



Microsoft Ignite



BRK 147: What's new in Azure Local

Dean Paron, Microsoft

Meena Gowdar, Microsoft

Mohamed Khalid, GSK



Agenda

1

Azure's adaptive cloud in action

2

Expanding Azure Local in function

3

Voice of the customer: GSK



Azure's adaptive cloud in action

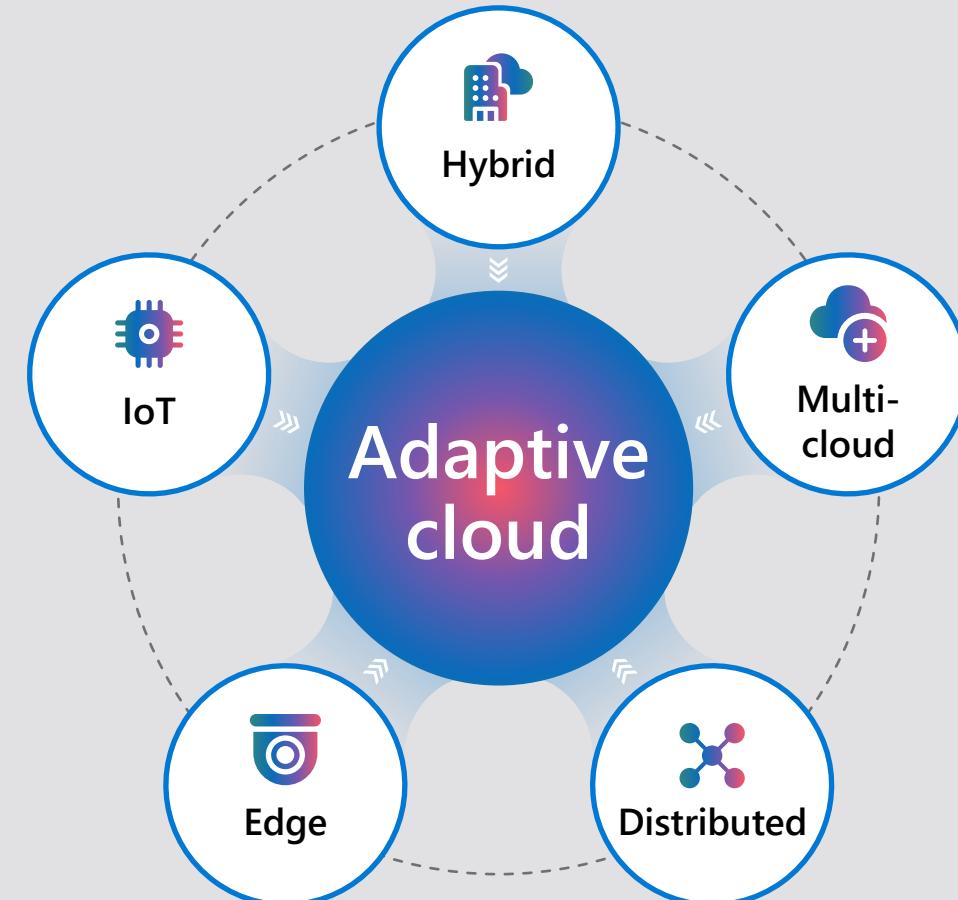


Azure's adaptive cloud approach

Enabled by Azure Arc

Unify hybrid, multicloud, and edge environments with a common operations, innovation, and insights platform.

