DartPoints
DATA CENTERS

INDUSTRY GUIDE

WHY MANUFACTURING RESILIENCE DEMANDS MORE THAN A CLOUD STRATEGY

How manufacturers should think about hyperscale, regional infrastructure, and on-prem environments when uptime, recovery, and production continuity are on the line.

Manufacturing infrastructure decisions are no longer just IT decisions. They affect plant uptime, recovery speed, production continuity, and operational risk across the business. This guide looks at three core infrastructure approaches and where each one fits in a modern manufacturing environment.

Who this is for

For manufacturing IT, operations, and infrastructure leaders evaluating resilience, workload placement, and modernization strategy.

> **“Manufacturers aren’t trying to solve infrastructure problems. They’re trying to protect production continuity while modernizing operations without putting them at risk. That requires a completely different way of thinking about infrastructure.”**
>
> **Matthew McKee**
> VP, Solution Architects, DartPoints

Why Infrastructure Risk Looks Different in Manufacturing

In manufacturing, IT and OT systems need to operate flawlessly together because even minor disruptions can quickly escalate into production delays, shipment issues, and significant financial loss. In fact, Siemens’ *The True Cost of Downtime 2024* report found that:

> **“Manufacturers can lose up to $253 million annually due to unplanned downtime.”**

This is very different from traditional enterprise IT, where downtime may mean temporarily unavailable applications or reduced employee productivity.

“In a knowledge-worker environment, a 20% loss in productivity might still be good enough because employees are still functioning at 80%. In manufacturing, a 20% loss in production is a hard money cost. You’re not producing, inventory gets backed up, shipments get delayed, and suddenly there’s a domino effect across the entire supply chain.”

This operational pressure forces manufacturing teams to evaluate infrastructure decisions through very different lenses:

- Operations teams focus on recovery tolerance, plant uptime, and the speed at which production systems can recover without disrupting throughput.
- Finance leaders focus on downtime costs, infrastructure predictability, and the tradeoff between short-term savings and long-term resilience.
- IT teams focus on workload placement, integration, security, resilience, and infrastructure control.

While each team evaluates infrastructure differently, those priorities ultimately converge around operational disruption. That growing risk is also changing how manufacturers approach security and compliance decisions.

How Security and Compliance Affect Workload Placement

As production environments become more connected, operational disruption becomes more valuable to attackers. Fortinet found that

> **“Manufacturing accounted for 17% of all targeted attacks.”**

That’s why security and compliance requirements now play a much more direct role in manufacturing infrastructure decisions. They influence where workloads can safely operate and how environments need to be segmented, secured, and recovered.

Poorly segmented environments don’t just create cybersecurity exposure. They can also complicate recovery efforts, expand ransomware blast radius, and increase the likelihood of operational disruption spreading across production systems.

In practice, that means workload placement decisions increasingly need to support requirements such as:

- OT segmentation that limits lateral movement across production systems
- Ransomware containment that reduces operational impact and recovery time
- Auditability and evidence readiness for investigations and reporting
- CMMC and industry compliance requirements for sensitive data
- Recovery objectives for production-critical workloads & operations

These requirements are changing how manufacturers evaluate infrastructure environments. While some workloads benefit from hyperscale scalability, production-critical systems often require tighter operational control, geographic proximity, and stronger segmentation.

Resilience Is Not the Same as Continuity

Manufacturing environments cannot evaluate resilience through uptime metrics alone. In manufacturing, operations can be disrupted long before systems go offline. Systems may still appear available while degraded latency, delayed telemetry, or inconsistent performance interfere with production.

> **“On average, organizations experience 86 outages per year.”**
>
> **Cockroach Labs,** State of Resilience 2025

For example, an air-quality monitoring system may technically remain connected during a network degradation event. But if telemetry delays prevent operators from receiving accurate environmental readings in real time, the system is no longer operationally reliable, even though it still appears “online” from an infrastructure perspective.

This highlights the difference between technical continuity and operational resilience. Manufacturing environments increasingly need infrastructure strategies built around graceful degradation: *the ability to continue operating safely and predictably as conditions deteriorate.*

Technical Continuity Focus	Operational Resilience Focus
Was the system available?	Did operations continue safely?
Did connectivity remain active?	Was telemetry still usable in real time?
Did applications stay online?	Could production continue without disruption?
How quickly was uptime restored?	How much degradation could operations tolerate before failure occurred?

This focus on operational resilience requires manufacturing leaders to evaluate workload placement through the lens of operational impact, including how systems fail, how quickly critical functions can recover, and how dependencies between environments could increase disruption during an incident.

The Three Infrastructure Models Manufacturers Are Balancing

Once manufacturers account for operational resilience, recovery tolerance, security requirements, and production dependencies, the infrastructure conversation changes entirely.

Successful manufacturers are not trying to force every workload into a single environment. Instead, they are evaluating which systems require scalability, which require tighter operational control, and which need to remain closest to production operations to minimize risk and disruption.

That shift is leading manufacturers to balance workloads across three primary infrastructure models. Each environment solves different operational problems, so each should be considered for the workloads it handles best.

