How to Scale Intelligent Factory Initiatives Through an Adaptive Cloud Approach

Discover the underlying architecture that helps manufacturers gain the value of AI across all factories and build strong collaboration between IT and OT

By Christoph Berlin and Kam VedBrat

Abstract

The journey to industrial transformation has progressed from individual production assets connected to computers, to networked, on-premises systems, to hybrid cloud systems. In the drive to innovate, manufacturers have spun up dozens of initiatives, each with its own approach. The result is sprawling systems, siloed teams and operations, and costly short-term solutions that often don't work together but, instead, trap data in inaccessible systems. The rising demand for artificial intelligence (AI) solutions is heightening concerns among industrial enterprises about the most effective way to scale to meet these challenges.

It's clear that there's a need for a unified underlying architecture to support new intelligent factory initiatives, while optimizing existing investments. One way to achieve that is through an adaptive cloud approach. Such an approach unifies siloed teams, distributed sites, and sprawling systems into a single operations, security, application, and data model. This approach enables manufacturers to use cloud-native and AI technologies to work simultaneously across hybrid, multicloud, edge, and Industrial Internet of Things (IIoT) systems.

An adaptive cloud approach will help both information technology (IT) and operational technology (OT) teams maximize resources with widely adopted standards, optimize production with repeatable solutions, and efficiently manage their complex environments with a common architecture.

Whether a company is already on its AI journey or is exploring how and where to begin, this paper provides insights on an underlying architecture that can help manufacturers efficiently and effectively apply AI from the shop floor to the C suite. Factories around the world are starting to use this framework to adapt IT solutions to meet manufacturing production requirements that support scalable, automated, resilient, and secure factories.

Table of contents

Executive summary

The factory of the future is an intelligent, connected, and flexible manufacturing ecosystem that uses technologies such as AI, robotics, IIoT, cloud computing, and advanced communications to fundamentally change how manufacturers operate. Pursuing this vision can help manufacturers address some of their most pressing challenges, such as cybersecurity, regulatory compliance, core talent recruitment and retention, and centralized insights and control, while becoming more agile, collaborative, resilient, and sustainable.

However, manufacturers still have to contend with diverse and complex manufacturing environments and the need to achieve mission-critical, core production metrics such as improving quality, reducing defects, increasing yield rates, and reducing per-unit cost. Most manufacturers assemble solutions from in-house developers, partners, and industry-specific proprietary solutions to meet specific business requirements.

Managing this heterogeneous environment has become increasingly challenging. Over the past decade, there have been multiple attempts to transform manufacturing with unique, factory point solutions. However, rather than streamlining operations, in many cases this has produced technical sprawl and inefficiencies that create major obstacles to being able to scale solutions—including AI—to multiple production lines and geographically dispersed factories.

Many manufacturers struggle with vendor-specific and scenario-specific solution silos, where data and analytics are isolated and managed through different interfaces. This leads to a lack of global visibility, high management costs, and limited solution scalability. Moreover, there's often a misalignment between IT and OT environments, which hampers seamless data integration and operational efficiency. Enterprises currently can harness data, derive insights, and automate decision-making, but only in isolated use cases and at a high cost, instead of realizing this transformation globally.

A new framework for resource management, data integration, and application development is required to accelerate digital transformation so that manufacturing enterprises can benefit from AI insights on a global scale. Applying AI allows enterprises to maintain both a single production-line view and a global view that compares and benchmarks similar lines across all factory sites worldwide. Pairing historical data with generative AI creates a powerful analytics tool for IT and OT users and decision-makers.

Microsoft believes that an adaptive cloud approach establishes a solid foundation for achieving these benefits. An adaptive cloud approach enables integrating advanced technologies across globally distributed factories, while addressing complexity challenges. It centralizes management and security, deploys and updates applications at scale, and builds a unified data estate that spans edge and cloud. The adaptive cloud approach also uses the best of both IT and OT tools to maximize resources and optimize factory efficiency, while enabling many different types of equipment and software used at the edge and global unification.

