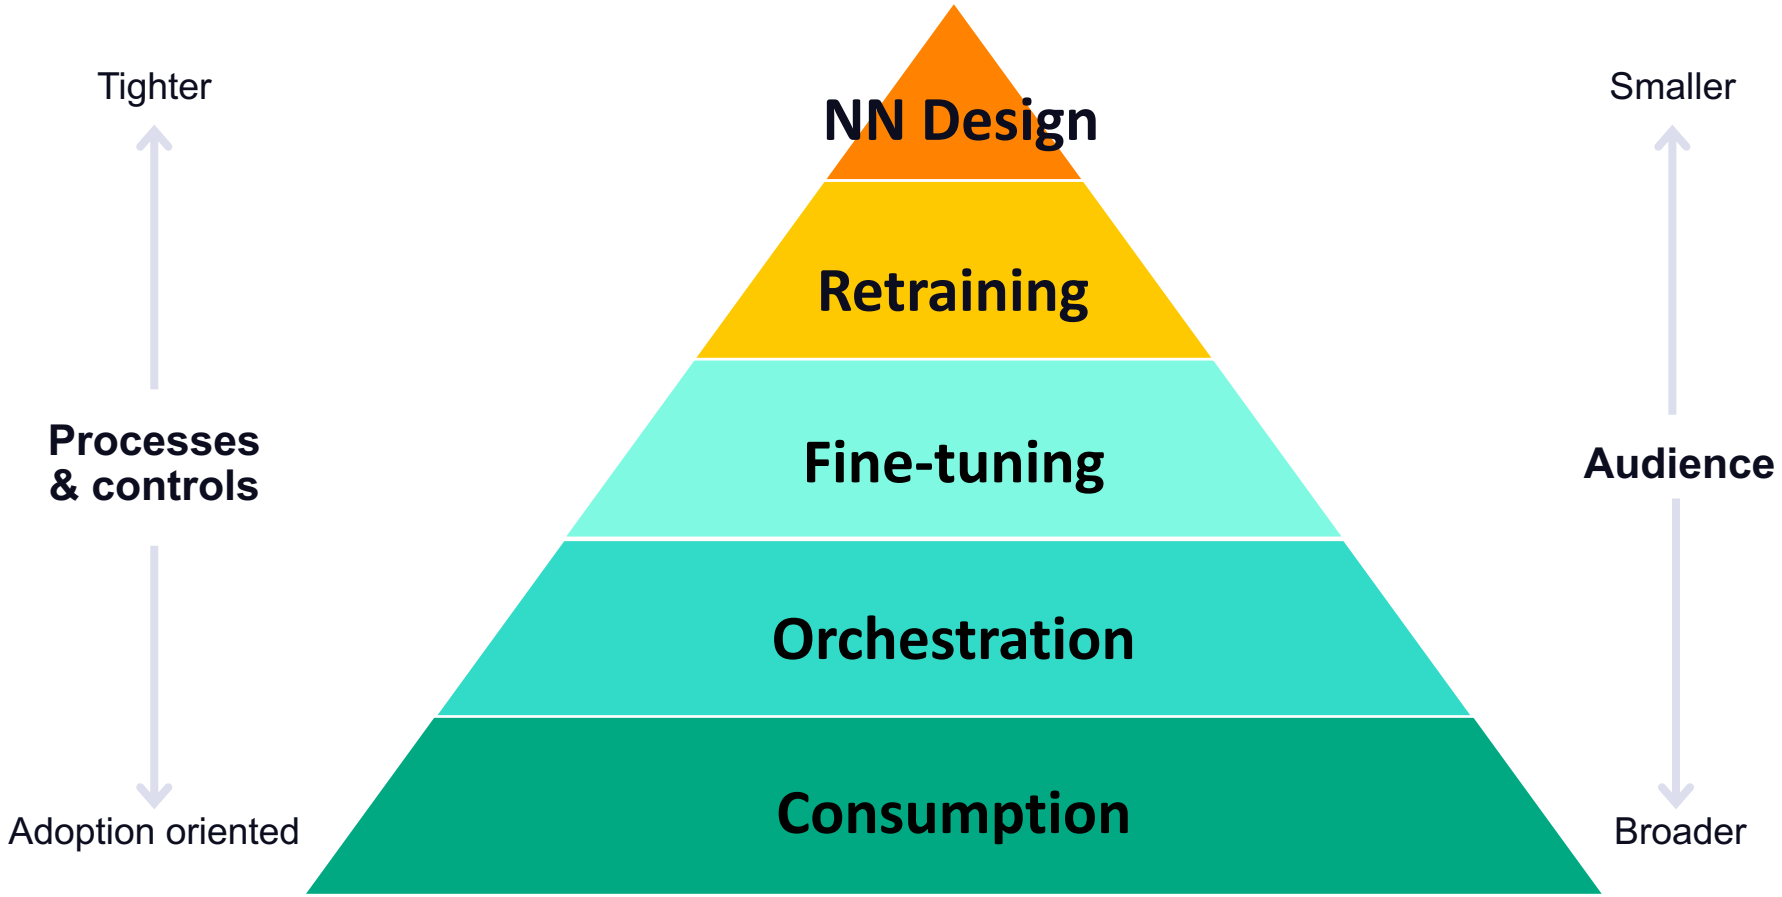

What Is Required? I Depends On Where You Want To Play

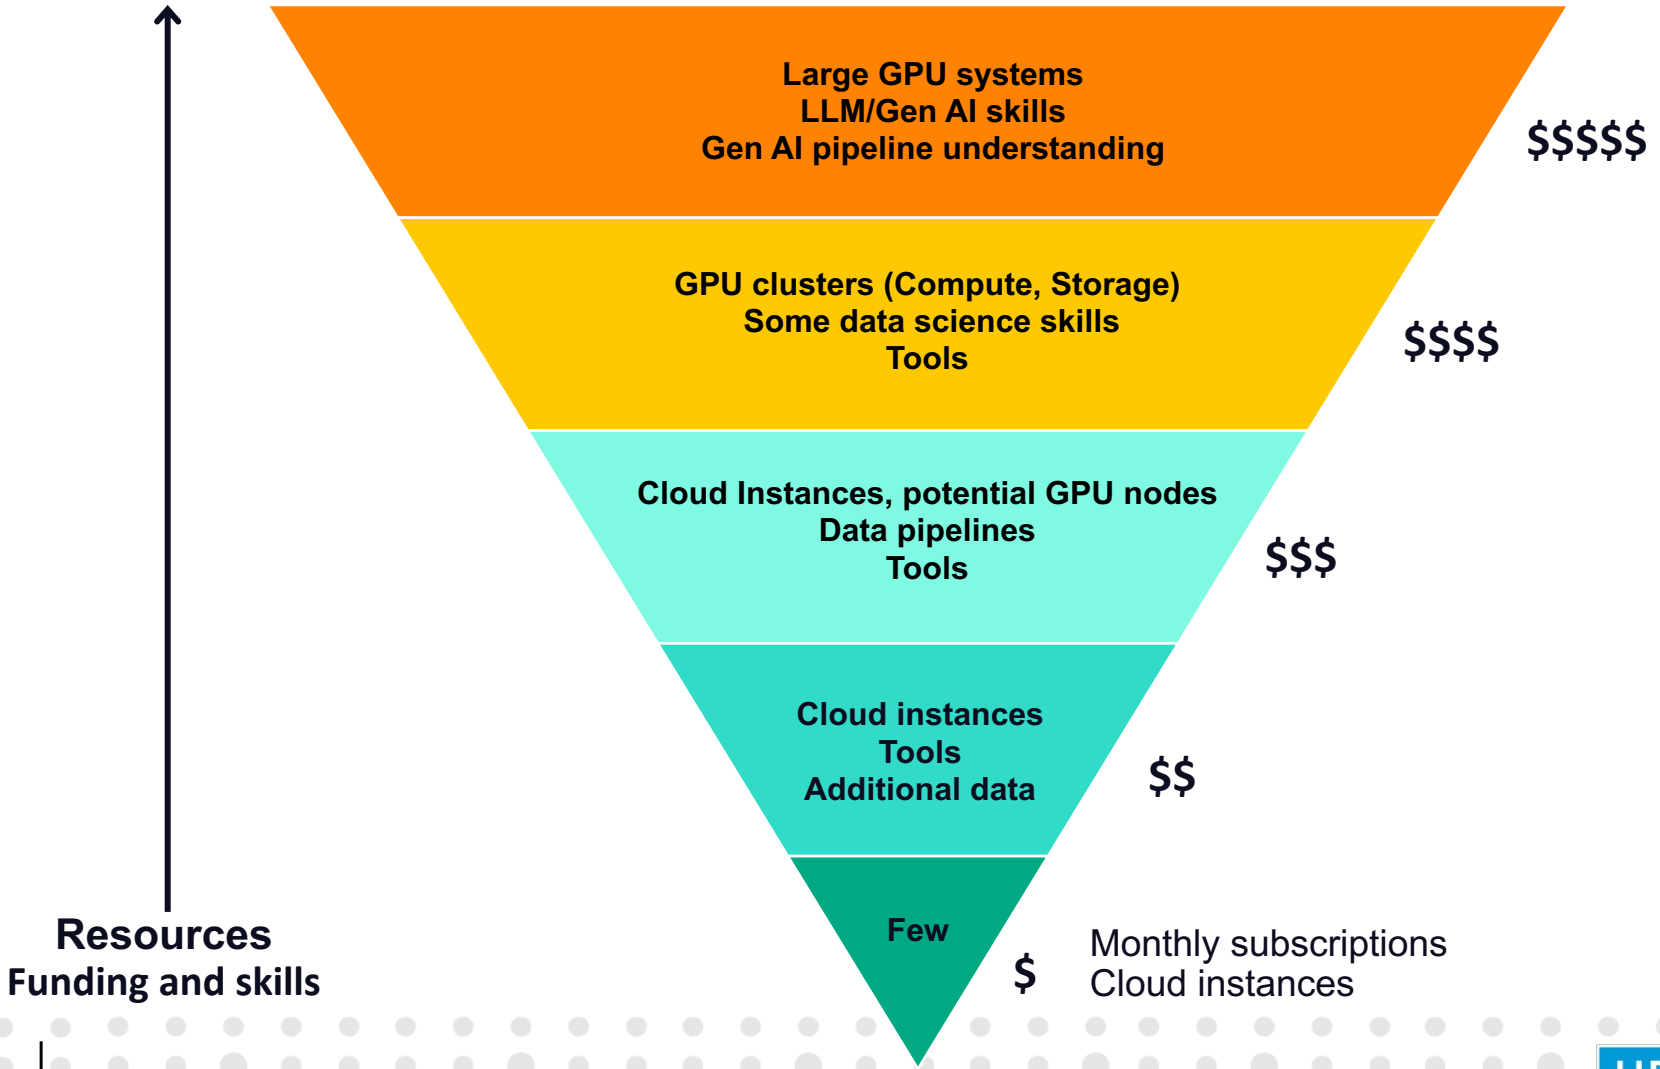

Artificial Intelligence (Gen AI and Beyond)

HPC + AI WALL STREET

Models and Services

Open-source models

Commercial partnerships

Services to better drive AI models

+

AI Platform

Data management

Model training

Model inference

+

Infrastructure

From bare metal to containers

From training to inference

From supercomputer to edge

One platform | Vendor neutral | Cloud neutral | AI accessible for all

What Are Your Key Considerations?