1. **Hyperscale cloud**
2. **Regional colocation** or **private hosted cloud**
3. **On-prem infrastructure**

Consideration	Hyperscale Cloud	Regional Colocation	On-Prem Infrastructure
Best for	Elastic scale, innovation velocity, global services	Low-latency workloads, data locality, predictable performance	Real-time control systems, air-gapped environments, system isolation
Strengths	Broadest service catalog, rapid provisioning, global reach	Performance consistency, data sovereignty, dedicated resources	Complete control, no external dependencies, works offline
Considerations	Internet dependency, variable latency, egress costs	Capacity planning, fewer instant scale options, regional footprint	Higher CapEx, longer refresh cycles, ongoing maintenance

1 When Hyperscale Cloud Makes Sense in Manufacturing

Manufacturing environments may operate under stricter operational constraints than traditional enterprise IT, but that does not reduce the value hyperscale cloud platforms can provide when workloads are placed appropriately.

Platforms like Amazon Web Services, Microsoft Azure, and Google Cloud continue to play an important role in manufacturing infrastructure **by providing scalability, centralized management, and global accessibility.**

That often includes:

- **Backup and archival storage**
- **Collaboration platforms**
- **Scalable, non-critical business applications**
- **Enterprise analytics**
- **Disaster recovery environments**

The Real Challenge Is Workload Placement

Too many infrastructure decisions are still driven by standardization goals, migration targets, or short-term infrastructure savings. A better approach is to **evaluate workloads operationally before deciding where they should live**.

Answering these questions helps manufacturers determine which workloads benefit from centralized infrastructure and which become more difficult to manage at a distance.

Moving to the cloud doesn't automatically translate into better savings either. Flexera's 2026 State of the Cloud Report found that managing cloud spend remains the top cloud challenge for enterprises, so it's still something manufacturers have to monitor and manage as they scale.

- Does latency affect production timing or operational responsiveness?
- How quickly does the system need to recover before operations are affected?
- Does the workload generate large volumes of telemetry or real-time operational data?
- Would connectivity disruption create downstream production consequences?
- Can the system tolerate operating at a distance from the production environment?

2 Regional Colo: The Middle Ground Between Hyperscale and On-Prem

Some manufacturing workloads cannot operate predictably at a distance. Systems that depend on real-time responsiveness, continuous telemetry, or rapid recovery coordination can struggle in highly centralized environments where latency and operational distance introduce additional complexity.

At the same time, keeping every workload fully on-prem can be hard as manufacturing environments grow more distributed, interconnected, and data-intensive across facilities.

How Data Gravity Changes Infrastructure Decisions

The more connected operations become, the harder it becomes to ignore data gravity. Moving large volumes of telemetry, analytics, and operational data across distant centralized environments can be slower, more expensive, and harder to manage consistently at scale.

That's why regional infrastructure remains valuable. It provides a middle ground between hyperscale scalability and on-prem control by allowing manufacturers to keep infrastructure geographically closer to operations while still supporting hybrid cloud architectures, managed services, and scalable interconnection across sites and providers.

Regional Infrastructure Supports Operational Responsiveness

These environments are often well-suited for workloads that require:

- **Lower latency and more consistent responsiveness**
- **Faster recovery coordination during disruptions**
- **Greater operational visibility across production environments**
- **Stronger segmentation between critical systems**
- **More predictable data movement and infrastructure costs**

That proximity becomes especially valuable for latency-sensitive systems that depend on consistent real-time responsiveness, where distributing infrastructure across multiple regions can reduce network latency by **as much as 35%.**

For manufacturers, the value of regional infrastructure lies not simply in proximity but in operational control. By reducing the distance between infrastructure and production systems, manufacturers can improve resilience, simplify recovery, and contain operational disruption more effectively when systems degrade or fail.

3 On-Prem Infrastructure Still Plays a Critical Role

Some manufacturing workloads cannot tolerate operational distance. To maintain safe and consistent throughput, production systems often rely on real-time communication between machines, sensors, controllers, and monitoring environments where even small delays can affect operational performance. In these environments, timing matters more than centralized efficiency.

For example, if AI systems are helping to drive logic-controller behavior on a production line, moving those workloads to a distant hyperscale environment can introduce numerous opportunities for latency, connectivity disruptions, or communication failures between the infrastructure and the production system.

When Predictability Matters More Than Centralization

Factory-critical workloads create additional constraints. Systems tied to PLCs, life-safety operations, and real-time production environments often cannot tolerate degraded responsiveness because delays can directly affect throughput, production timing, product quality, or safety conditions.

Recovery expectations also look different for these systems. Some workloads cannot wait for traffic rerouting, cloud failover processes, or large-scale recovery orchestration during disruptions because recovery delays immediately affect production continuity.

Why Telemetry-Heavy Systems Need Proximity?

Telemetry-heavy environments introduce another operational challenge. Some production systems generate such large volumes of real-time operational data that continuously moving telemetry into distant centralized environments becomes expensive, bandwidth-intensive, and increasingly difficult to manage efficiently at scale.