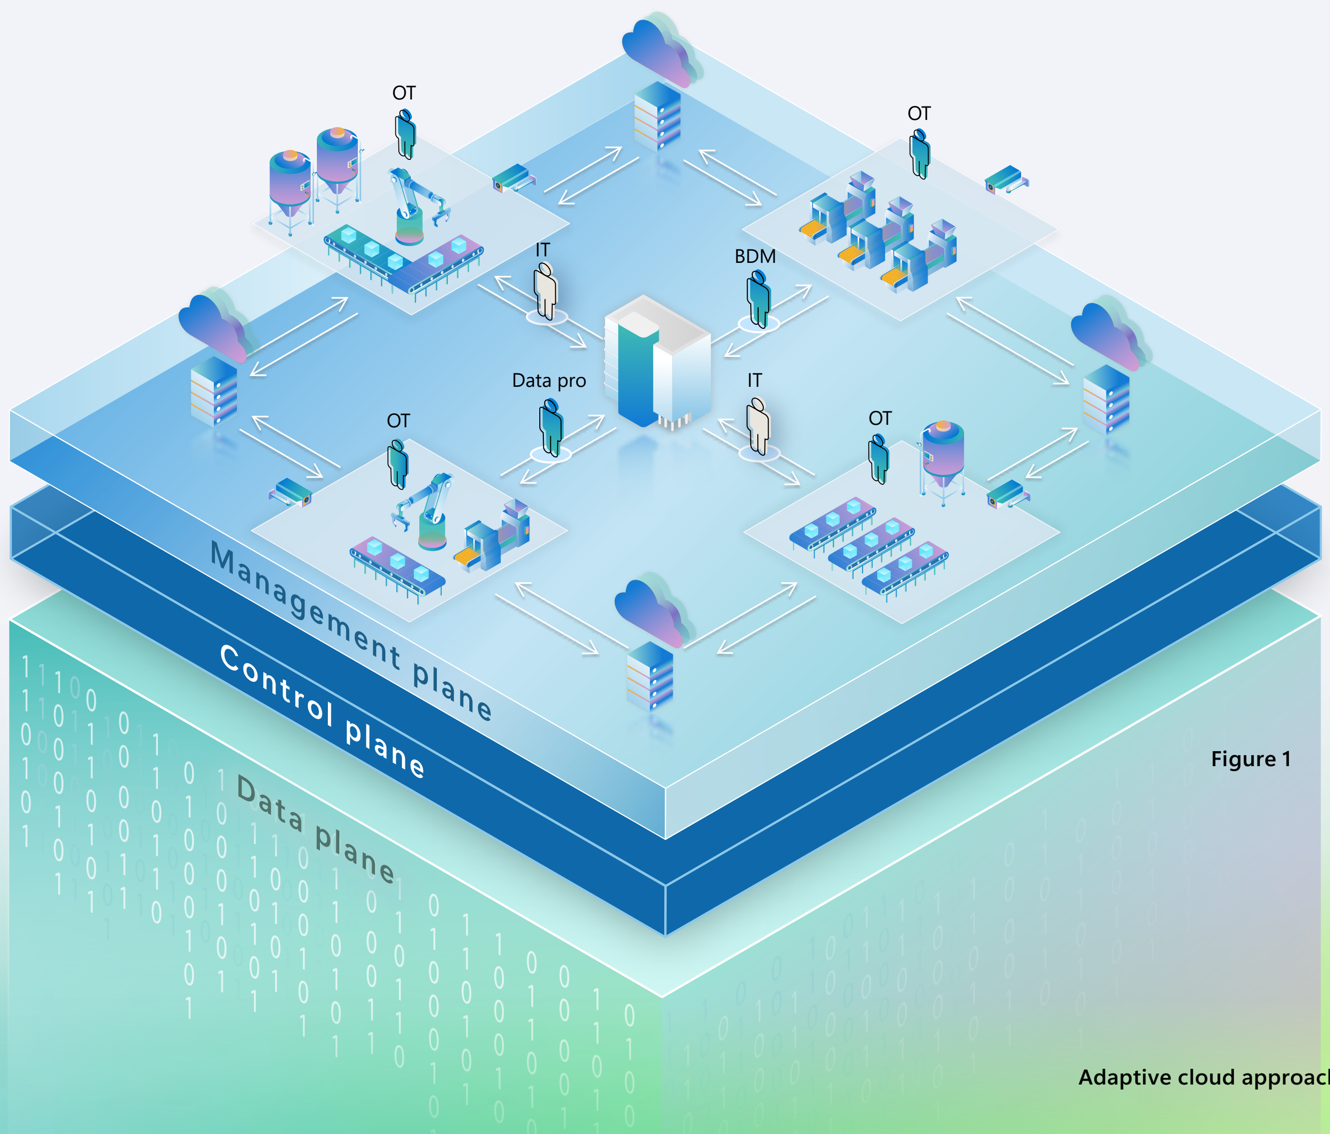

Figure 1

The following section details the scalability of an adaptive cloud approach and emphasizes Azure services that ensure repeatability across different lines and factories:

Management plane:

- Provides repeatable, centralized management for distributed IT and OT environments in both the cloud and on the edge, using the Azure portal, Azure Resource Manager, and Azure Arc.
- Provides a unified view of physical sites, processes, and assets, thus minimizing the number of interfaces they have to work with and providing them the ability to apply policies globally.

Control pane:

- Extends application deployment and management consistently from cloud to edge with Azure Kubernetes Service (AKS) or other Kubernetes orchestration services.
- Uses Kubernetes and Azure Arc to act on local resources such as hybrid, multicloud, edge and IoT, enforcing policy set by the Azure management plane.
- Ensures resilience to cloud-connectivity interruptions and maintains local control for OT change management.

Data plane:

- Supports building a unified data estate that encompasses local and global requirements—for example, meeting edge data needs with Azure IoT Operations, and cloud data needs with Microsoft Fabric or other edge-to-cloud data and analytics services.
- Supports processing and using more data closer to mission-critical operations, and interoperability between custom-built and partner solutions, while ensuring insights can be sent to the cloud for further contextualization and analysis.

EXECUTIVE SUMMARY

Microsoft is committed to helping customers realize the promise of Industry 4.0 through an adaptive cloud framework. The goal is to enable industrial enterprises to gain the most value with the least disruption to operations. As part of that commitment, we have built a robust partner ecosystem that includes original equipment manufacturers (OEM), system integrators (SI), independant software vendors (ISV), connectivity partners, and industrial automation partners.

These partners take an adaptive cloud approach to deploying value-added services that help customers lower deployment costs and simplify infrastructure and data integration. This approach allows customers to deploy any application or workload across multiple orchestration and edge platforms and enable AI at scale to unlock operational efficiencies and savings.

Microsoft also provides a comprehensive set of security and compliance tools spanning cloud and edge, such as Microsoft Defender, Microsoft Purview, Microsoft Sentinel, and the Microsoft Entra Suite, to help protect OT environments from threats and vulnerabilities.

This white paper examines the core concepts of an adaptive cloud approach to help manufacturers understand how to use this framework at every factory to lay the groundwork for applying AI at scale.

Six ways an adaptive cloud approach helps manufacturers to scale intelligent factory initiatives

1. Evolving IT and OT convergence

IT and OT convergence describes a pattern where IT and OT environments and platforms begin to merge with unified business processes and orchestration rules, to achieve the same operational efficiencies. This approach has proven to be difficult because IT and OT environments have different priorities and follow their own individual practices. These differences include:

IT software life cycle management fundamentally differs from OT hardware requirements

Generative AI in the manufacturing industry will have the greatest immediate impact on the IT adopts a proactive, centrally initiated, update-push model, while OT prioritizes preventing unplanned production stoppages. These differences require careful timing of updates with rollback options. OT also defers manufacturing equipment-related updates until absolutely necessary, requiring robust change management before initiating any IT updates. Understanding these differences is key to effective IT and OT collaboration.

Change management is an OT priority

Production uptime, with zero unplanned manufacturing stoppages, is an OT mission-critical requirement. While IT is aligned with this goal, it's necessary to adopt robust change management with OT as the final approval to avoid having changes negatively affect production. In other words, orchestration control and supervision are more important in the OT space.

The current OT hardware life cycle is significantly longer than the IT software life cycle

OT hardware is typically specified for a minimum 10-year replacement cycle, but it's common for OT hardware to actually run productively for many years longer. The IT software life cycle is typically 5 to 10 years. The potential risks involved in patching or other changes to production equipment firmware or software lead manufacturers to adopt a "set it and forget it" approach when deploying production equipment. Once installed and validated, the underlying production equipment software is rarely touched before the equipment is retired. To address these differences, the industry is moving toward adopting open standards and central management platforms that can accommodate the lengthy hardware refresh cycles inherent in the OT environment.

Modern security requirements challenge the traditional factory isolation approach

OT systems have typically approached security through network layer isolation (Purdue network/ISA95 hierarchy model), full network or internet disconnection, and/or completely locking down physical access. The need to connect factory production data to the cloud for global production optimization is causing a fundamental change in how OT systems are secured from cyberthreats. A further constraint is the difficulty of updating production-line equipment firmware and software. Enterprises must accommodate on-premises factory challenges while managing security, software updates, and changing regulatory reporting requirements. These challenges are encouraging enterprises to embrace a connected security model across the global digital estate.

Although a complete unification between the IT and OT environments goes too far in one direction, keeping them separate goes too far in the other direction. In times of continuous security threats, the need for improved operational efficiencies and increased centralized goals make it impractical for IT and OT to always be separate islands. A middle-ground framework is needed. According to a recent study by IoT Analytics,[1] 52% of respondents ranked having an IT/OT data platform as very or extremely important to enhancing operational intelligence. An adaptive cloud approach combines the best of both IT and OT, whereby OT maintains control and consistency as needed, but uses the IT framework for scale.

1. "How manufacturers prepare shopfloors for a future with AI,"

2. Adapting IT tools for OT scenarios

Since IT and OT processes have many differences, IT tools must adapt to OT requirements in specific ways to achieve the degree of digital transformation needed to support intelligent factory initiatives. Just using IT tools and processes for OT needs won't help OT-specific objectives, standards, protocols, and culture. Instead, manufacturers can adopt IT processes that use existing tools adapted for OT domains.

These adapted IT tools capitalize on the proven advantages of IT business processes and modern strategies, especially in areas such as security and observability. These IT business processes include using AI, scalability, agility, and modern strategies such as continuous integration and continuous deployment configuration as code, in addition to infrastructure fleet management. These processes and strategies not only reduce costs, but also can increase global manufacturing yield rates. In essence, factories need to maintain OT control to meet on-premises needs, while using IT practices for scalability.

Example: Instead of relying on USB sticks or third-party service technicians accessing on-premises manufacturing equipment for piecemeal installation of unverified software updates, use a more secure and efficient approach. This involves modifying the existing IT update delivery process to grant OT the final independent evaluation, timing, and approval rights. If the OT update process implements a few crucial changes to the policies associated with how a software package is applied, the same IT process can work well for OT, using established IT practices and tailoring them to fit OT requirements.