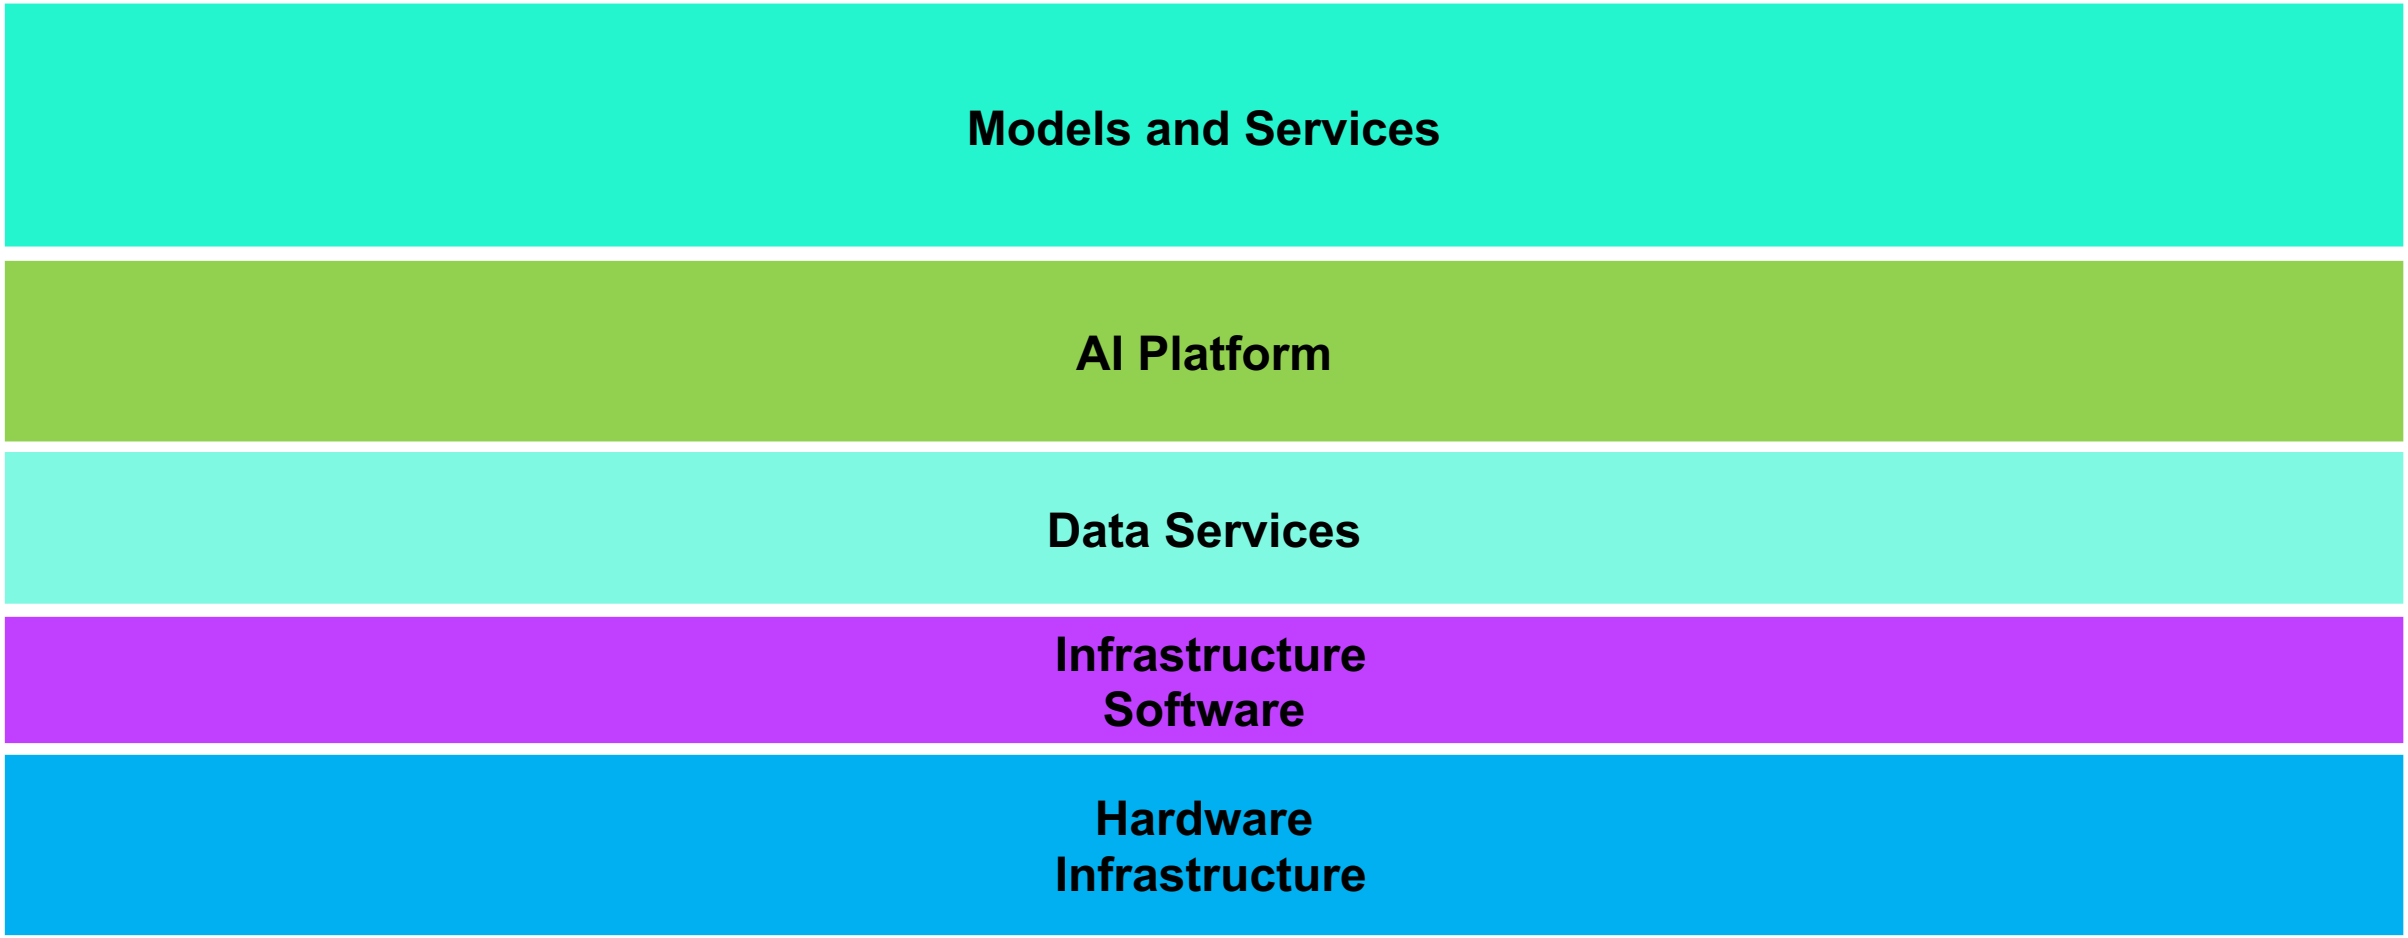

Infrastructure | Good models require good data and also a lot of compute

Completion	Base—64 GPUs for 37 days—56,832 GPU hours
make sure the economy is strong.	

Completion	Extended—256 GPUs for 22 days—135,168 GPU hours
declare war on the dog.	

Infrastructure I Where to Deploy?

	Pros	Cons
Private Cloud	• Enhanced Security & Privacy • Customization • Predictable performance • Compliance adherence • Reduced Risk of vector lock in	• Higher Initial Investment • Limited Scalability • Resource underutilization • Complexity
Public Cloud	• Cost Efficiency • Scalability • Global Accessibility • Maintenance and Updates • Innovation	• Security and Privacy Concerns • Dependency • Limited Control • Potential Compliance Issues • Vendor Lock-In