Azure anywhere



Adaptive cloud approach enabled by Azure Arc

Cloud services
and tools

 Management & security  Copilot  Policy  Defender  Monitor

 Applications  Azure Kubernetes Service  Visual Studio  GitHub

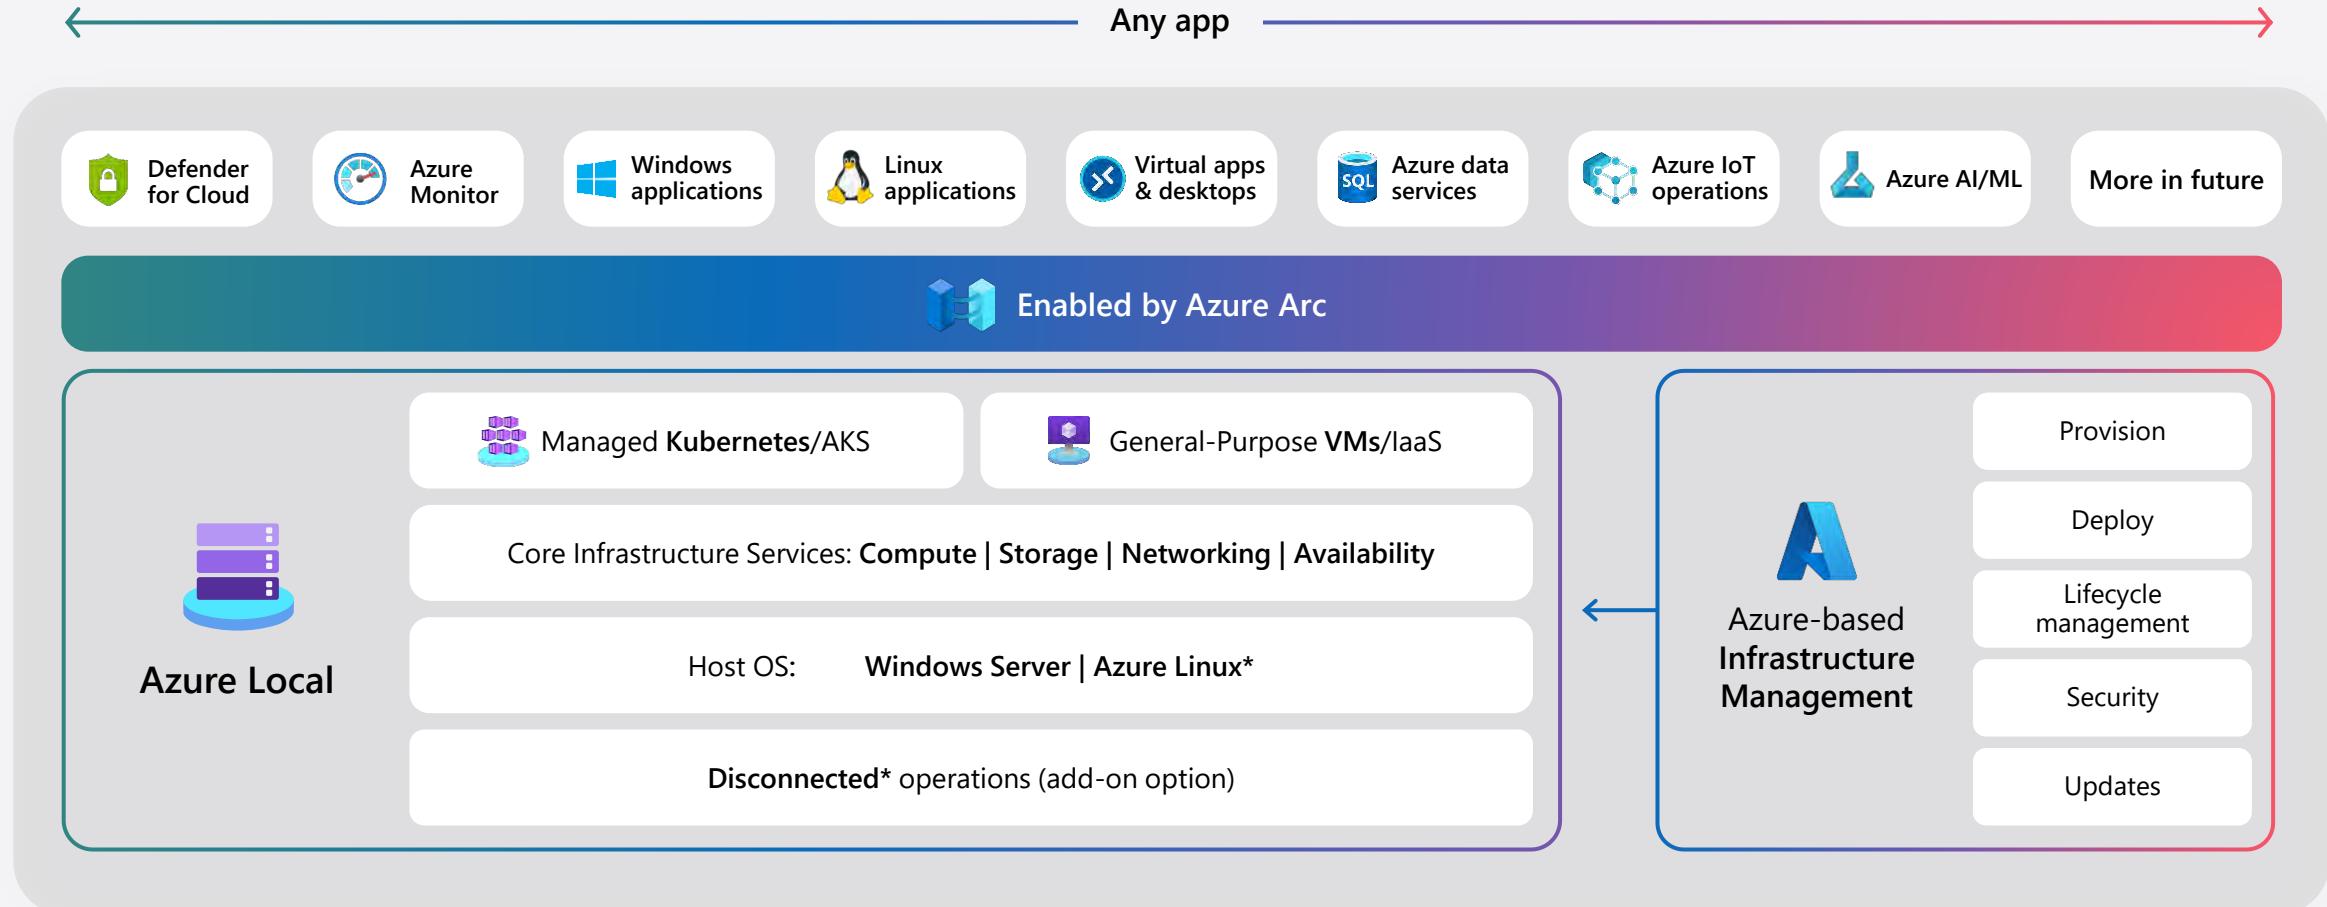
 Data and AI  Fabric  Data services  AI/ML  IoT operations

Global
infrastructure

 Infrastructure  Azure Local | Existing infrastructure |  aws  Windows Server  Google Cloud  Linux  Kubernetes  vmware by Broadcom

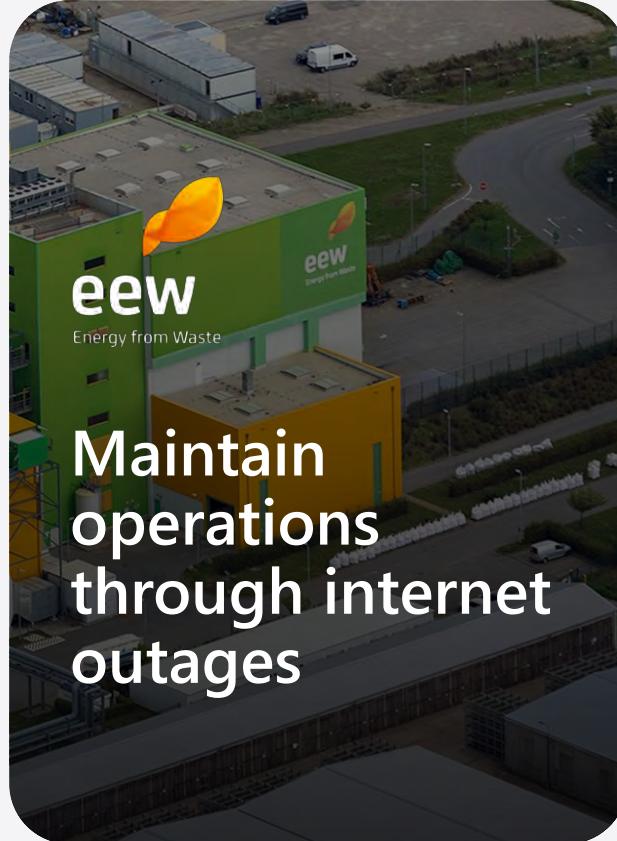


Cloud-connected infrastructure solution for on-premises locations



* in preview

How customers are innovating



Microsoft Sovereign Cloud

Most comprehensive set of sovereignty solutions across AI, productivity, security and cloud platform workloads