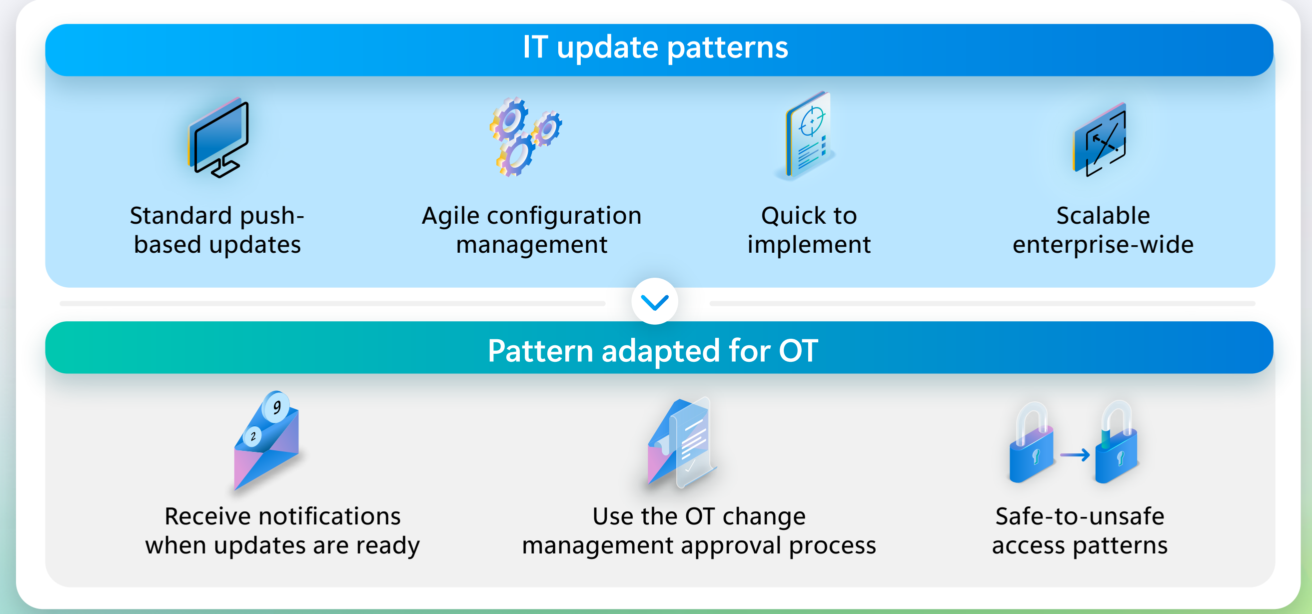

Figure 2

3. Creating a scale template with separate management, control, and data planes

To scale digital transformation, IT has established scalable practices for various aspects of the software life cycle by separating management, control, and data planes. This architecture design makes it easier for manufacturers to scale intelligent factories. Each plane has elements that run both locally at production sites and centrally in the cloud to accommodate local operational requirements on a global scale.

Management plane

The management plane provides global, centralized visibility and resource governance capabilities. It maintains a global inventory of resources under management, enables policies for resource utilization by applications, consumes health information about resources for reporting and analytics, and facilitates the resources' use of security and identity services.

When deploying applications, the management plane enforces common standards, access controls, and policies across distributed environments. It provides the ability to set global policies and necessary configurations that apply across all applications, APIs, and microservices. Governance policies can be applied by application group, type, geolocation, and so on.

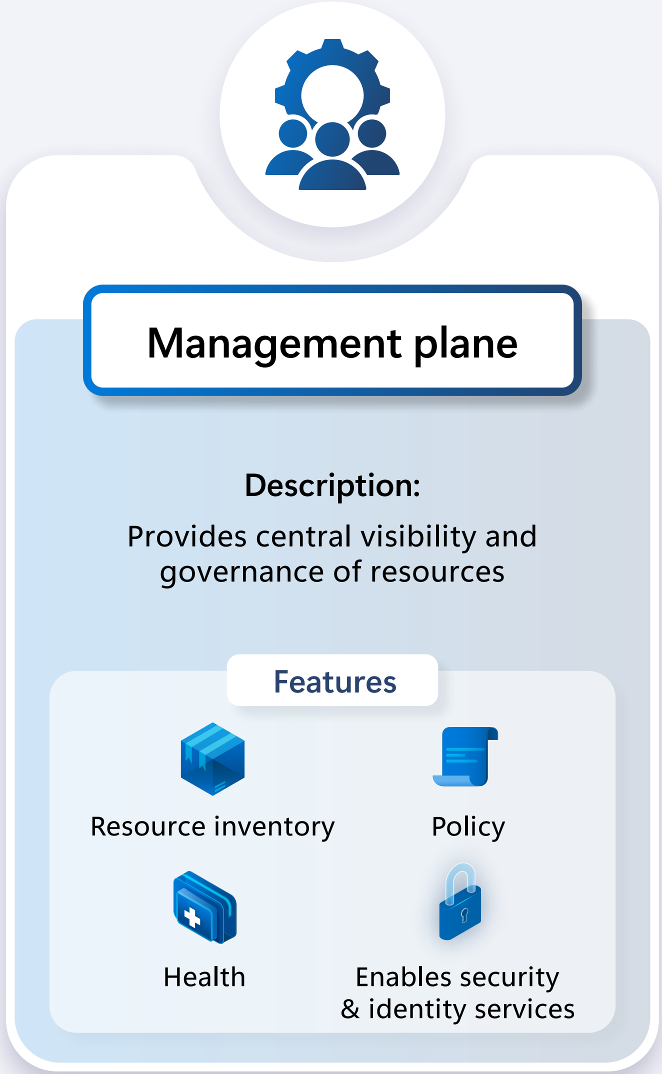

Figure 3

The management plane simplifies complex combinations of control plane operations, provides visibility and insight into application performance, and streamlines the configuration of the control plane for easier scaling, observability, security, and resilience. In the era of modern applications, many groups can benefit from a robust management plane, including app development teams, network operations, security, compliance teams, and teams responsible for specific line-of-business functions and operational capabilities.

Control plane

The control plane acts on policies, applications, and resources, based on what's specified in the management plane, and enables automation. In this taxonomy, the control plane includes resources used by applications (such as compute nodes, virtual machines, containers, and storage), manages application lifetime and workload scheduling, and provides platform-level resilience to system events that might otherwise interrupt workloads. The control plane supports critical inner-loop functions for applications running in an environment. In cloud-based IT environments, management and control planes are typically designed and delivered as one system operating in a single, central location.

To meet manufacturing environment requirements, control plane capabilities must operate locally at the edge without relying on a persistent connection to a global enterprise-wide resource. Ensuring that a control plane can operate without an internet connection for up to 72 hours seamlessly reconnect helps prevent data loss.

Data plane

The data plane ingests, processes, normalizes, and analyzes data on the edge, in the cloud, or both. As the primary resource for enabling analytics and insights for use by all levels of an organization, the data plane is responsible for collecting information across various resources, normalizing and enriching that information with essential context, and providing a foundation for AI scenarios. The data plane collects, stores, and flows information in all directions—across resources within local sites and published from local sites to global resources for visibility. To deliver enterprise-wide use cases—from the production floor to the boardroom—elements of the data plane need to exist and run at local operational sites and central locations to provide messaging, visibility, context, and insights as appropriate.