Platform | A Very Complicated Landscape

HPC + AI WALL STREET

Data

Pipelines
Pachyderm

Processing
RAY
dask
RAPIDS
Pachyderm

Versioning
LFS
DVC
DELTA LAKE
Pachyderm

Labelling
scale
Labelbox

Data Sources
LOG
CSV
aws
Azure
SQL
snowflake
CLOUDERA
databricks
lustre
MAPR
S3

Development/Training

Foundation Models
GitHub
arXiv

Packages
HOROVOD

IDE *
jupyter
PC

Evaluation
TensorBoard
gradio

Experiments
Determined AI
mlflow

Scheduling
kubernetes
slurm

Compute
SHASTA
intel
NVIDIA
AMD
habana An Intel Company
cerebras
Qualcomm

Deployment/Inference

Logs, Metrics
Prometheus
Grafana

Optimization
tvm
TensorRT

Testing
great_expectations
Buildkite

Trustworthy AI
ALIBI EXPLAIN
truera
modzy
fiddler
ALIBI DETECT
TROJ.AI

Serving (Central)
KServe
BENTOML
SELDON

Serving (Edge)
Triton Inference Server
Custom Docker

Runtime Environment
TensorFlow Extended
ONNX
TORCHSERVE

Models | How Do You Choose the Right One? It Depends

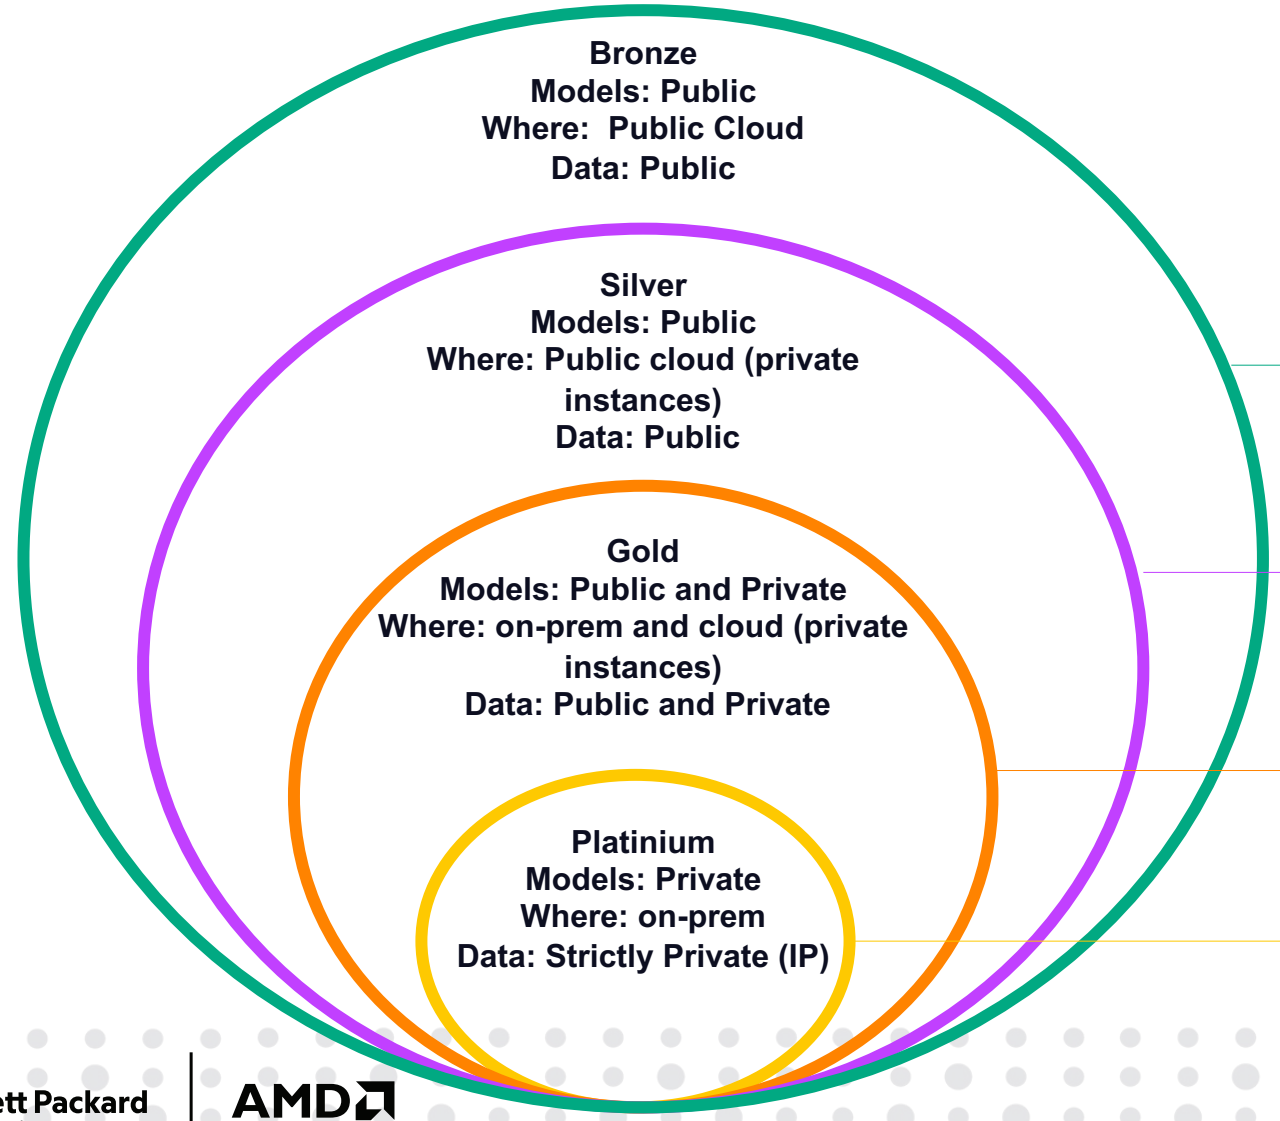

Enterprises need hybrid environments to support multiple safety, control, and privacy requirements.

ChatGPT like models used, either direct access or via APIs.

Models fine tuned in the public cloud, public models API end points used in private instances, own models built

Models trained on prem.
Models fine-tuned in the public cloud or on prem.
Inference running on prem and public cloud

Models trainined and running on prem

Models | Choosing the Right Foundation Model

HPC + AI WALL STREET

- Select the best AI foundation model for your use case
- Evaluation categories:
  - Project Requirements
  - Model Capabilities
  - Operational & Ethical Considerations
  - Ethical & Governance Considerations
- Align model choice with project goals and organizational policy