Sovereign Public Cloud

Operations and access controlled by European personnel

Data stays in Europe, under European law

AI model choice and AI processing that stays in Europe

National Partner Clouds

Access in-country expertise and nationally owned organizations

Germany: Delos Cloud designed to meet BSI Cloud requirements

France: Bleu designed to meet SecNumCloud requirements

Sovereign Private Cloud

Continuous operations across hybrid or disconnected environments

Integrated cloud and productivity with Microsoft 365 Local

AI capabilities and advanced GPU support with Azure Local

← Consistent infrastructure management and development platform →

Sovereign Private Cloud

Powered by **Azure Local**

Expanding Azure Local: Innovation in function



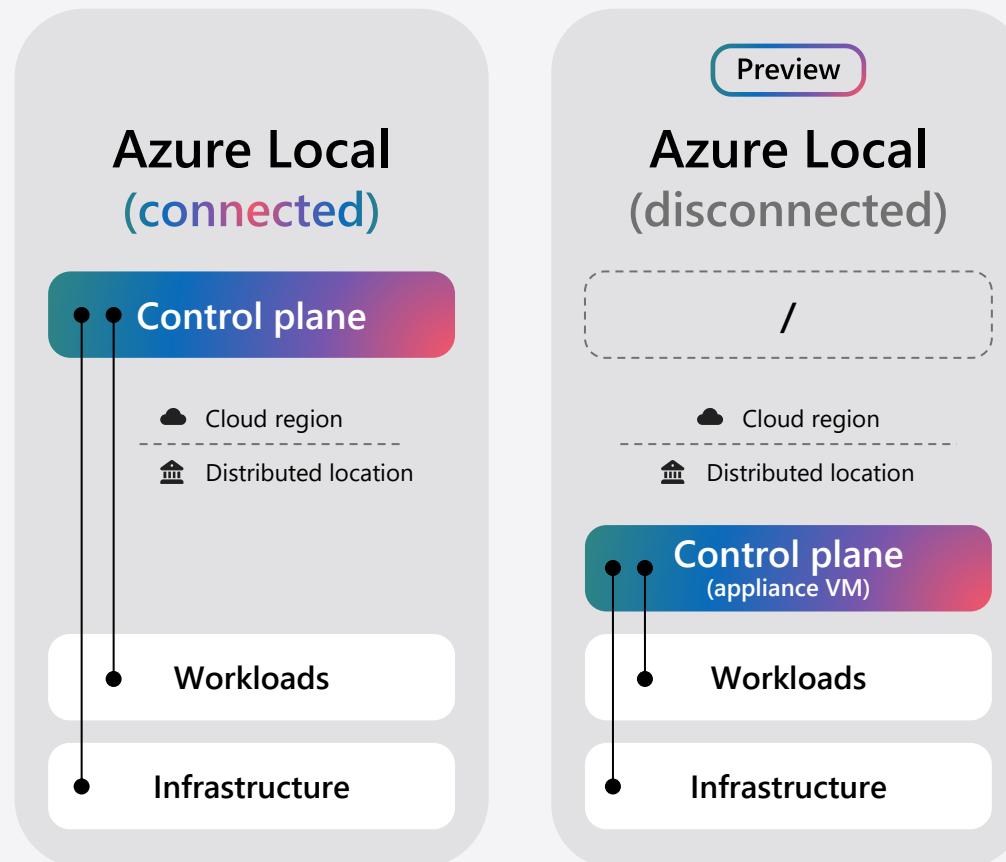
Your path to the sovereign cloud

Public regions support geographic locations worldwide,
for Azure Local clusters deployed around the world



Connected or disconnected flexibility

Operate seamlessly when connected to Azure or fully autonomous in air-gapped



Delivering composable scale points

AI-ready infrastructure, anywhere

Connected to Azure region



Enabled by Azure Arc

Disconnected to Azure region



Enabled by Azure Arc

Managed operations
Day 0 to Day N



Managed Kubernetes/AKS

General-Purpose VMs/IaaS

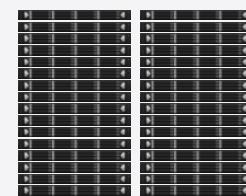
Core Infrastructure Services: Compute | Storage | Networking | Availability

Host OS: Windows Server | Azure Linux

Azure-based infrastructure management

Provisioning, deployment, lifecycle management, security, updates

Composable scale point for sovereign AI from single nodes up to multiple racks clusters



Preview

some scale points and disconnected for multi-rack scale offering will be coming in the future

