

# How to Build a Sovereign AI Environment

accenture



Red Hat

# Accenture Innovation Hubs with Red Hat Technology that are open for onsite Visits



**Large number of implemented Red Hat use cases on data sovereignty and AI, e.g.,**

- Red Hat for Sovereign Cloud technology blueprint
- VMware to RH OpenShift Virtualization: a migration strategy
- Ansible Automation for network and security
- Generative AI to reinvent and improve IT operations
- Operational Resilience with Red Hat technologies (Antifragility)
- Red Hat Openshift AI to support Private AI scenarios
- Confidential AI
- ACIC Private AI Platform
- Confidential Containers on OpenShift on Baremetal



**The Lab aims to solve the most complex issues for governments and citizens using data & (Gen)AI in a secure and responsible way**

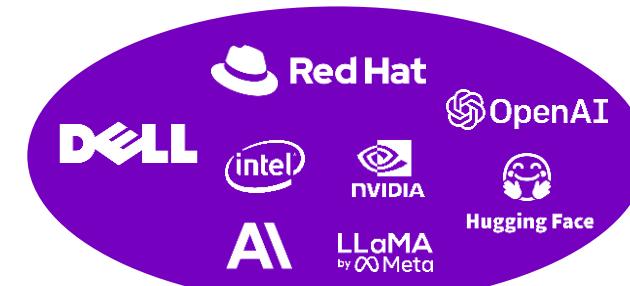
**190+** curated public & commercial datasets

**80+** Demo for rapid AI & GenAI prototyping

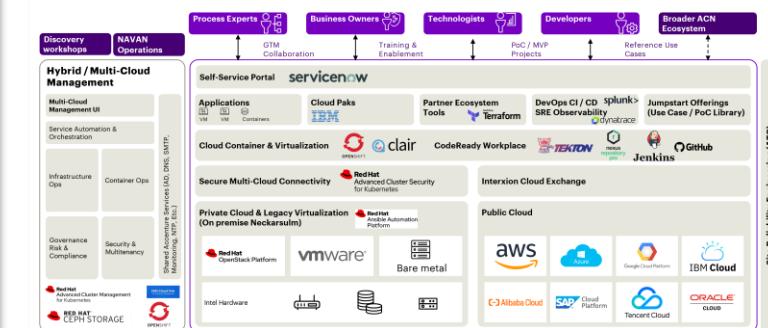
**30+** Client workshops in 2025

**15+** Project delivery in 2025

Harnessing cutting-edge technology through strategic partnerships



**Red Hat reference architecture demo lab within the STACKIT datacenter.**



**STACKIT**  
A Brand of Schwarz Digits

# Cornerstones for Sovereign AI Environments

Self-hosted large language models (LLMs) offer enhanced control, privacy, and customization for specific use-cases and data security requirements



## Model & Data Sovereignty

- Data must not be processed outside the company due to compliance or strategy
- Ensure model exclusivity to avoid unintentional knowledge transfer via shared models
- Where AI models are uncooperating critical IP, a zero-trust approach can ensure long-term success



## Customization and Control

- Full transparency of models and data are supporting transparent decision-making processes and enhance reporting to comply with regulations
- Customization with model fine-tuning and FM training are key to generate the most value and having the best results



## Pricing

- SaaS providers usually price their services in amount of business transactions, amount of data transferred, and amount of data stored
- The cost of the current LLM solutions can add up quickly, with climbing consumption
- Private AI and an open-source approach can be very cost effective