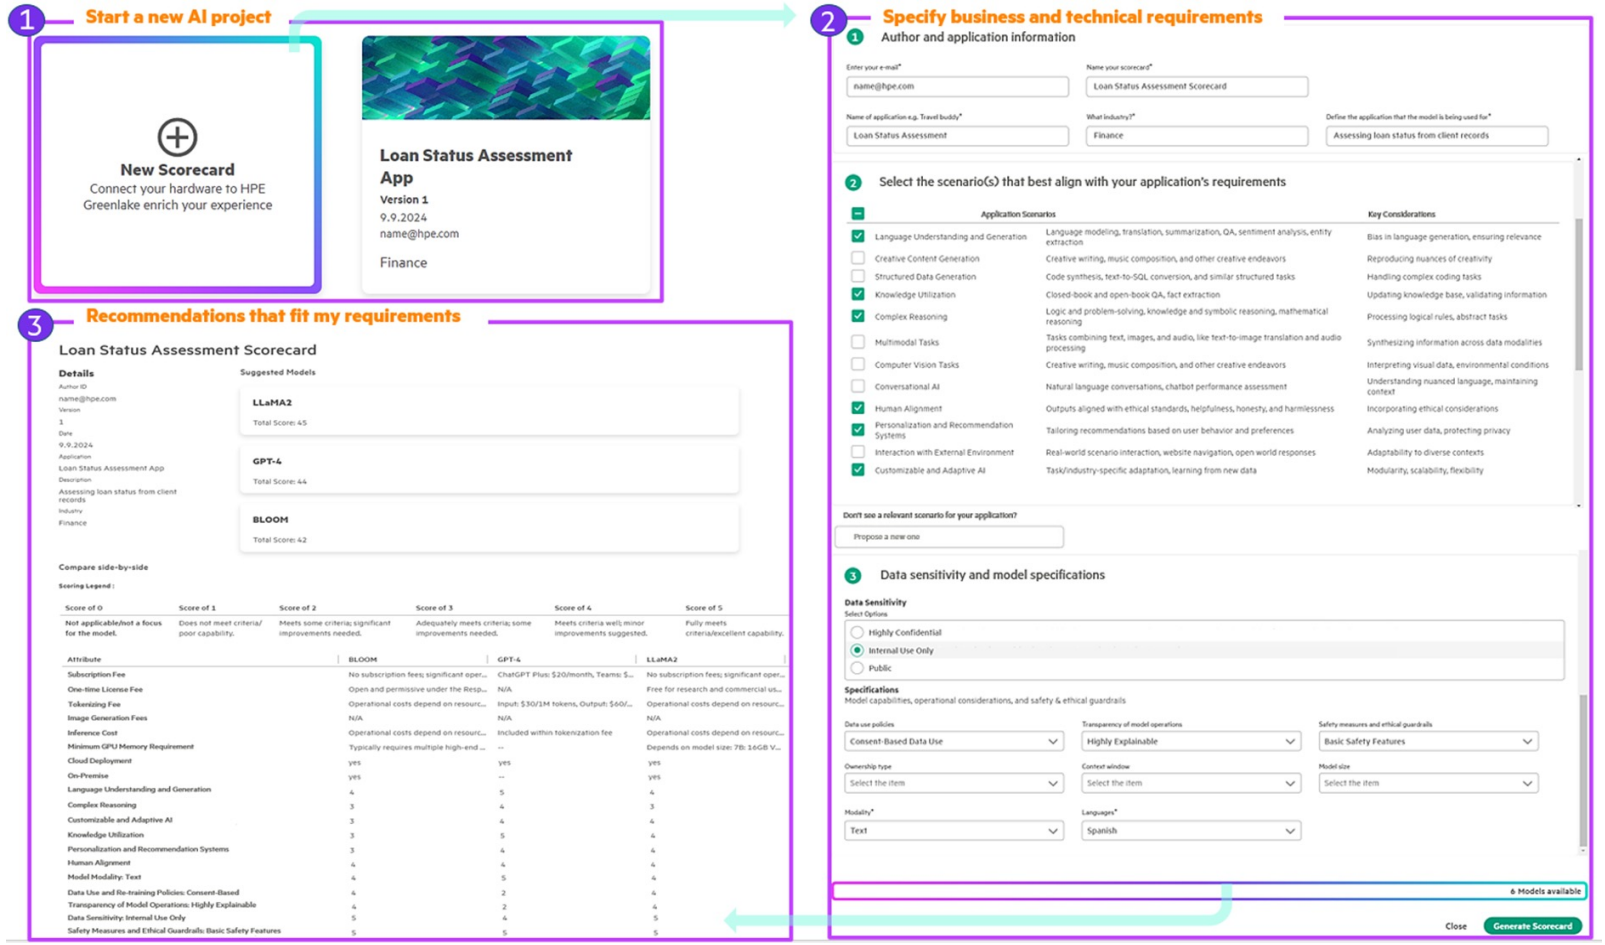

Hewlett Packard Enterprise | AMD

HPC | QUANTUM | DATA | AI

Implementation | What About Security?

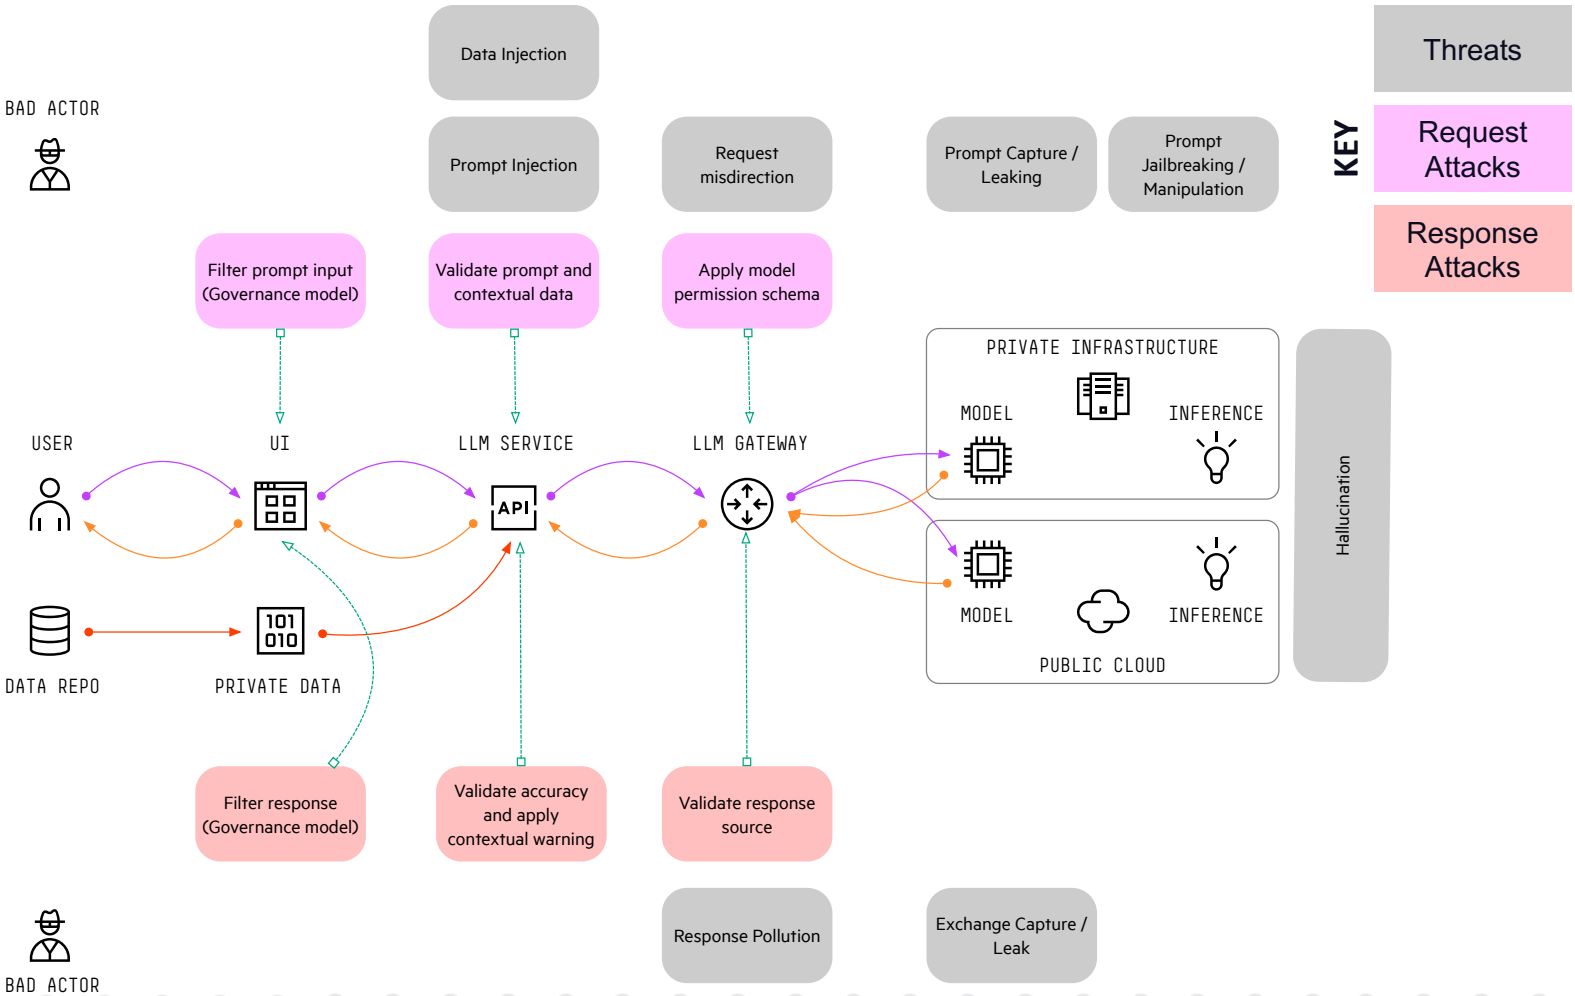

- Fail gracefully and secretly
- Audit user authorization
- Secure access to external resources
- Parameterize and validate all inputs and outputs
- Avoid persisting changes when possible
- Adversarial testing

Implementation | Governance

HPC + AI WALL STREET

- Allow employees to **securely and auditably** use new technology to assist with researching, writing, and code generation.
- Keep your **sensitive data on premise.**
- Deliver the ability to **choose your model.**
- Enforce **company policy**.
- *KEY DELIVERABLES*
  - *AI guardrails*
  - *Model Token Management*
  - *Access to models either on-prem or public cloud*