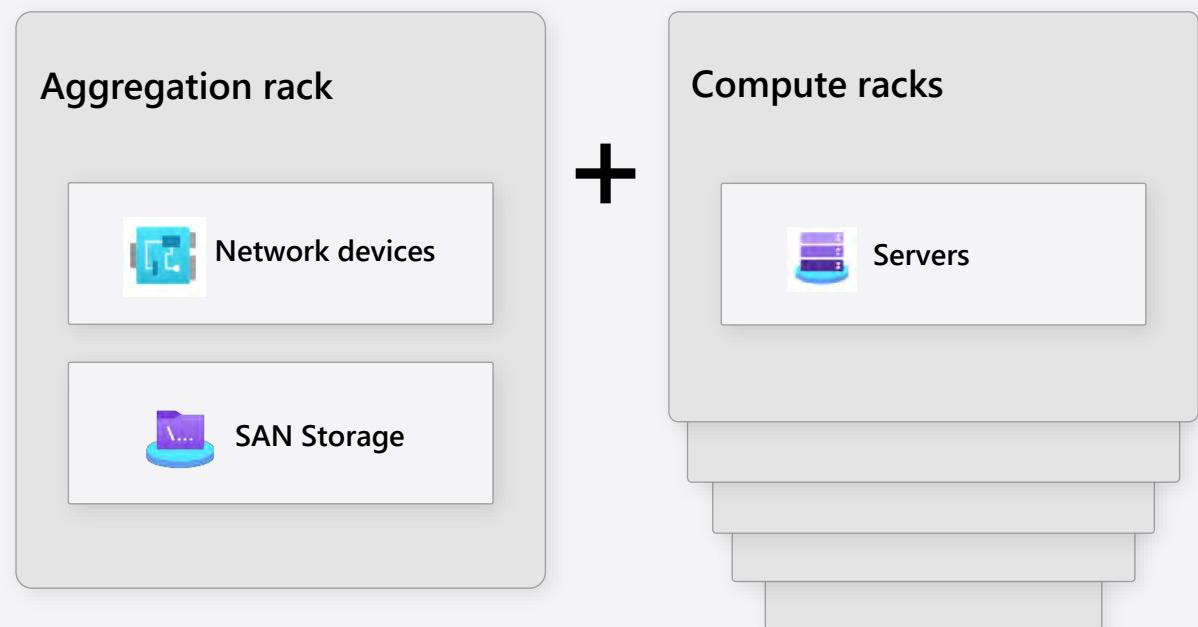
Preview

Multi-rack deployments of Azure Local

Expands the maximum scale point of Azure Local up to hundreds of nodes per instance

Delivered as integrated racks of compute, storage, and networking

Fully managed compute, storage, and networking via Azure



Preview

External storage for Azure Local

Now available with



Bring your own SAN storage to keep data local with full operational control



Choose trusted storage partner solutions to meet sovereign needs

Deliver enterprise-grade performance with Fiber Channel SAN side-by-side with S2D

Microsoft 365 Local

Supports Exchange Server, SharePoint Server, and Skype for Business Server



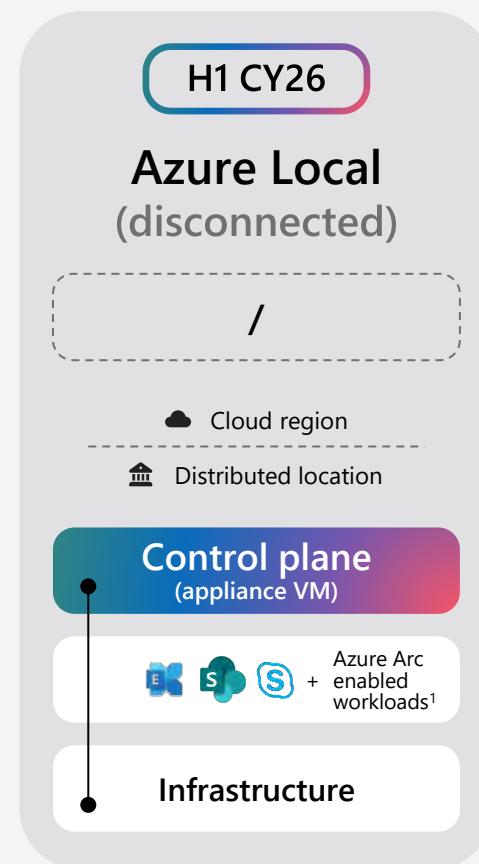
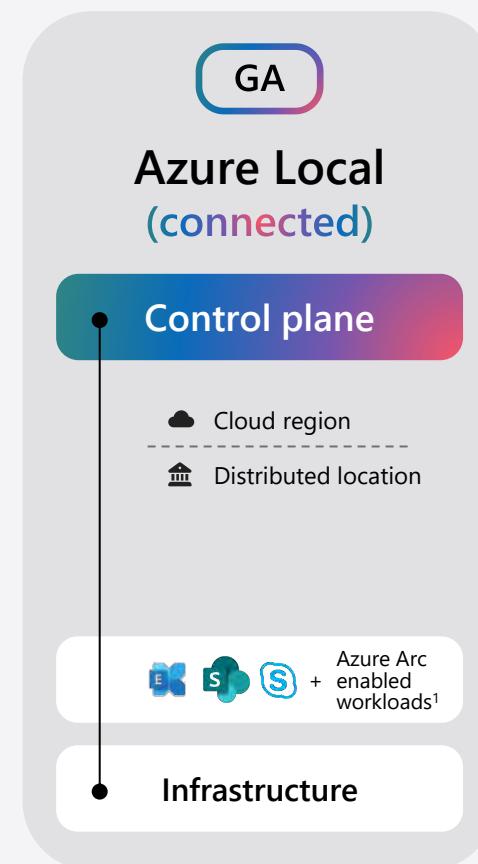
Familiar productivity suite, sovereign-ready, locally deployed for data residency and sovereignty needs



Deploy and run workloads and services in a private cloud through Azure-consistent experiences



Simplify operations and manageability by running Microsoft 365 on Azure Local



Coming soon

Nvidia RTX Pro 6000 Blackwell Server Edition

AI Acceleration at the Edge



High-performance AI at the Edge

Provides data center-level GPU acceleration for AI tasks on Azure Local, minimizing latency and supporting compliance



Designed for generative AI and LLM inference

Runs large language models and multimodal AI tasks locally for instant responses without relying on cloud access.



Advanced Visualization & Simulation

Ideal for rendering, digital twins, and immersive 3D experiences (e.g., Omniverse) in industries like manufacturing, energy, and defense.



Secure, Sovereign Compute

Combines NVIDIA's GPU performance with Azure Local's jurisdictional control, meeting strict data residency and security requirements.

Coming soon

Nvidia RTX Pro 6000 Blackwell Server Edition on Azure Local

Powerful universal GPU for AI and visual computing on-premises



Azure
Virtual
Desktop



NVIDIA
Omniverse
Enterprise



Azure Local
Virtual
Machines



Azure
Kubernetes
Services



Edge
RAG



NVIDIA
NIM

Coming soon
on a variety of platforms



DELL Technologies

AX-770



Lenovo

MX650a V4



HPE

DL380 Gen 12

Designed to run 1000+ models such as GPT OSS, DeepSeek-V3, Mistral NeMo, Llama 4 Maverick, RTX Pro 6000 BSE enables organizations to accelerate their AI initiatives directly within a sovereign private cloud environment.