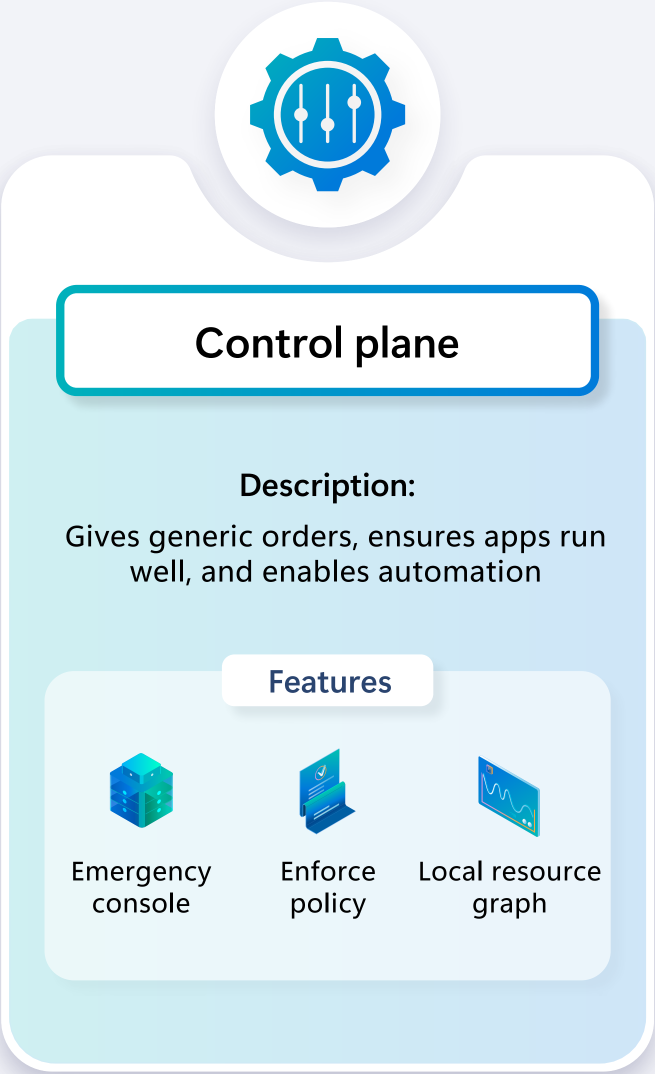

Figure 4

Data plane continued

User experience, latency, and all other key application performance metrics depend on a responsive, reliable, and highly scalable data plane. Policies defined using the management plane, such as service-level agreements (SLAs), scaling or behavior triggers (such as retries and keepalives), or horizontal scaling information, are all held within the data plane. This fundamentally makes the data plane the layer that applications depend on for instructions on what to do and how to behave. To effectively serve distributed manufacturing use cases, the data plane must make this information reliably available on the edge at operational sites without relying on a persistent connection to a global enterprise-wide resource.

From [a recent IoT Analytics survey](#),[2] the head of production planning and quality control at an electrical equipment manufacturer stated the following:

> **A strategic approach to digital transformation involves feeding AI applications from a common data layer. This can provide vast improvement to manufacturing processes with AI use cases like AI-based predictive maintenance.**

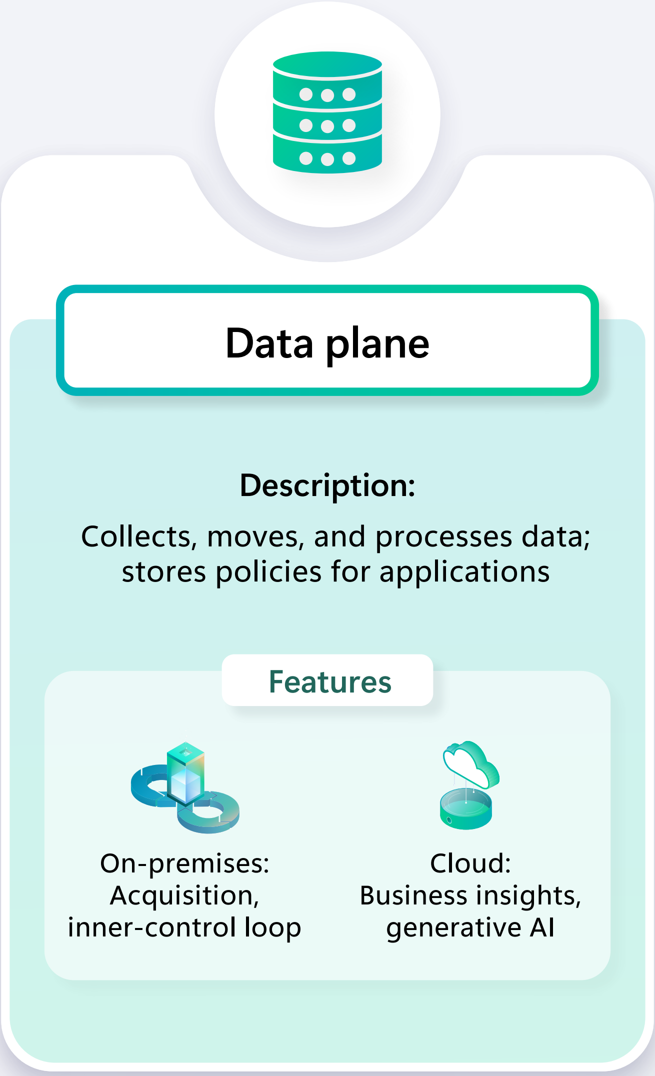

Figure 5

2. "How manufacturers prepare shopfloors for a future with AI,"

4. Applying a scale template to the factory floor

The approach of layering management, control, and data planes has been successful for service workloads in traditional IT environments. Still, that approach is difficult to apply in OT environments due to the technology gaps required to support the unique physical and organizational constraints of OT environments. Microsoft is working to build a global template that can be implemented at all sites to allow operators to create a single, consistent point of view for all sites that provides quick insights into status, risks, and workload rollout.

To visualize this pattern, Figure 6 shows how manufacturing lines and sites can achieve a local, independent, optimal state in the context of a connected group. At the same time, a pattern for a global management experience with local runtime control emerges alongside a data pipeline to deliver valuable insights at every level of the organization.

The benefit of plane separation increases when the management plane can easily scale from a single site to multisite, global deployments, using the same orchestration framework for all sites, while respecting site-specific nuances such as data sovereignty and data plane workloads.