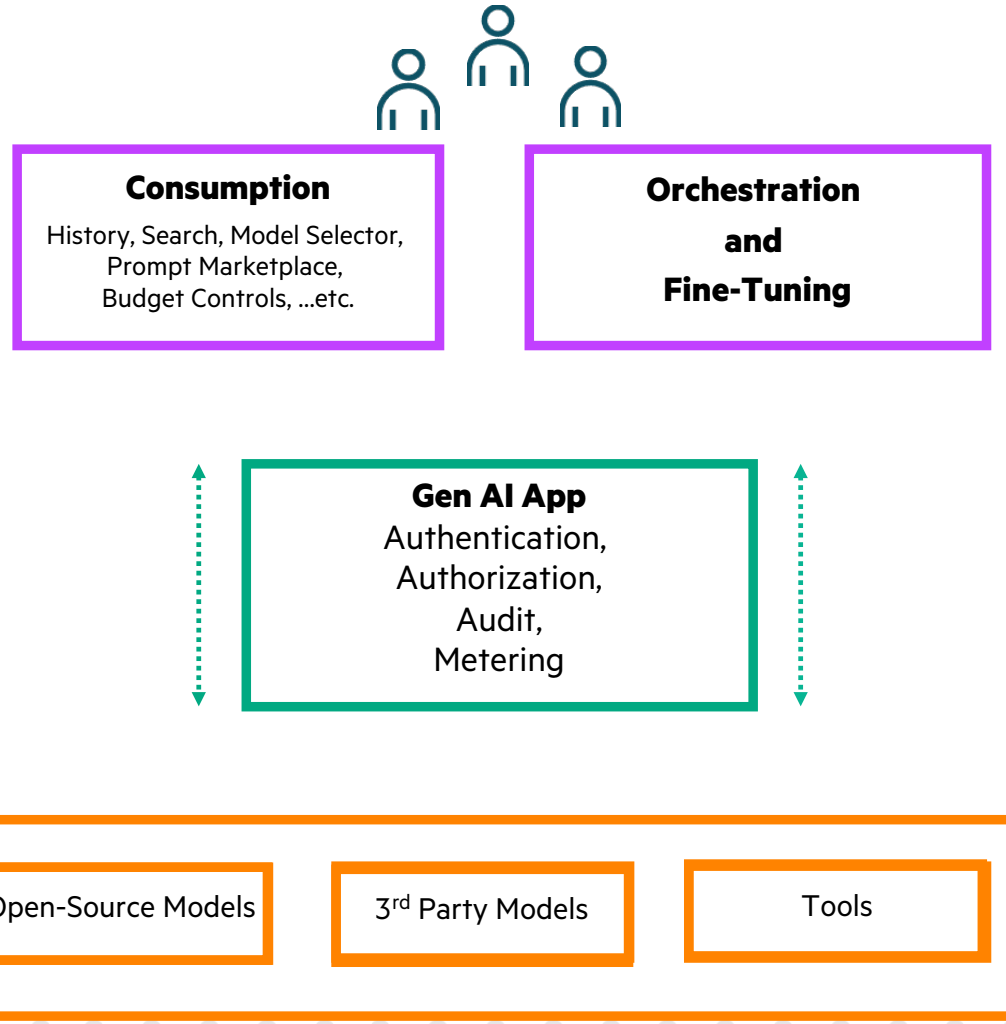

What Else? | Observability | Sustainability | FinOps

HPC + AI WALL STREET

Clouds

Datacenter

GPU

Containers

Virtual Machines

Bare metal apps

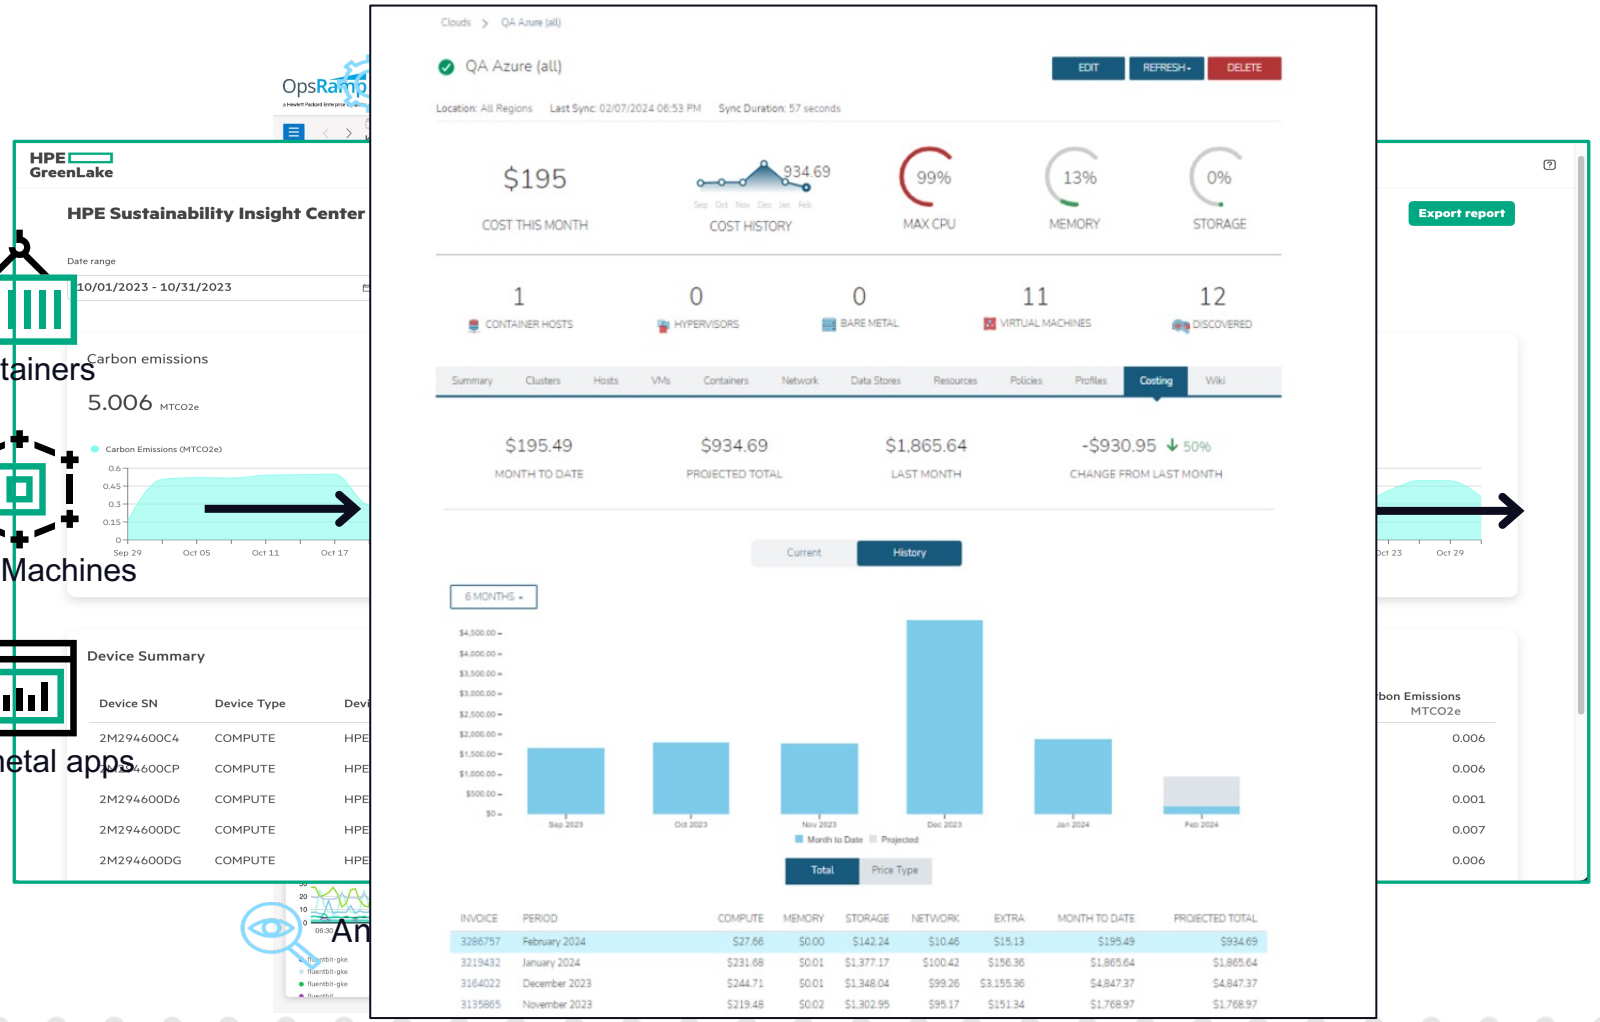

Dashboards

Cost Savings

Automate

Scale up / down

AI Strategy on One Page

Models and Services	Accelerate the journey to Gen AI value for the Enterprise	**Model Hub** **HPE A&PS Services**	**Model-aaS**
AI Platform	Drive collaborative AI model development from Data preparation to Deployment	Machine Learning Development Environment Machine Learning Data Management HPE AI Essentials	**AI PaaS**
Data Services	Manage data governance, availability, and security	Data Fabric Data Observability Machine Learning Data Management	
Infrastructure Software	Allow AI models to run efficiently anywhere	Private Cloud for BMaaS, VMaaS, and CaaS.	**IaaS**
Hardware Infrastructure	Optimize the performance of AI workloads with AI native infrastructure	Accelerators: Heterogenous support for Nvidia/AMD/Intel Compute: Cray, Proliant Storage: GL4FL, Networking: Slingshot, Mellanox Best in Class DLC	

Observability
Sustainability
FinOps

AI Use Case Families in Financial Services: Spanning Front, Middle, and Back Office

HPC + AI WALL STREET

FRONT OFFICE

MIDDLE OFFICE

BACK OFFICE

Credit Scoring	Customer Acquisition	Customer Experience	Customer Retention
Trading	Process Automation	Knowledge Management	Risk Management
Agent-Based Modelling	Stress Testing	Regulatory Compliance	Fraud Detection
Synthetic Data	Data Architecture	Data Capabilities	Infrastructure Optimization

Getting Started: Pick Your Use Case

HPC + AI WALL STREET

	Regulatory Compliance	Credit Scoring
End-User Group	Internal Users Employee domain SME's (CPA's, compliance officers, Accountants, auditors, (CAO/CFO) Small Userbase, Highly Specialized Able to provide feedback on early iterations Low exposure to bad actors	External and/or Internal Users Customers or customer-facing employees (Loan officers, salespeople, brokers) Large Userbase, Experience levels may vary High exposure to bad actors
Use Case Regulations	Regulations mandate what to report to who and when, with specific instructions on preparation methodology "Automate as much as possible to minimize human error" -regulators	Fair lending, anti-discrimination, industry regs High-risk Use Case- AI Act (EU) GDPR/CCPA/Data protection regulations "AI can't do anything a human wouldn't be allowed to do" -regulators
Source Data (Unstructured)	Public information (Regulatory text, instructions, etc.) Trusted source- directly from regulators Low potential for wrongdoing if leaked	Non-public, Personally Identifiable Information Various sources- trustworthiness may vary High potential for wrongdoing if leaked Data sovereignty & data gravity concerns
Models Used	Open-source foundation model (Llama3 based) Out of box, no customization	Proprietary vendor or internally developed models Highly customized
Hardware Infrastructure	Small-scale inference environment or cloud API solution	Private, hosted, secure AI Inference at Scale cluster AI Factory training resources if building your own models

Development	Production
Security Guardrails Governance	Observability Sustainability FinOps

Internal Case Study: Hewlett-Packard International Bank DAC