## Mitigate vendor lock-in

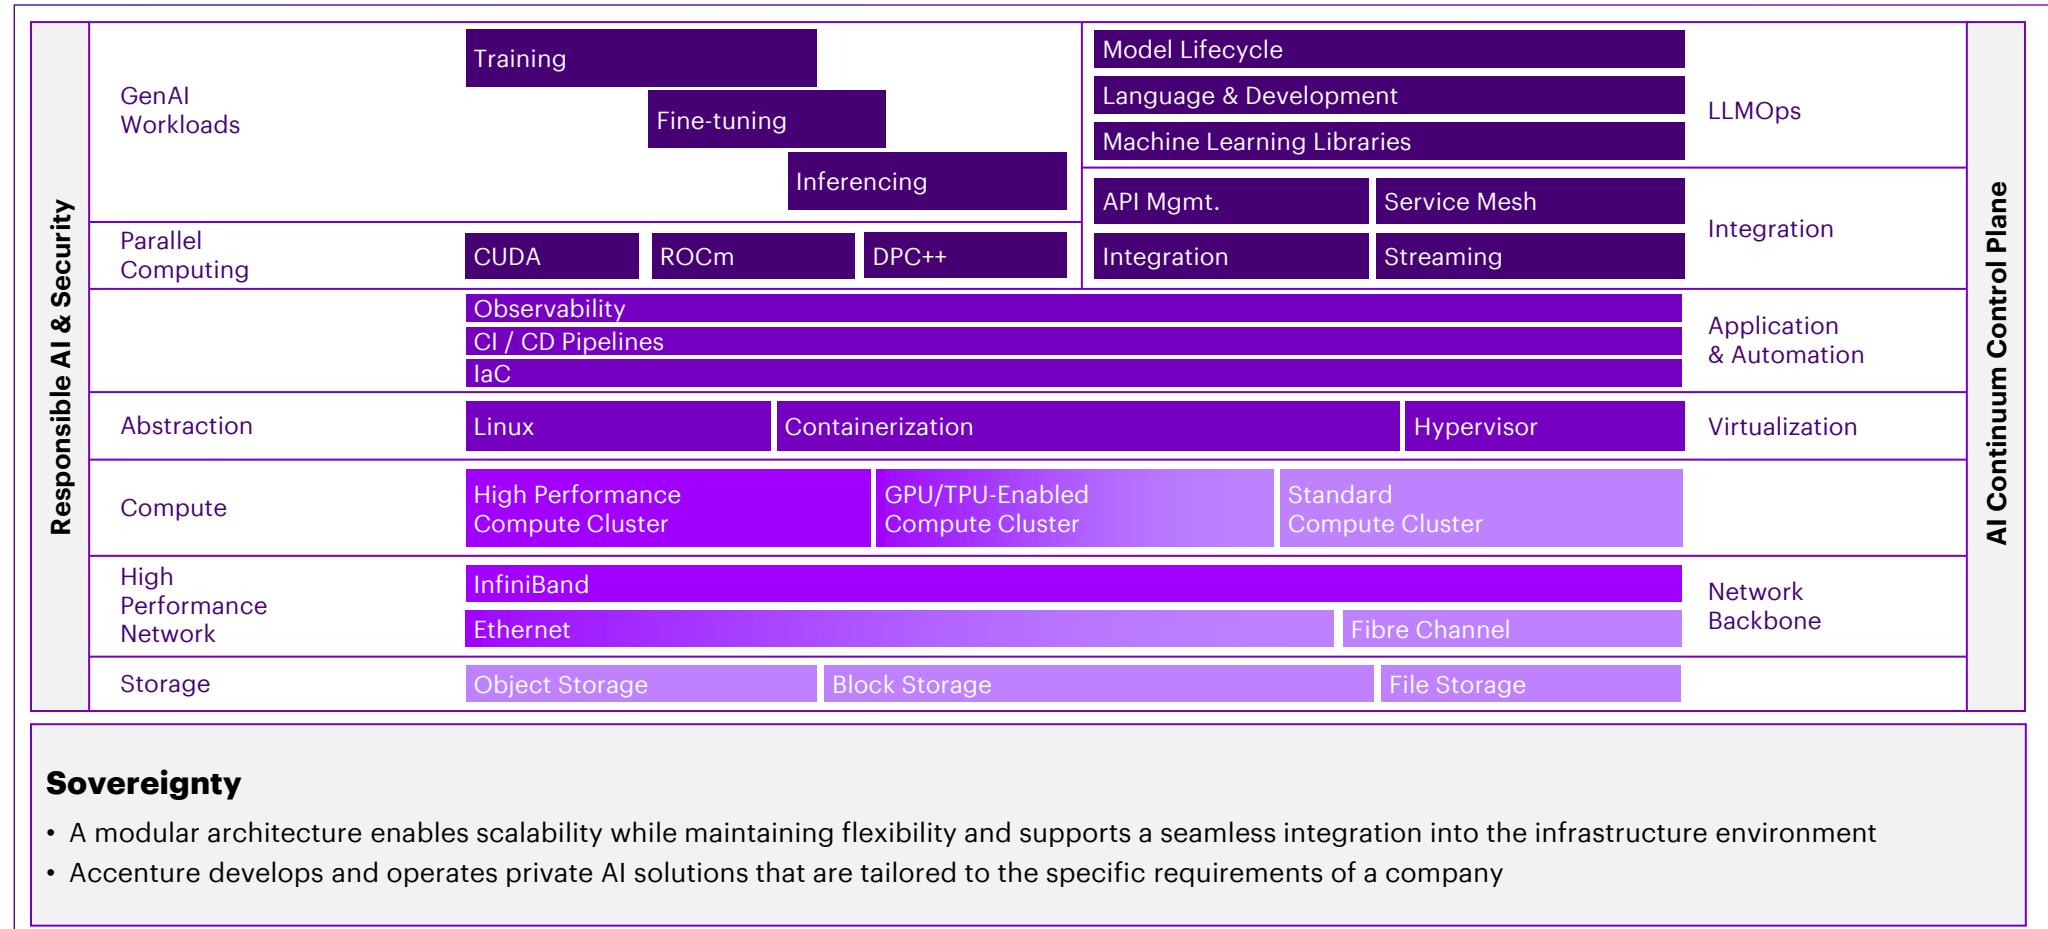
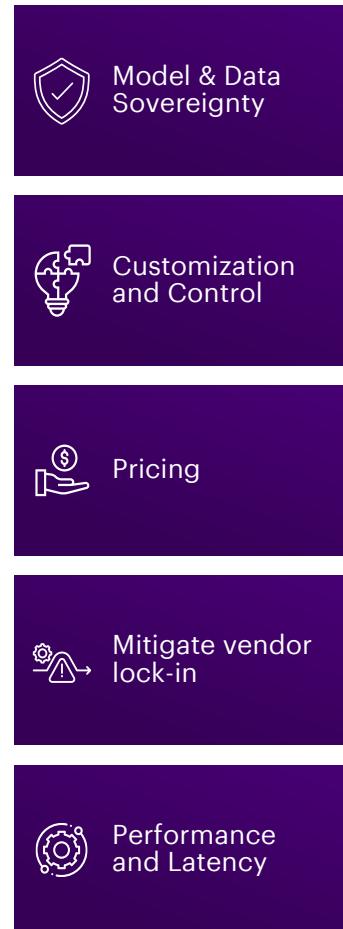
- Avoid proprietary APIs / complex model integrations
- Follow industry standards and your own strategy at your pace