Sovereign Private Cloud AI stack



BYO RAG

Bring your own RAG application and logic



AI Services



AI Video Indexer

AI video and audio insights from live* or recorded sources



OpenAI OSS



Llama



Mistral



DeepSeek



Microsoft

Foundry Model Catalog

Task-specific models

OCR, speech, vision, translation, and more

BYO models

Foundry Local Model Server

(inferencing and runtime deployment*)

AI Infrastructure

NVIDIA GPUs

Leading OEM Partners

Dell, HPE, Lenovo, and more

Small to large footprints

Single node to multi-rack deployments

← Connected and fully disconnected operations with Azure Local →

DEMO

Nvidia RTX Pro 6000 Blackwell Server Edition GPUs on Azure Local



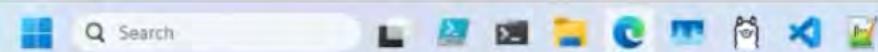
n42r0604-cl | Machines

Azure Local

- Overview
- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems
- Resource visualizer
- Settings
- Infrastructure
 - Machines
 - Resources
 - Virtual machines
 - Kubernetes clusters
 - Logical networks
 - Disks
 - VM images
 - VM network interfaces
 - Network security groups
 - Storage paths

[Refresh](#)

Name	Status	Manufacturer	Model	Serial number	Cores	Memory	OS version	OS bu
ASRR1N42R06U29	Connected	Lenovo	MX650 V3 PR w/ Ctrl GPU	JZ001RVN	48	512 GiB	24H2	26100.
ASRR1N42R06U31	Connected	Lenovo	MX650 V3 PR w/ Ctrl GPU	JZ001L6E	48	512 GiB	24H2	26100.

Add or remove favorites by pressing **Ctrl+Shift+F**

Edge RAG

Intelligent Knowledge Handling



- LazyGraphRAG for high-quality retrieval
- OCR-based parsing of documents, tables, and images, with SharePoint integration.



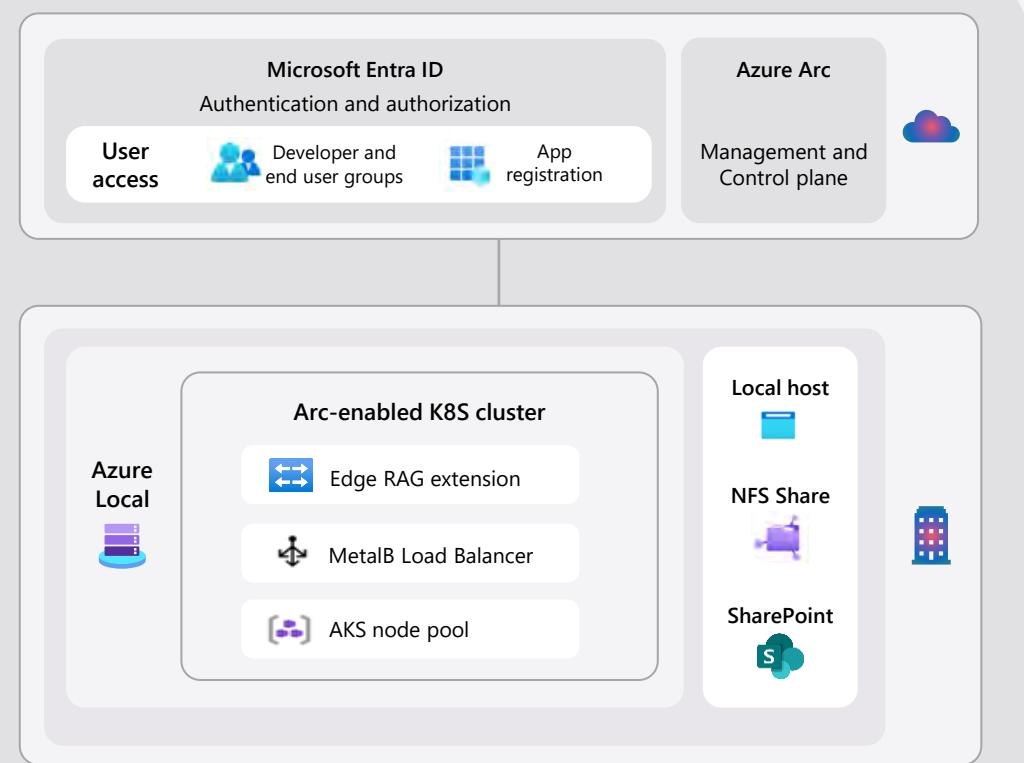
- 100x faster ingestion of live-streamed images
- 5x faster query performance for Hybrid Search



- Multimodal search, image retrieval & image-rich outputs with chat UI upgrades
- Connected and fully disconnected scenarios



- Agentic workflow with OCR & LLM integration focusing on CEQA (California Environment Quality Act) applications



Preview

Vision AI: Azure AI Video Indexer

Enable Real-Time Video Intelligence



Create custom AI models with natural language and apply object detection in restricted areas.



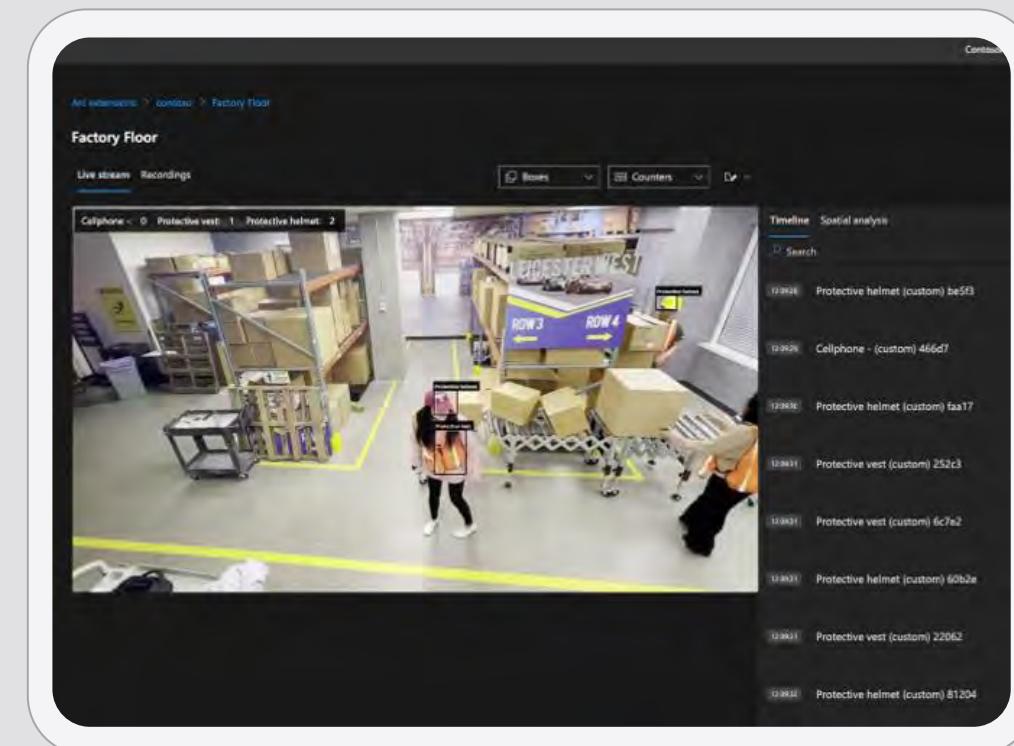
AI video assistant chat: Ask questions in natural language and get real-time answers