- Hewlett Packard International Bank DAC, (“HPIB”) is a credit institution authorized by the Central Bank of Ireland and owned by HPE.
- The Capital Requirements Regulation 3 (CRR3) is a new regulation spanning 646 pages that goes into effect for banks in the EU on January 1, 2025. It is part of the wider “Basel 3.1”/ “Basel 3 Endgame” regulatory risk framework for global financial institutions.
- The CRR3 requirements have been consistently modified leading up to implementation, and portions still lacked fully defined requirements with less than 1 year to go.
- Risk jargon often requires a CPA with specific experience in the financial sector to understand these standards.

Lengthy and complex risk regulations

+

Highly technical subject matter

=

Complexity, Confusion, Delays, Increased risk

Problem Statement

HPC + AI WALL STREET

- The financial services industry complies with a complex and fast-changing web of regulatory requirements.
- International firms are forced to adapt to the pace of change in the standards and laws of multiple regulators, jurisdictions, and governments
- Failure to comply can lead to civil (and even criminal) liability. Fines and sanctions can amount to **billions of dollars** for large institutions
- Each regulation or report change has the potential to have no impact, low impact, or very high impact and a long adoption process. Expert interpretation is required to determine “Does this change affect us?” and “What do we need to do differently than before?

Initial Experiments

HPC + AI WALL STREET

Experiment 1

Bulk analysis of documents on regulator website via Python to produce categorized CRR3 Implementation Report

Experiment 2

Analysis triggered by frequent regulator website updates, which can provide automated guidance and alerts for impactful changes and opportunities to comment

Experiment 3

Chatbot experience for real-time self-service with follow-up questions

Key Benefits

HPC + AI WALL STREET

Transparency: End-to-end tracking of new proposed regulatory changes, opening and closing of public comment periods, adoption of final changes, and phased implementation deadlines.

Operational efficiency: Swift document analysis and reporting. Give users opportunity to ask questions on complex subject matter in natural language and receive proactive updates on what they need to know.

Cost savings: Minimize burden of adapting to new legislation, manual reviews and potential fines for non-compliance, and consulting spend.

Increase collaboration: Create wider awareness of upcoming changes.

Part of a Holistic Compliance Strategy

HPC + AI WALL STREET

Generative AI analysis is one part of a comprehensive strategy

AI-Powered Analysis:

- Swift document review and extraction of essential data.
- Predictive modeling for potential compliance risks.
- Real-time notifications on regulatory changes.
- Simplification of complex regulatory texts using NLP.

Training & Development:

- Regular workshops to educate staff on compliance matters.
- Role-specific training modules for targeted learning.
- Case studies and simulations for hands-on experience.

Human Oversight:

- Expert panels for manual review and validation.
- Routine audits to ensure AI-generated results align with compliance requirements.
- Feedback loops for continuous AI model refinement.