As manufacturing environments become more connected, data gravity becomes harder to ignore. Keeping telemetry processing and operational infrastructure closer to production environments can improve responsiveness, simplify data management, and reduce unnecessary infrastructure overhead.

On-Prem Infrastructure Still Matters

That is why on-prem infrastructure continues to play a critical role in manufacturing environments.

For manufacturers, the value of on-prem infrastructure is not simply ownership or control. It is operational predictability. Keeping critical infrastructure physically and operationally close to production environments helps manufacturers reduce latency exposure, maintain throughput consistency, and preserve operational continuity even if systems degrade or fail.

These environments are often best suited for workloads that require:

- **Deterministic, low-latency system responsiveness**
- **Real-time communication between production and controllers**
- **Local processing for telemetry-heavy operational environments**
- **Fast recovery for production-critical systems**
- **Tighter operational control over safety-sensitive workloads**

What Resilient Manufacturers Are Doing Differently

The manufacturers navigating modernization most effectively are treating infrastructure less like a IT standardization exercise and more like an operational dependency map.

They are asking harder questions about how workloads behave under pressure

What happens if connectivity degrades during production?

Which systems can tolerate latency, and which cannot?

How quickly does recovery need to happen before operations are affected?

Which workloads create downstream disruption if they fail?

This operational lens is changing how manufacturing IT leaders approach their infrastructure strategy. A workload that supports production scheduling carries different recovery requirements than a collaboration platform. A telemetry-heavy operational system creates different infrastructure challenges than a business application with limited real-time dependencies.

And a safety-adjacent workload requires a very different tolerance for latency, degradation, and recovery delays than a non-critical enterprise system.

That operational mindset changes how manufacturers evaluate resilience, workload placement, and infrastructure risk across their business. The conversation is moving away from “where should infrastructure live?” and toward “what does this workload require to keep operations resilient?”

Why “All-In” Infrastructure Strategies Often Fail

Consolidating workloads into a single infrastructure environment can be appealing for many reasons. Fewer environments usually mean simpler management, fewer vendors, and lower operational overhead. While this can simplify infrastructure management, it can also concentrate operational risk into a single failure domain. In manufacturing environments, that tradeoff can be costly.

Saving $10,000 per month by moving workloads into a hyperscale environment may look efficient from an infrastructure cost perspective. But if that decision increases production downtime by even a fraction of a percent, the resulting operational losses can quickly outweigh the infrastructure savings many times over.

Production continuity is often far more valuable than incremental infrastructure efficiency. In this context, “all-in” infrastructure strategy rarely makes sense in manufacturing.

> **“The moment infrastructure decisions become philosophical instead of operational, resilience suffers. In manufacturing, simplifying environments for the sake of simplicity often makes them more brittle, not more resilient.”**

AI Is Accelerating the Shift Toward Distributed Infrastructure

AI is adding even more pressure on manufacturers to modernize infrastructure environments. But the most important infrastructure decisions right now are not about pushing autonomous AI deeper into production operations, but about becoming operationally ready for the workloads AI will require over time.

> **"Today, as many as 96% of manufacturing firms are using AI to improve operational visibility across production environments."**

These production environments include:

- Predictive maintenance systems that detect equipment issues early
- Quality monitoring models that identify anomalies faster than manual inspection
- Forecasting models that help teams respond earlier to supplier and demand shifts

What connects these use cases is that AI is supporting operational decision-making rather than directly controlling production-critical systems. While AI may already be "good enough" for lower-risk business tasks, manufacturing environments still require a much higher standard for reliability before AI-driven decision-making can safely move deeper into zero-tolerance production operations.

For manufacturers, the smarter near-term goal should be becoming AI-ready rather than accelerating AI adoption without full understanding of all operational risks. That means consolidating operational data, building telemetry pipelines, and supporting AI workloads closer to production environments.

> **"The goal isn't to rush AI-driven autonomous decision making into the business today. It's to make infrastructure AI-ready so that when AI becomes reliable enough for mission-critical, zero-tolerance manufacturing environments, the business is ready to adopt it."**

Those requirements are also reinforcing the need for more distributed infrastructure strategies. AI workloads continuously generate and consume large volumes of operational data, making latency, bandwidth costs, and real-time responsiveness harder to manage once infrastructure moves too far from production systems.

The result is an infrastructure model built less around centralized standardization and more around keeping data, processing, and operational responsiveness aligned with production requirements.

Manufacturing Resilience Is Becoming a Workload Strategy

The manufacturers building the most resilient operations are rethinking infrastructure less as a technology stack and more as an operational system. The real challenge is not deciding whether cloud, colo, or on-prem infrastructure is "best," but understanding how different workloads behave when production conditions become unstable.

Some systems benefit from centralized scale and coordination. Others become harder to operate once latency, recovery dependencies, or data movement introduce too much operational distance between infrastructure and production environments.

And certain workloads remain inseparable from the factory floor itself because even small disruptions in responsiveness can affect throughput, timing, or safety conditions.

In this context, the most resilient environments will not necessarily be the most centralized or technologically uniform ones. They will be the ones designed to contain disruption, preserve operational responsiveness, and keep production systems functioning reliably as conditions change.