Video agents for retail, service, security, and sales to monitor safety, experience, and performance



Capture key events and insights using natural language to generate summaries



Generally available

Azure Migrate

Move workloads seamlessly from VMware to Azure Local



Full-stack alternative with Azure Local



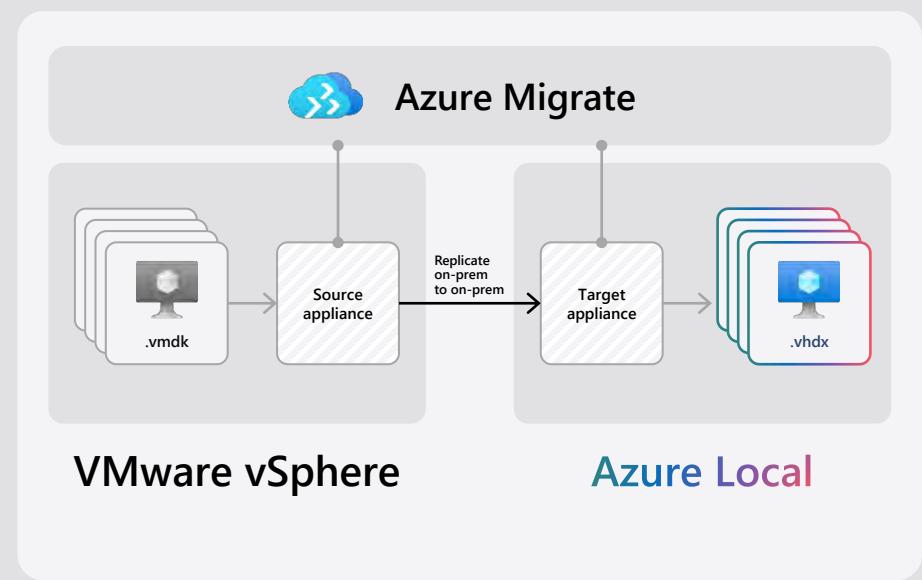
No need to change/rewrite apps



Copy and convert VMware VMDK to Azure Local VM enabled by Azure Arc entirely on-premises at no cost



Guided, agentless workflow in Azure Migrate



aka.ms/AzureLocal/Migrate

Generally available

Network Security Groups

Protect and isolate your network traffic



Protect network access with inbound and outbound allow and deny rules



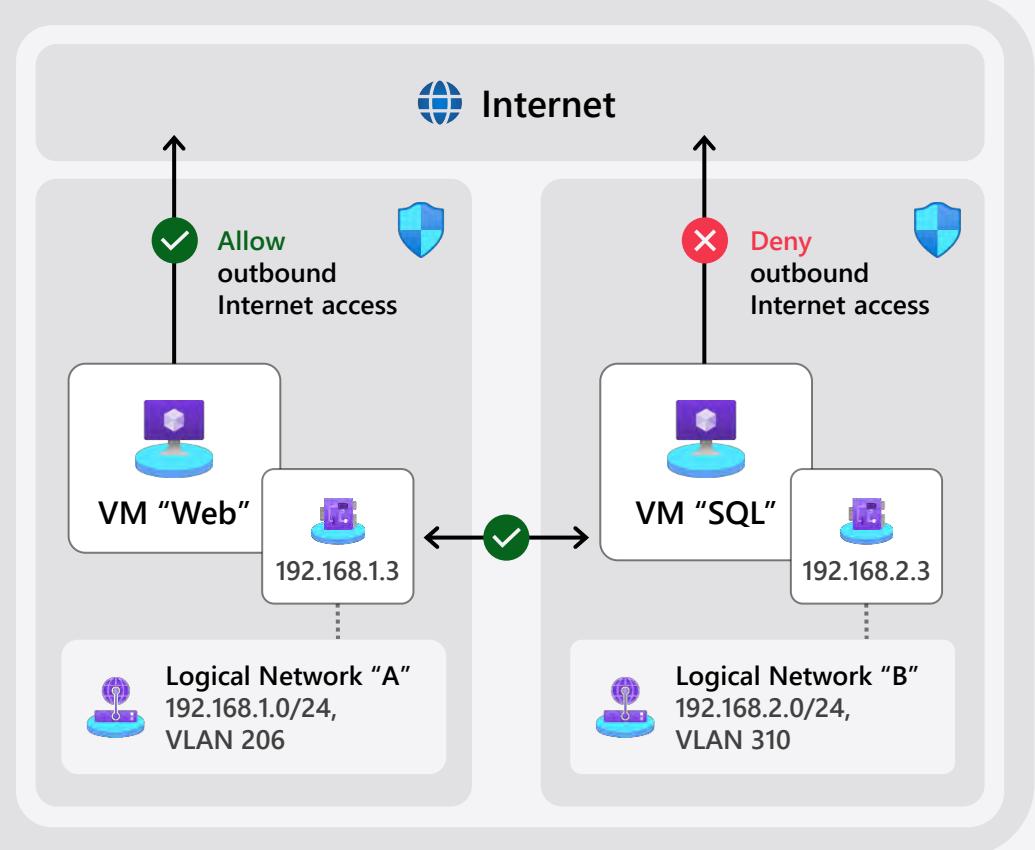
Complete 5-tuple control: source and destination IP, port, and protocol