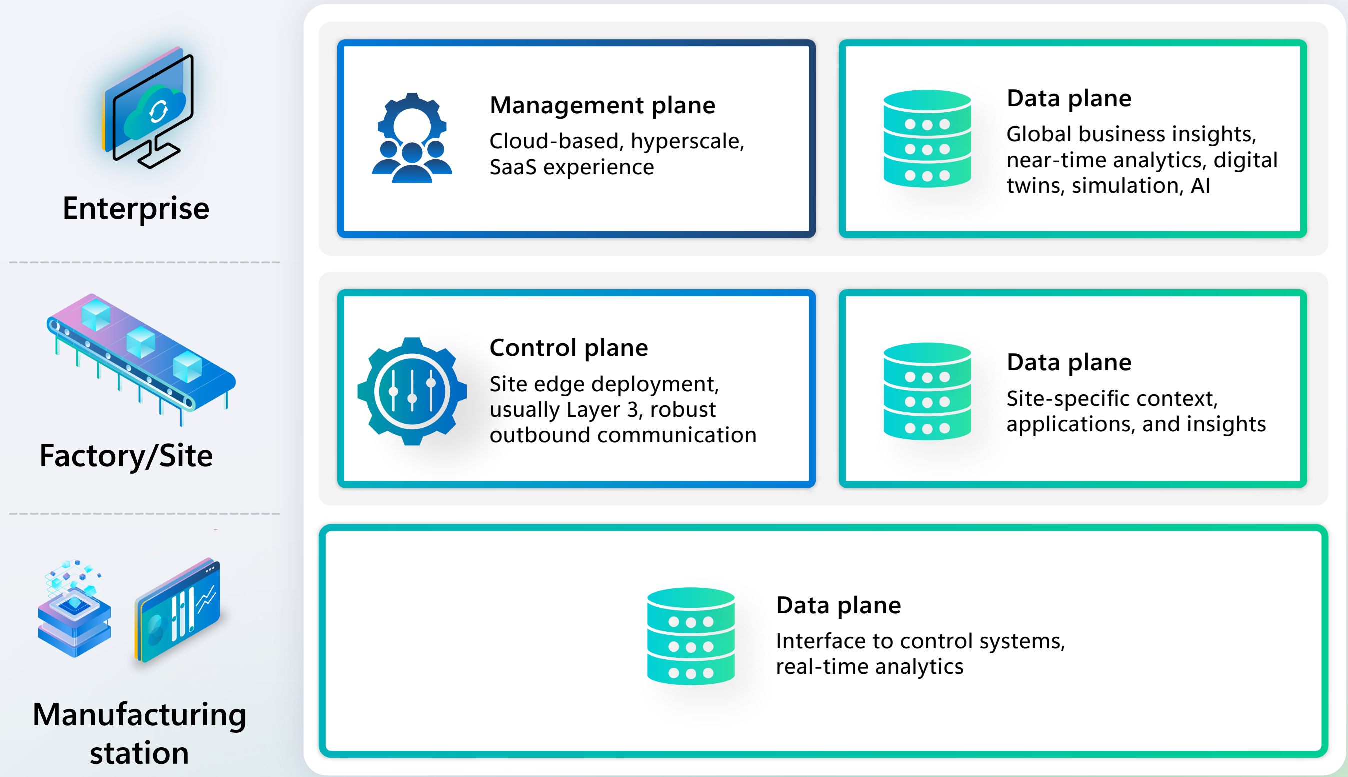

Figure 6

Managing sites as individual entities introduces the risks of inadvertent misconfiguration, poor software hygiene, and hidden inconsistencies, which create security vulnerabilities and drags on efficiency. Instead, manufacturers can use the plane separation approach to meet site-specific needs without creating silos.

Customers should deploy workloads on-premises that support core operations that need to be hosted and executed locally. Typically, these are workloads where safety and latency are critical, such as those providing functional safety or resilient business operations. At the same time, workloads that focus on scheduling, planning, or other complex tasks that pertain to the future operation of the site but aren't relevant to production currently in progress at the site, should be moved to a centralized environment.

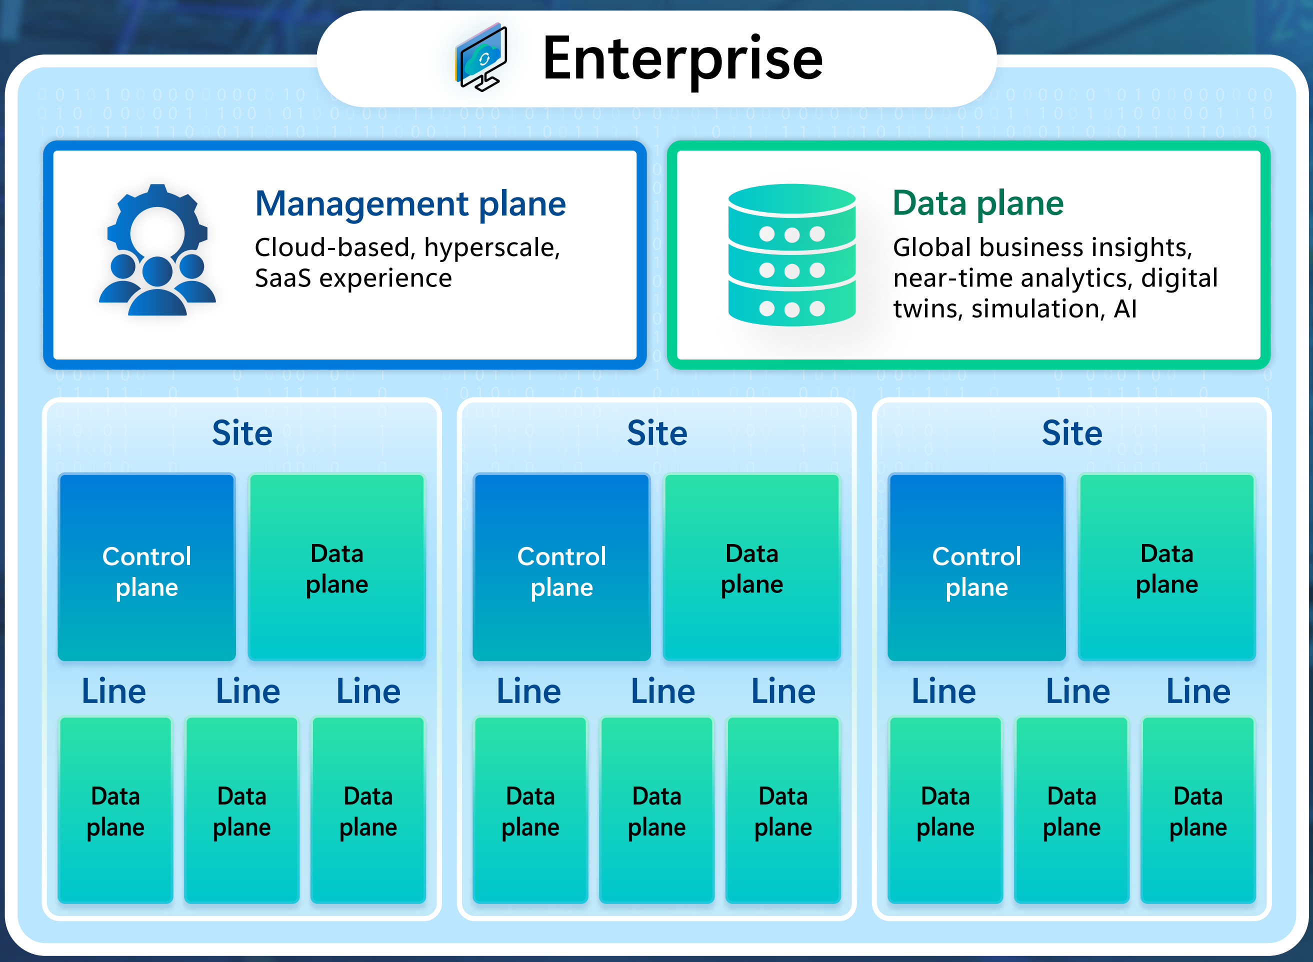

Figure 7

5. Adopting open standards for interoperability

All three planes must embrace open standards, such as OpenTelemetry (OTel), Message Queuing Telemetry Transport (MQTT), and Open Platform Communications Unified Architecture (OPC UA), for data collection, communication, and modeling. This will enable interoperability, compatibility, and innovation across a diverse set of industrial products, protocols, and solutions. Kubernetes, the industry standard for container orchestration, should be adopted as a local control plane standard for consistent and unified application and resources.

According to the recent study by IoT Analytics,[3] there's a strong preference for Kubernetes-based containerization in manufacturing, and Azure Kubernetes Service (AKS) is the most widely used container orchestration standard in both cloud and edge environments. Some 55% of respondents indicated that containerized software could significantly or extremely mitigate reliability and uptime challenges, while 53% indicated that it could do the same for cybersecurity challenges.

Software open standards enable interoperability, compatibility, and innovation among different software products and platforms.

These open standards specifications are publicly available and developed by consensus among multiple stakeholders, allowing software developers to scale faster and helping customers reduce costs, avoid vendor lock-in, and increase choice and quality. A scalable, future-proof strategy will embrace open standards to gain benefits such as:

OTel for scalable observability

OTel offers a set of technologies and practices for collecting, processing, and analyzing data from distributed systems. Its goal is to provide a consistent and comprehensive view of the performance, reliability, and behavior of complex applications and services. OTel supports various data types, such as metrics, traces, and logs, and integrates with multiple platforms and frameworks.

Using OTel, developers and operators can gain more visibility and insight into their systems, troubleshoot issues faster, and optimize their performance. Integrating software deployments into the control plane is easier with OTel because it inherits the management plane. Deployments can also run on multiple platforms without affecting visibility and operational overhead. Microsoft is working with the OTel community to add support for asset health models, in which each entity has health status and alert levels.