## Performance and Latency

- Hosting the model on-premise has reduced latency and is independent internet connectivity
- The computational demand of state-of-the-art (open source) foundation models can be met with out of the shelf hardware solutions
- The trend towards task specific and smaller LLMs is easing the overall computation demand for productive systems

# The Accenture end-2-end view on the building blocks of a sovereign AI offering



DELL Technologies

Hewlett Packard Enterprise

NVIDIA

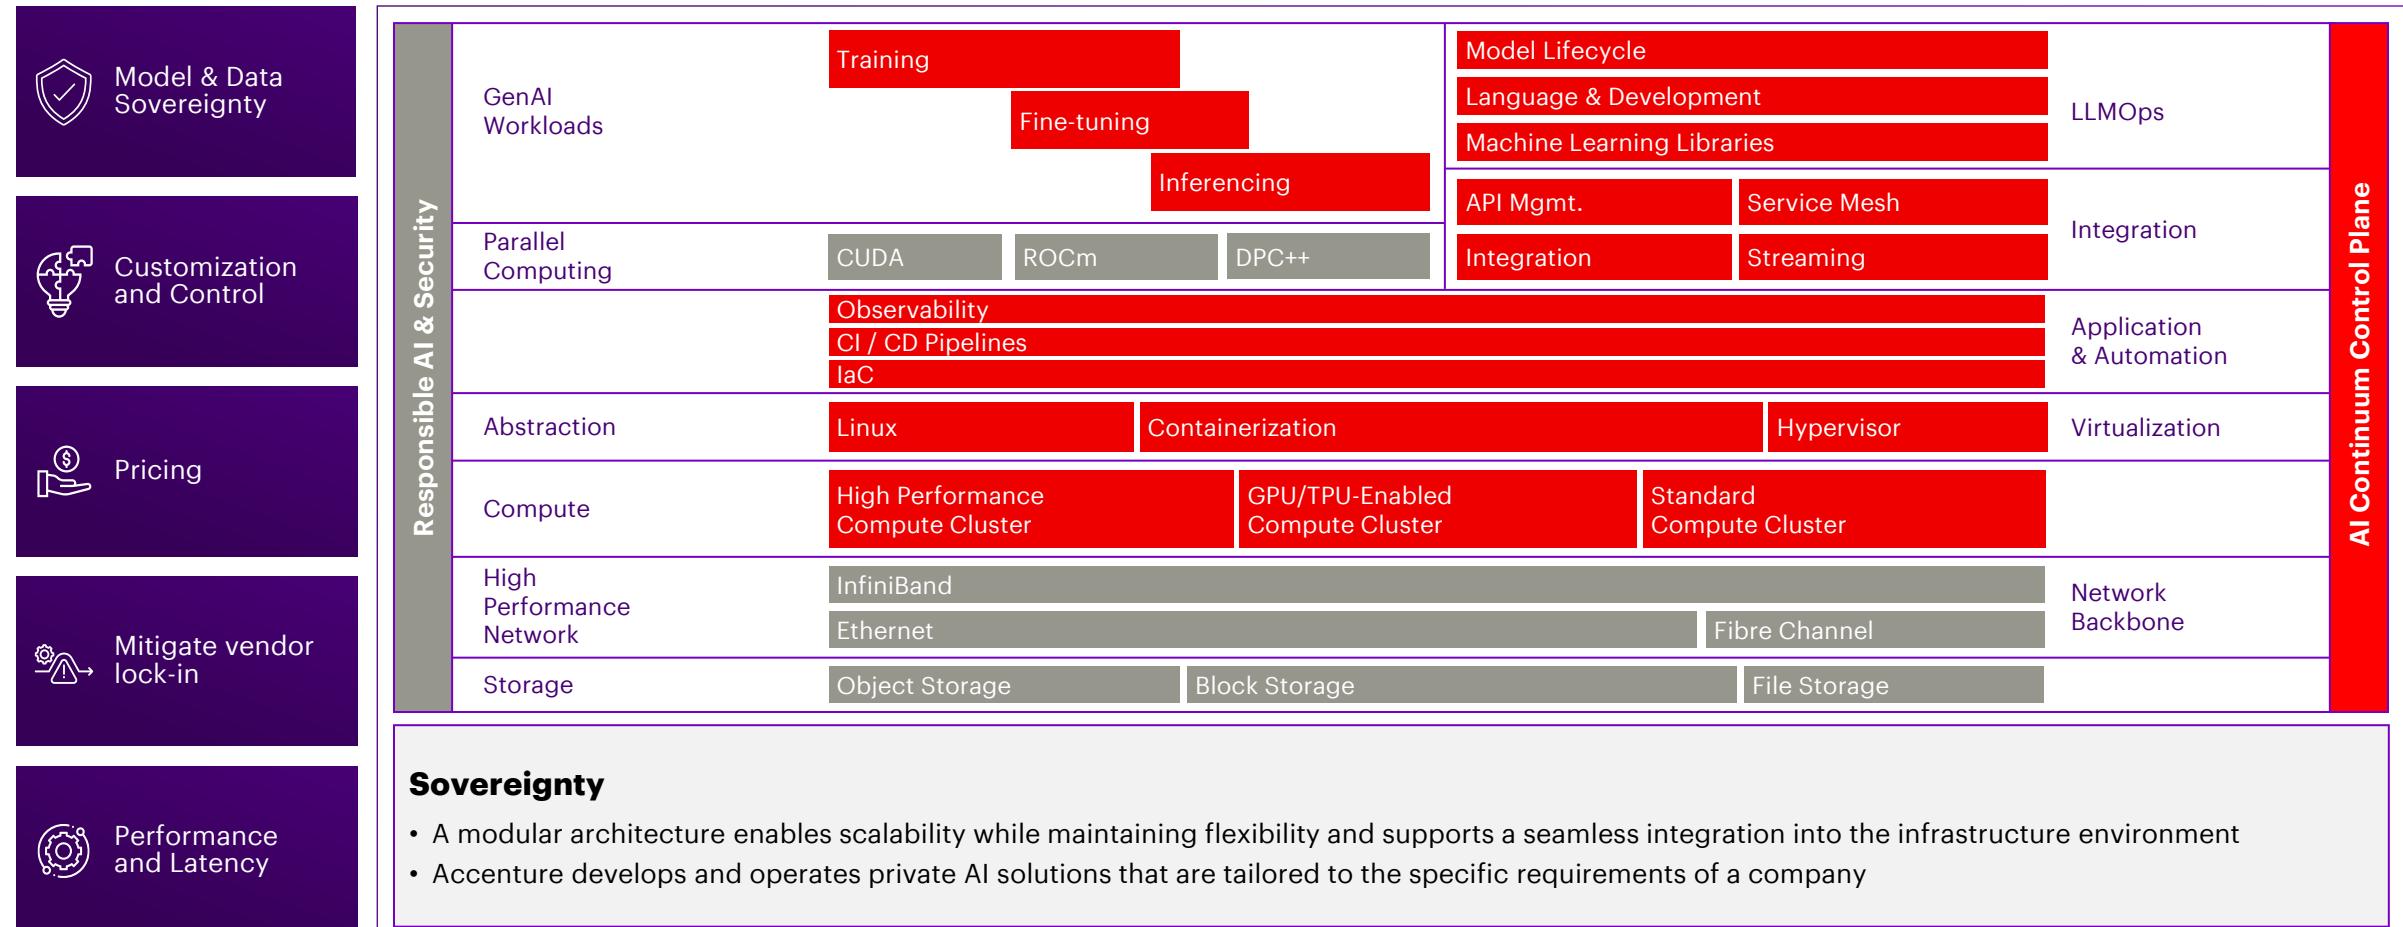
IBM

Red Hat

intel

ORACLE

# Red Hat offers a comprehensive set of functionality