Assign NSGs to individual VM interfaces or whole logical networks



Enforced within the virtual switch at the virtual port level



Local Identity without Active Directory

Simplified identity model with Azure Key Vault



Deploy without Active Directory – hosts don't need to be domain joined



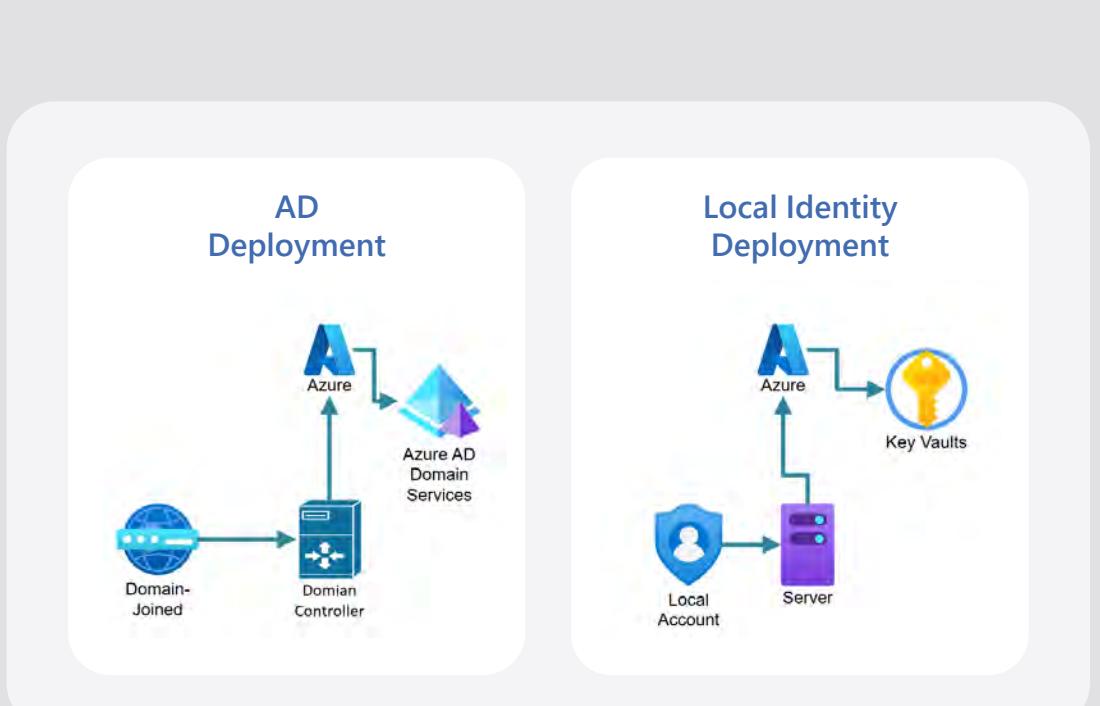
Full feature parity with AD-based deployment, including live migration and other fabric operations



Secrets automatically backed up to Azure Key Vault - local user passwords and BitLocker recovery keys



Built-in DNS Support – customers no longer need to provide their own DNS



Rack aware clusters

Improve fault tolerance and workload resilience



High availability (HA) between two rooms in the same location



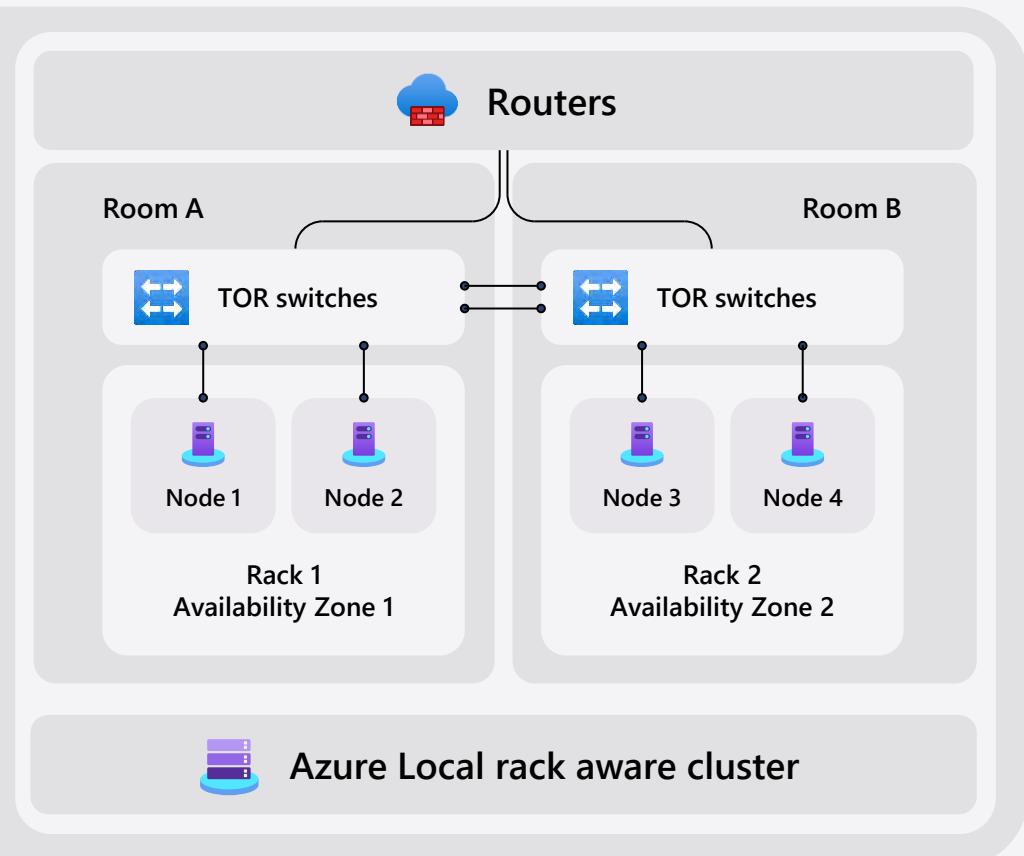
Synchronized S2D storage replication between local availability zones