3. "How manufacturers prepare shopfloors for a future with AI,"

ADOPTING OPEN STANDARDS

MQTT with standardized message formatting for data and service communication

MQTT is an industry standard maintained by the Organization for the Advancement of Structured Information Standards (OASIS) consortia. It serves as a lightweight, publish-subscribe and queuing protocol for communication between edge services and the cloud. MQTT is widely used in edge and edge-to-cloud scenarios because it's efficient, secure, and scalable. Interoperability is only possible when message formats and message payload encodings rely on open standards to structure message data for simplified parsing and common understanding. Formats like JSON, Apache Avro, and Apache Parquet are open-source licensed and used in many different solutions for formatting and persistence of service communication.

OPC UA as a common language for asset-based business insights

OPC UA is an industry standard for industrial interoperability that offers many benefits for data exchange, data modeling, security, and scalability. OPC UA enables assets and services from different vendors and platforms to communicate seamlessly and reliably using a discoverable data model and interface. Insight generation becomes significantly easier to establish and maintain when using this standard. Not every asset can communicate via OPC UA, so it's essential to use messaging brokers—components that convert messages from a variety of proprietary asset interfaces and data models into standard IT protocols. Messaging brokers can help enable communication with the broad range of assets in a manufacturing environment.

Kubernetes as the control plane standard

Kubernetes, a Cloud Native Computing Foundation (CNCF) project, is the de facto enterprise standard for container orchestration because it offers several benefits for deploying and managing applications at scale. Kubernetes allows developers to package their applications into application bundles ("containers") that can run on any supported platform and provides authoring for deployment, scaling, and high availability of these containers. Kubernetes also enables service discovery, load balancing, networking, security, monitoring, and logging.

Using Kubernetes as the local control plane provides a common, vendor-agnostic, and scalable path for system management. Combined with the open-source Akri CNCF project for managing downstream assets securely, Kubernetes provides an enterprise-ready framework for managing remote locations. It's widely adopted by enterprises, cloud providers, and open-source communities. And it has a rich ecosystem of tools and extensions that enhance its functionality.

By employing standards in a consistent framework, organizations can more easily roll out assets and equipment, flow data to the right places, more quickly develop new use cases, reduce costs, and more rapidly scale across multiple operational sites.

6. Building thoughtful user experiences for IT and OT

Experiences must be tailored to fit how individuals in specific user roles participate in an organization across both OT and IT. OT roles such as plant managers, production engineers, and operations staff focus on local decision-making at the factory or plant level. In contrast, regional production managers or operations executives work at a regional or enterprise-wide level. IT roles, such as enterprise IT professionals or cloud architects, typically focus regionally or company-wide. However, their skills are increasingly needed to support IT scenarios in factory environments or production lines.

OT-focused experiences

Local OT roles, primarily responsible for manufacturing at a factory site, require experience that focuses on improving operational efficiency, quality, and throughput. Factory managers will be interested in the overall output and efficiency of the site. Production engineers may need to coordinate how various machines produce critical data to feed production reports. Line operators require near real-time performance information. As the number of solutions these roles work with continues to increase, OT user experiences should remain focused on local operational requirements without burdening local OT roles with the technical complexity of the underlying systems.

This can be accomplished through protocol-based communication between OT applications and applications consuming or processing data from physical equipment. The latter applications are built by industrial equipment vendors, ISVs, integration partners, or customers.

While most OT applications run within a local site using data from on-premises machines or applications, enterprise or regional OT roles should be able to gain controlled access to that data in appropriate ways. Given the multisite nature of enterprise OT tasks, these user experiences are most appropriately served by cloud-based analytics and by tools that combine local data with enterprise-wide context to inform broader decision-making.

IT-focused experiences

Enterprise IT roles are proficient at operating with cloud-based functions and tools, and can help manage local site systems and workloads. For example, they can use enterprise-wide identity and access management tools, such as Microsoft Entra ID or Azure RBAC, to control local-site user access to applications and data.

IT-focused experiences continued

IT is an increasingly strained resource at the local level as infrastructure technology, application solutions, and use cases expand quickly.

Technologies like Kubernetes may be familiar to cloud developers, but IT professionals with skills to manage Kubernetes deployment are less common in local environments. Therefore, it's critical to provide experiences that enable local IT scenarios in a scalable fashion, empowering the same IT organization responsible for supporting the security and health of enterprise-wide technology to support a growing number of OT use cases.

IT is best enabled by consistent, unified experiences and tools to work with both global and local resources in a manner appropriate to their scope. These tools should provide the ability to update and maintain computing infrastructure, applications, and runtime environments using consistent views of site-specific health for repeatable management across multiple locations.

Partnering with Microsoft and an ecosystem that scales intelligent factory initiatives

Microsoft invests in product experiences that deliver operational efficiency and integrate ISV solutions effectively. The Azure adaptive cloud approach provides the framework for all the services and solutions in a complex manufacturing environment to use the same management plane for global visibility, the same control plane for local control, and the same data plane to balance both global and local needs. As shown in Figure 8, this open, modular approach helps manufacturers consolidate disparate teams, decentralized environments, and complex systems into a unified, edge-to-cloud-based framework of edge and cloud services.