to establish a production-ready sovereign AI offering



DELL Technologies

Hewlett Packard Enterprise

NVIDIA

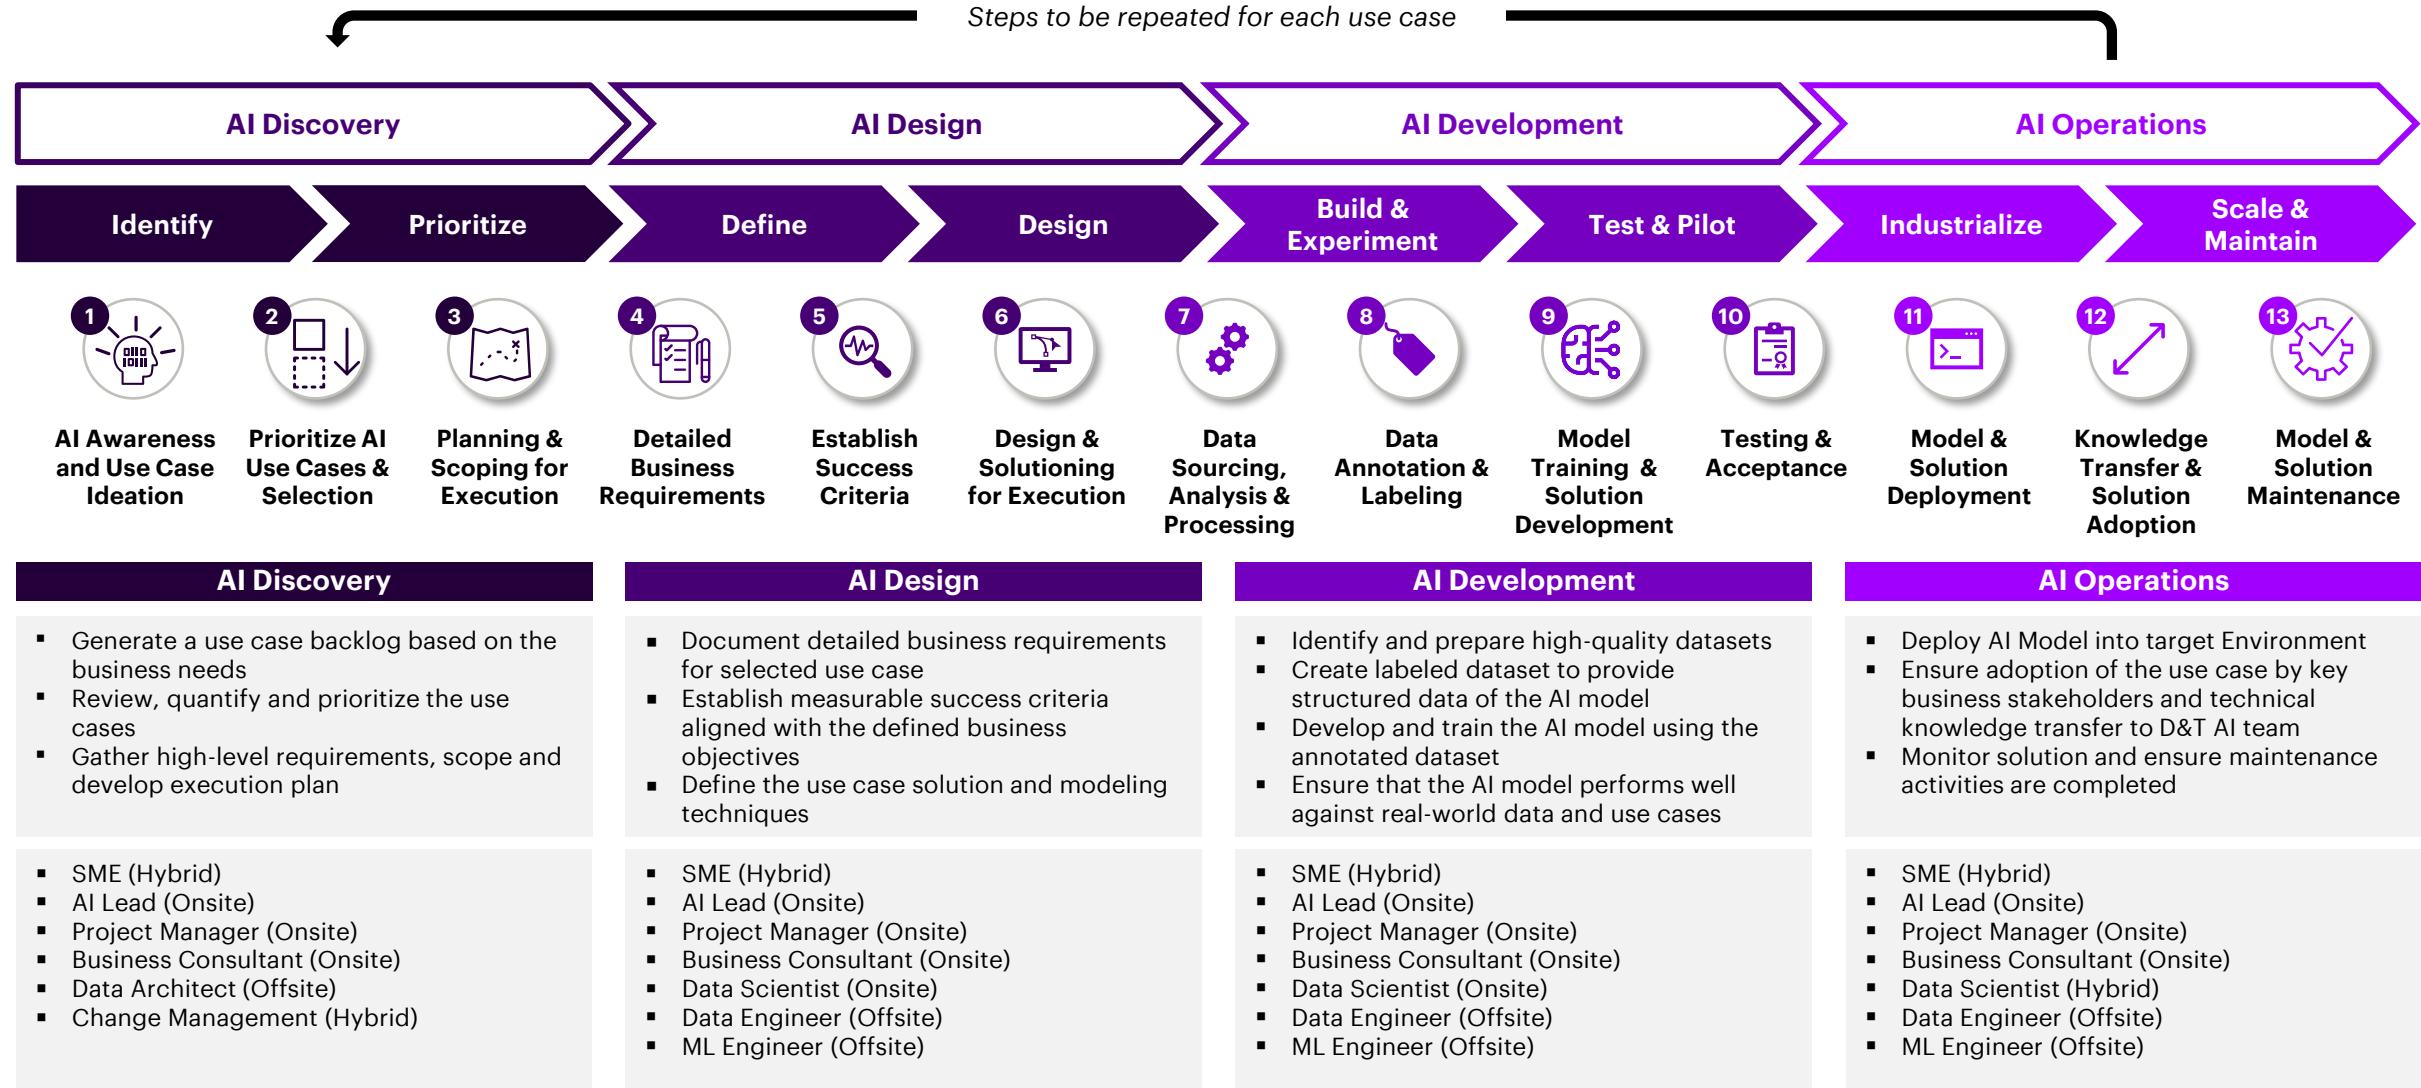
IBM



intel

ORACLE

# Accenture has the means to bring a sovereign AI platform to life through its scalable Application Factory



A nighttime aerial photograph of a city, likely Shanghai, showing a complex network of elevated highways and a large cable-stayed bridge spanning a river. The city lights are reflected in the water, and the overall scene conveys a sense of urban development and infrastructure.

Take the next step –  
learn more now!

accenture



Red Hat