HPE Has Been Preparing for Some Time

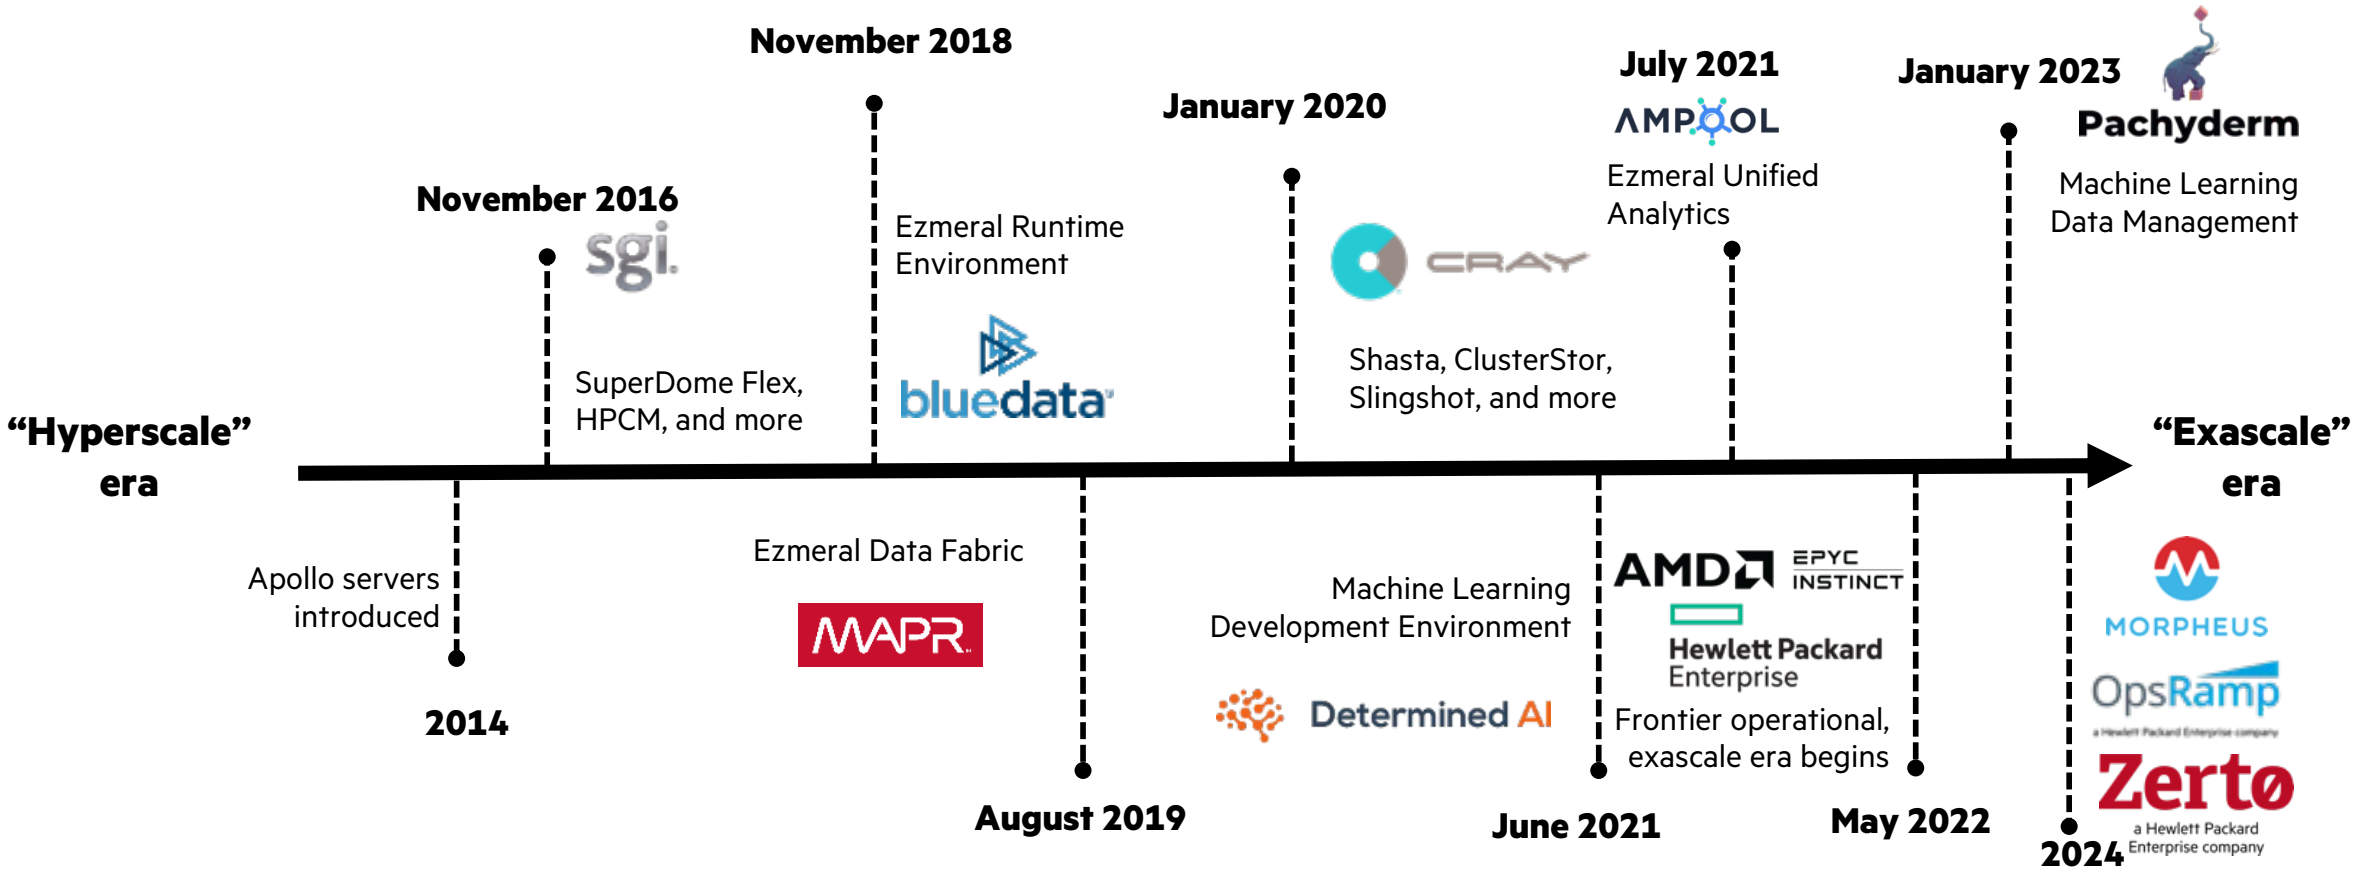

Enabling Artificial Intelligence

HPC + AI WALL STREET

HPE GreenLake

Models and Services

Open-source models

Commercial partnerships

Services to better drive AI models

+

AI Platform

Data management

Model training

Model inference

+

Infrastructure

From bare metal to containers

From training to inference

From supercomputer to edge

One platform | Vendor neutral | Cloud neutral | AI accessible for all


HPC + AI
WALL STREET
2024
Thank You
HPC QUANTUM DATA AI
Hewlett Packard Enterprise
AMD