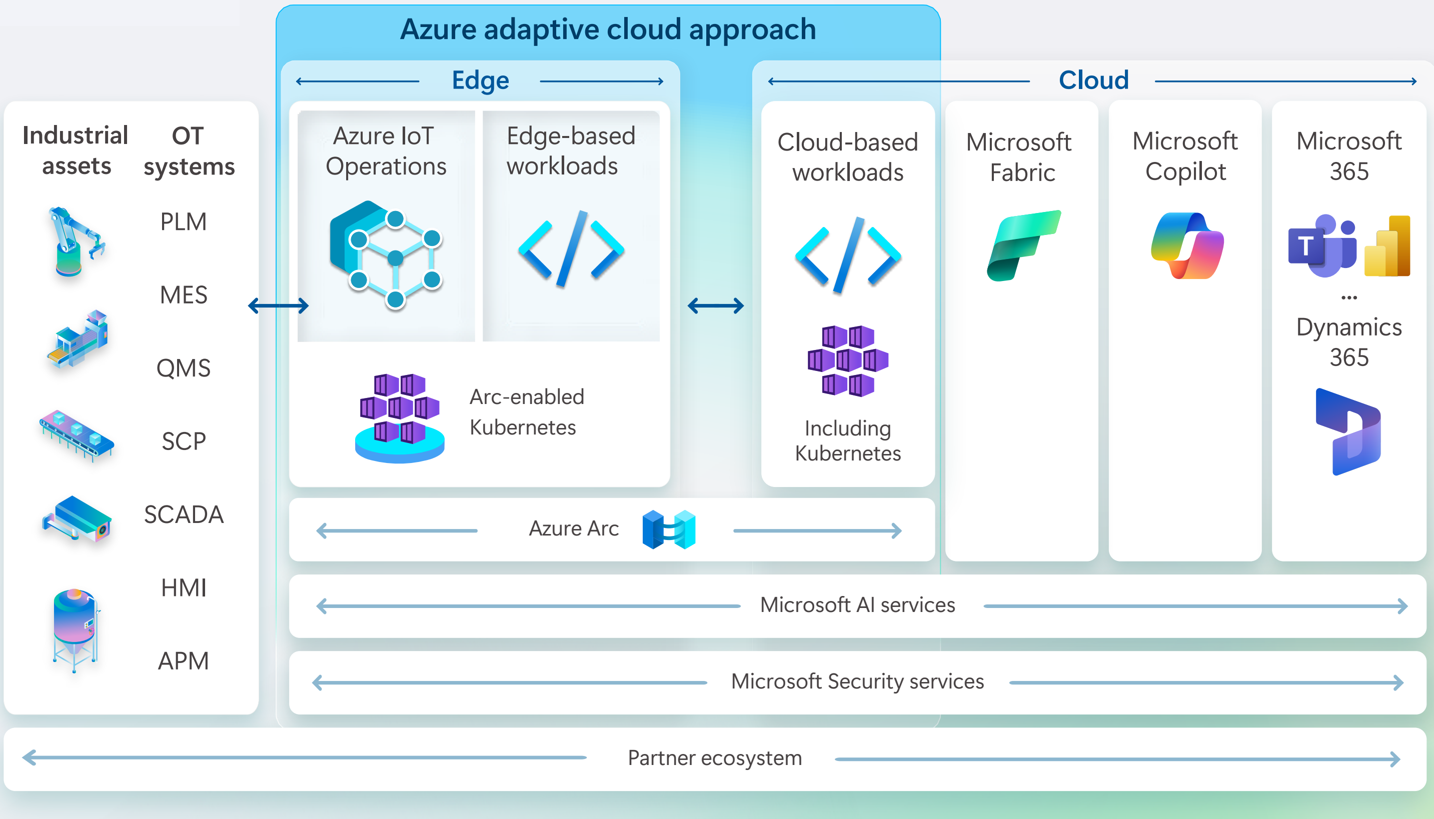

Figure 8

Azure is the foundation for an adaptive cloud approach. Every workload deployed or managed by Azure, such as virtual machines, databases, networks, or applications, is represented by a resource object that defines its properties and relationships. These resources are stored in a hierarchical structure in Azure Resource Manager. This acts as the single source of truth for all Azure resources and provides a consistent and unified way to create, update, delete, and query any resource across all Azure services and regions. It also enables features such as role-based access control, tags, policies, and templates that simplify and secure the management of complex solution environments.

Azure Arc operates similarly to a bridge between cloud and edge. Using Azure Arc, customers can extend the reach of Resource Manager to include resources hosted outside of Azure datacenters, including servers, workloads, and IIoT solutions running on the shop floor. Because manufacturers can define these resources outside the datacenter within Resource Manager, even for multicloud environments, services such as system health monitoring, security, and others can be easily repeated and applied across a globally distributed digital estate. Combining the global management plane with the control plane orchestration allows a production site to become as easily manageable as resources in a hyperscale datacenter. The management plane can govern a global list of resources and users, even when some resources reside locally at operational sites.

Kubernetes is a key technology for enterprises that need to manage software containers at scale. Software containers package applications uniformly for easier and faster delivery, similar to how shipping containers are used to package manufactured goods. Kubernetes orchestrates these containers to run workloads in any environment.

Azure offers a managed Kubernetes service—Azure Kubernetes Service (AKS)—and supports third-party Kubernetes distributions. AKS simplifies setting up, deploying, and connecting Azure Arc–enabled Kubernetes clusters. Azure Arc–enabled Kubernetes extends Azure management capabilities by connecting Kubernetes clusters to Azure. As a result, the Kubernetes workloads can centrally govern policies, roles, and identities while each site maintains local control. Any resource can be registered to the Kubernetes control plane for consistent management. This framework provides a scalable, vendor-neutral system for managing remote locations.

Azure IoT Operations ingests data for multiple applications running at the edge. It enables near-live asset data to be normalized and used at the edge, and then transfers only what's needed to a centrally managed data lake for global analysis and insights. Microsoft Fabric is a comprehensive, software as a service (SaaS) data integration and analytics platform that enables organizations to ingest, process, and analyze data in the cloud from various sources, including Azure IoT Operations. Together, these two technologies create a data plane that runs from edge to cloud. Not every dataset must or should be uploaded to the cloud; instead, the data plane determines which data to upload in order to balance local resiliency with global scale.

To learn how each of these services is implemented in a real-life scenario, refer to the extended technical version of this white paper.

Microsoft is committed to helping manufacturing customers realize the promise of Industry 4.0 through adaptive cloud technologies. The goal is to deliver the most value with the least disruption to operations. As part of that commitment, we have built a robust partner ecosystem with specialized skills and expertise to help industrial customers benefit from an adaptive cloud approach. This ecosystem includes:

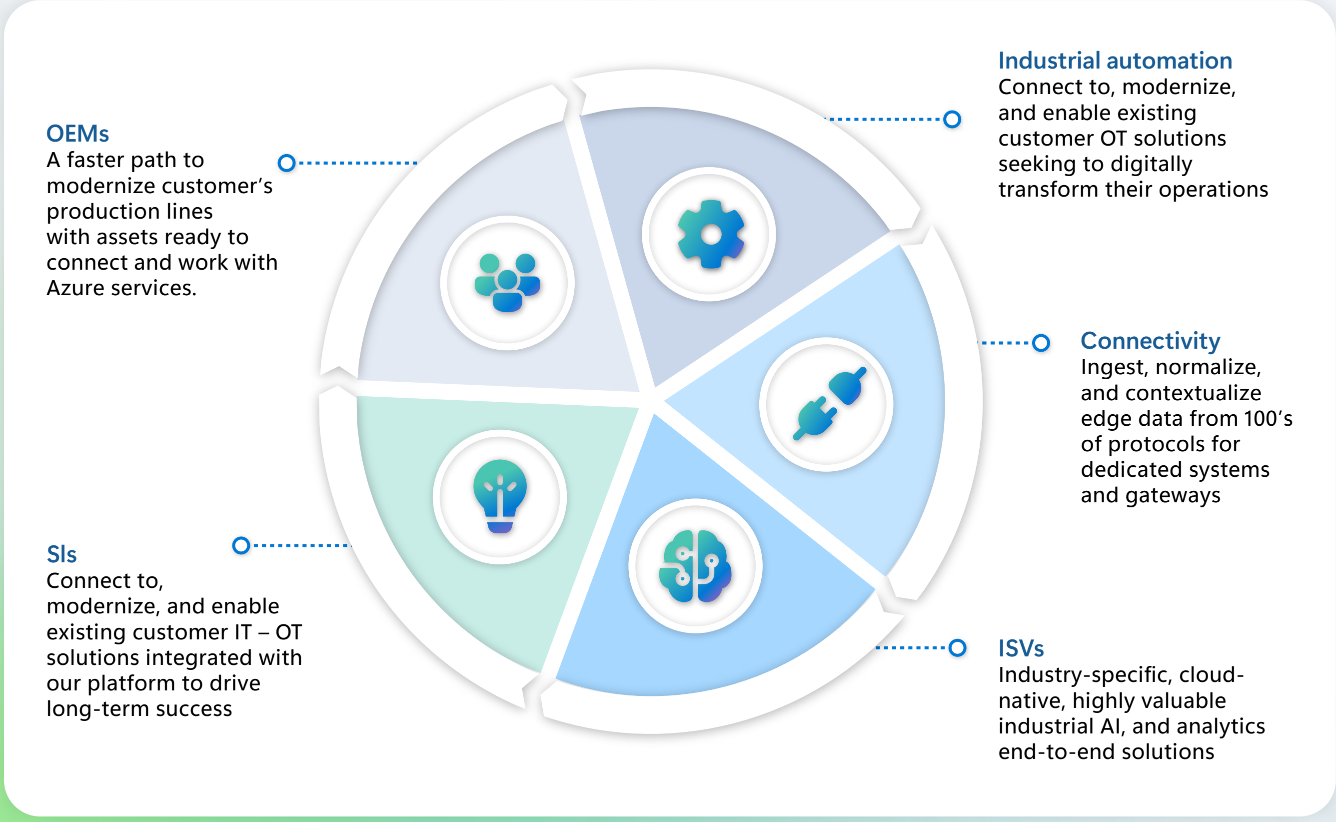

Figure 9

By taking an adaptive cloud approach, these partners can efficiently deploy value-added services and solutions that help customers:

- **Lower deployment costs and simplify infrastructure integration** by using a single control plane for all edge infrastructure, thus simplifying edge provisioning, management, monitoring, and updates.
- **Simplify data integration** by using out-of-the-box data services to achieve edge and cloud data interoperability.
- **Deploy any application or workload** across multiple orchestration and edge platforms and increase adoption of repeatable software templates and solutions on the shop floor.
- **Enable production AI at scale** to increase uptime, throughput, and quality and to unlock operational efficiencies and savings.