Provision workloads to a specific local availability zone



Automatic workload and volume failover and fallback



Coming soon

Sovereign Private Cloud

What's coming

Enhanced Change Control

Configurable policies
and approval workflows
with explicit oversight

Site to Site Disaster Recovery

Business continuity by
keeping workloads
running during outages

Hybrid to Fully Disconnected

Flexibility to seamlessly
transition from hybrid
to fully disconnected

Edge transformation at GSK: Powering manufacturing sites and R&D labs with Azure Local

Mohamed Khalid
Senior Director, Solutions Architect



GSK: A focused biopharma company



**Our purpose: We unite science, technology and talent
to get ahead of disease together**

Expertise in developing
innovative specialty
medicines and vaccines

Leading portfolio of
vaccines, specialty and
general medicines

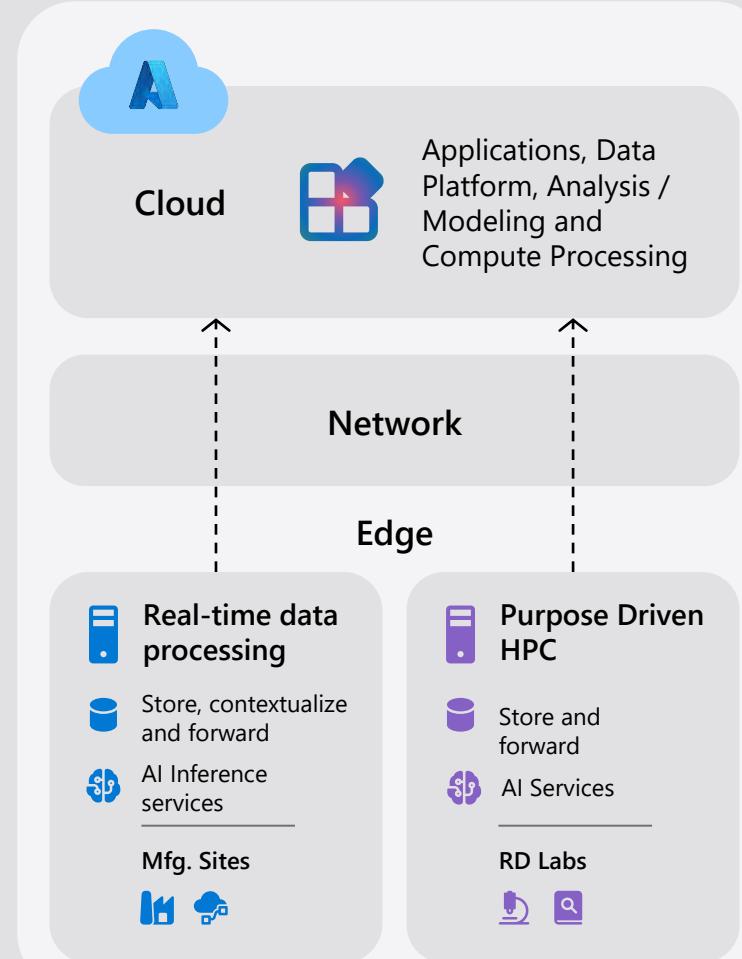
Almost 70,000
people and a presence
in more than 75
countries worldwide

The challenge



GSK has a global presence of manufacturing sites and R&D labs that rely on on-premise systems (Edge)

- Low-latency requirements: Need for sub-millisecond feedback for certain business applications require local processing.
- Inconsistent architectures, legacy systems, and siloed operations hinder automation, increase cost, and slow modernization.
- Lack of unified visibility across distributed infrastructure makes it hard to manage performance, compliance, and incidents consistently.
- On-prem environments lack access to the advanced data, ML, and automation services available in the cloud, limiting AI adoption and innovation at the edge.
- Regulatory and security compliance necessitate local solutions.



Bringing cloud innovation and AI to the edge with Azure Local



Azure Local extends cloud capabilities into manufacturing sites and R&D labs, enabling real-time data processing, contextualization, and AI inference directly at the edge.



Azure Local delivers a single management plane and consistent tooling for deployment, policy enforcement, observability across cloud and edge.



Enables compliant, self-service deployment of apps, data and AI via standardized GSK qualified architecture



Integrated Factory & Lab Operations
Connects Edge data and systems into Azure data and AI services, improving efficiency, quality, and productivity.



Standardized, future-ready architecture — delivers a secure platform with consistent patching, OS updates, and lifecycle management.



Lessons learned



A skilled technical team is essential to design, deploy, and operate Azure Local with compliance, security and resiliency at the core.



Seamless integration with network, identity, security and monitoring services must be designed upfront.



Success depends on strong coordination between internal teams, Microsoft and hardware vendors.



Ensure standard operation procedures are documented, Operations team trained before go-live.



Manage Cost via optimization of resources and manage Capacity strategically.

Building the foundation for the future



Faster service enablement, from months to hours

What once took months, like standing up a Kubernetes cluster, **now takes hours** with Azure Local's automated, cloud-consistent deployment.

Foundation for convergence and accountability

Established a **consistent IT/OT model** and enabled FinOps-based chargeback, **improving visibility, cost allocation, and operational efficiency**

Innovating for future proof architecture

Data extends from edge to cloud for GSK to leverage intelligent applications — AI, GenAI, and Digital Twins — supporting **GSK's Smart Manufacturing and Future of Labs vision**

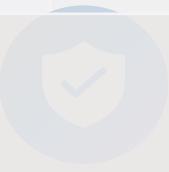
Q&A





Azure Accelerate

Fuel transformation with experts and investments



Access trusted experts

Specialized partners | Cloud Accelerate Factory

Factory



Unlock Microsoft investments

Partner funding | Azure credits

Investments



Succeed with comprehensive coverage

Spans migration, modernization, and AI scenarios

aka.ms/AzureAccelerate

Spans migration, modernization, and AI scenarios

Includes new zero cost deployment benefit with

Cloud Accelerate Factory

Includes new zero cost deployment benefit with Cloud Accelerate Factory