Real-world examples of how partners are delivering these benefits to their customers include:

DXC Technology, a trusted Microsoft partner for over 30 years, excels in both enterprise and operational technology with extensive global coverage. They understand that customers need future-ready, scalable architectures to quickly realize value from data, which is essential for accelerated industrial transformation. DXC builds fully integrated solutions and managed platforms, leveraging their industry-leading modular accelerators to help customers save costs and time while creating secure solutions. By adopting an adaptive cloud approach, DXC is at the forefront for driving momentum in AI-driven use cases, pushing the boundaries of innovation and performance.

Rockwell Automation is a global leader in factory automation and industrial digital transformation. With over 100 years of industry expertise, Rockwell excels at delivering scalable solutions for modern manufacturing operations. Partnering with Microsoft, Rockwell enables a unified technology stack that bridges IT and OT, creating agile, replicable, and scalable solutions for the factory of the future. This adaptive cloud approach allows quick deployment, seamless management, and easy replication of high-value solutions across multiple factory sites and lines. Generative AI copilot agents empower manufacturers to make smarter, faster decisions, boosting throughput and efficiency while accelerating operational modernization.

Advantech is a global market leader in hardware, IoT intelligent systems, and edge software. Their solutions enable smarter machines, optimize production processes, and drive digital transformation across various industries. Recognizing customer need for seamless integration of digital and physical operations, Advantech leverages the Microsoft adaptive cloud to facilitate IT/OT convergence.

To learn more about partners taking this approach, watch these videos from Avanade, Tata Consulting, Mesh Systems, and MaibornWolfe.

Learning from customer examples

Here are just a few examples of how industry-leading enterprises are using an adaptive cloud approach to scale their intelligent factory initiatives.

Chevron

Chevron is committed to continuously improving safety and protecting the environment leveraging technology innovation. Minimizing leaks is a key focus area, which has been supported by on-site visits in certain locations or through centralized monitoring. Leaks can also cause production shutdowns. Therefore, the ability to improve leak detection is key to Chevron's overall goals of enhancing safety measures, safeguarding the environment, and reducing operating costs. To meet these business priorities, Chevron aims to develop a scalable, reliable, and cost-effective leak detection solution built on an Azure-based, standardized Edge Platform that increases automation and security. Here are the details of that solution in Chevron's words:

> **With assets and operations across the globe, Chevron is seeking a comprehensive, flexible platform for AI analytics and automated controls that can be managed from a single pane of glass. The global integration of many types and brands of devices to be monitored adds to the complexity of development. A solution is needed that supports open-source technologies and uses Kubernetes to sustain a "build once and deploy everywhere" strategy. Working with Microsoft, Chevron joined a preview of Azure IoT Operations, enabled by Azure Arc, to create an edge-to-cloud data plane that enables local data processing and analytics while at the same time providing a centralized approach to management and security. Using this solution, built with an Azure adaptive cloud approach, Chevron has the potential to enhance security and increase accuracy and reliability, thereby reducing the need to visit the sites as often. The capabilities may also assist in managing and monitoring more assets at a lower cost while gaining AI analytics to further improve operations safety and efficiency.**

ABB

ABB strives to continuously improve the technology it uses to optimize how things are manufactured, moved, powered, and operated to build a more sustainable and resource-efficient future around the globe. Remote condition monitoring is a key part of keeping those operations running efficiently, but the exponential growth in the number of sensors and data they collect has made it a challenge to extract value from the information in a timely way.

By moving its orchestration and operation of remote condition-monitoring applications to Microsoft Azure Kubernetes Service, ABB has a single end-to-end Kubernetes solution that helps unify edge and cloud assets to streamline application development and simplify deployment and management. As a result, ABB needs to create only one version of applications that can run on Windows, Linux, and in the cloud.

This significantly reduces development and deployment time and frees resources to create even more effective AI-driven data analysis solutions from all the rich data collected from large manufacturing operations. Reduced administration requirements also allow ABB to handle a rapidly expanding edge ecosystem with a shrinking number of available workers.

> **“We need to deploy only one application regardless of where it runs, whether on a minimally configured Windows IoT device, a large machine with Linux, or in the cloud.”**

Viswanathan Ramakrishnan, Vice President,
Release Engineering & SaaS Operations,
ABB

> **“By integrating modern technologies and advanced analytics into our devices, we are poised to provide customers with a forward-looking and seamless user experience in edge and cloud.”**

Anshul Arora, Lead Product Manager
Genix Datalyzer and NextGen Condition
Monitoring, ABB

Xignux

Xignux is a leader in the energy and food industries. Based in Monterrey, Mexico, it manages a variety of companies that energize life and society to contribute to a better world, thanks to the hard work and talent of more than 31,000 employees in Mexico, the United States, and Brazil. As part of its digital transformation, Xignux is seeking to optimize and standardize processes, improve overall equipment efficiency (OEE) and product quality, reduce waste and costs, and increase productivity across its diverse group of businesses.

It recently worked with Microsoft to create a proof of concept (POC), using Azure IoT Operations and Microsoft Copilot, to test cable quality in near real time into a business unit. When complete, the POC will guide the company's efforts to replicate and expand those capabilities across the enterprise to create a single source of truth for its production processes. Using these technologies, Xignux hopes to enhance operational efficiency, reduce the complexity and time required to manage on-premises infrastructure, and provide better insights into plant health and manufacturing processes across the enterprise.

Based on its initial experience with its POC using Azure IoT Operations and Copilot, Xignux sees the potential to:

- Enhance operational efficiency by providing frontline workers with timely, reliable data to expedite problem solving on the plant floor.
- Reduce the complexity and time required to manage on-premises infrastructure because OT teams will no longer need to go to every plant to add or upgrade components. They can make the change once and use Azure Arc–enabled Kubernetes to replicate it across all Xignux businesses.
- Gain better insights into plant health and manufacturing processes by breaking down information silos at the edge, thus facilitating better communication between OT and IT operations.

Once the POC is complete, Xignux hopes to create a standard architecture guide for devices running on the edge that can be replicated across the enterprise.

> **“Azure IoT Operations is a good starting point related to architecture running on the edge. We are expecting to use it to create a guide for architecture that we can replicate across the whole enterprise.”**

Jose Arturo Montiel Espinosa,
Industrial Architect for IoT,
Xignux

Conclusion

There's a need for a unified underlying architecture to ensure that technology decisions and investments made today will be sustainable over the long run. An adaptive cloud approach provides a framework that will guide manufacturers in efficiently and cost-effectively evolving their systems and operations to handle tomorrow's manufacturing demands. By unifying siloed teams, distributed sites, and sprawling systems into a single operations, security, application, and data model, manufacturers can use cloud-native and AI technologies to work simultaneously across hybrid, multicloud, edge, and the Industrial Internet of Things (IIoT) systems.

This standardized approach enables manufacturers to scale operations across the entire factory ecosystem to reduce complexity, increase security and agility, and free data from the silos of the past. This will drive better collaboration and decision-making and pave the way for manufacturers to apply AI across their enterprise to enable new business models and empower their workforce.

Learn more

More information about how manufacturers can use an adaptive cloud approach to accelerate and scale transformation is available at:

- Adaptive cloud approach website.
- A recent IoT Analytics study on the latest trends, technologies, and priorities for manufacturing companies wanting to efficiently manage their data to prepare for AI.
- Continue learning about the underlying architecture with the extended technical version of this white paper.
- Deploy, configure, and use a realistic connected factory solution on Arc Jumpstart Agora.