FNT

Confident but Wrong

Why Agentic AI for Digital Infrastructure Depends on Authoritative Knowledge and Data

CONTENT CERTIFIED
by Experts

Authors: Daria Batrakova & Oliver Lindner
June 2026

// simplify complexity

Table of contents

Executive Summary

Agentic AI – software that can reason over data, invoke tools, and carry out multi-step tasks with limited human intervention – is moving from experimentation into enterprise operations. In infrastructure and network domains, that shift is especially significant. Capacity planning, incident triage, change impact analysis, and compliance reporting are all candidates for increased automation. But these use cases depend less on model novelty than on the quality of the infrastructure data available to the system.

This paper argues that physical and logical infrastructure presents a grounding problem that general data platforms do not fully solve. Infrastructure decisions depend on current topology, validated relationships, and an accurate representation of the real world across facilities, racks, devices, ports, circuits, services, and dependencies. That information typically does not exist in a usable form inside a data warehouse alone. It is better represented in a maintained digital twin of the infrastructure itself: in DCIM, NRI, ITAM, and cable management systems that model the physical world with the precision that automated decision-making demands. In the telecommunications sector, TM Forum's Autonomous Networks initiative has formalized this dependency, establishing a widely adopted framework in which higher levels of network autonomy are explicitly predicated on the quality and completeness of the underlying network digital twin.

The strategic implication is straightforward: organizations that want to deploy AI safely in infrastructure operations need a governed, accurate, queryable infrastructure context layer. In practice, that usually means a digital twin supported by disciplined change processes, reconciliation, and integration.

This is not only an operations issue. Regulatory frameworks such as NIS2, DORA, the EU AI Act, NERC CIP, and the broader US federal continuous-monitoring model increase the value of current asset records, dependency visibility, audit trails, and traceability. These frameworks do not prescribe a digital twin. But for many organizations, a maintained digital twin is one of the most scalable ways to meet those needs while also preparing for more autonomous AI-enabled workflows.

FNT's position is that infrastructure AI will be adopted most successfully where organizations treat infrastructure data as a strategic asset rather than a documentation afterthought. Without that foundation, infrastructure AI will often be fast, confident, and wrong.

1. The Agentic AI Moment – and Why Infrastructure Is Different

The trajectory of enterprise AI has shifted. Early generative AI adoption focused primarily on assistive tasks: summarization, drafting, retrieval, and recommendations, usually with a human reviewer at the end of the process. Agentic AI extends that pattern. These systems can reason across multiple inputs, call tools, plan actions, and in some cases execute workflow steps directly.

In infrastructure operations, that points toward AI systems that can:

- **assess** available capacity across sites and recommend – or automatically execute – a resource allocation,
- **analyze** topology to identify bottlenecks and single points of failure, flagging risks before they surface as incidents,
- **trace** the dependency chain from a service degradation event back to a specific physical device or cable segment, and initiate a ticket or failover without waiting for a human to map the topology,
- **prepare or orchestrate** portions of a provisioning workflow,
- **assemble** compliance evidence from multiple operational systems, or
- **detect and escalate** infrastructure changes that create operational or security risk.

Each of these scenarios involves an AI system querying structured data about physical, logical, and virtual infrastructure, reasoning over it, and taking action. The quality of that action is bounded, absolutely and entirely, by the quality of the data the agent is reasoning over.

Recent enterprise guidance from Bain and others has emphasized that weak data foundations, governance gaps, and fragmented context remain major barriers to scaling agentic AI in production.[12] The issue is not unique to infrastructure, but infrastructure makes it harder because the data is highly relational, operationally consequential, and vulnerable to drift from physical reality.

What is not yet widely understood is how this challenge manifests specifically in the physical infrastructure domain, and why the data governance approaches designed for business systems cannot solve this problem – because infrastructure data has fundamentally different characteristics.

The AI shift in infrastructure operations

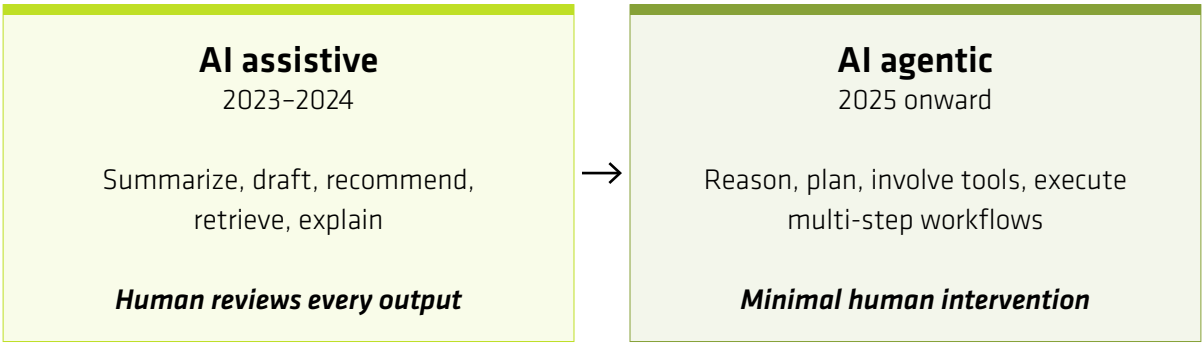

Infrastructure data quality determines whether agentic AI is reliable or hazardous

2. The Grounding Problem in Infrastructure AI

The term "grounding" refers to the process of connecting an AI agent's reasoning to accurate, current, real-world data. An ungrounded agent is one that reasons from assumptions, stale records, or incomplete models. In infrastructure operations, ungrounded AI is not merely inefficient – it is operationally hazardous.

Consider a concrete scenario. An AI planning agent is tasked with recommending additional rack space for a GPU cluster expansion in a colocation facility. The agent queries available infrastructure records, calculates space, power, and cooling headroom, and produces a recommendation: "Row G, positions 14–17 have sufficient capacity." The recommendation is generated in seconds, formatted as a change request, and routed for execution.

Except: Row G, positions 15–16, were reconfigured six months ago when a customer terminated their contract. The records were never updated. The physical space is now occupied by cross-connect panels, and the power circuit serving that row was rerouted. The "available" capacity exists only in the documentation, not in the facility.

A human engineer would have caught this in a site walk. An AI agent operating at machine speed, acting autonomously, does not catch it. The error is discovered when the technician arrives on-site with deployment equipment. The cost is not just the aborted deployment, it is the trust in AI-assisted planning that is set back by months.

This is the grounding problem for infrastructure AI. It is not a failure of the AI model. The model performed exactly as designed. It is a failure of the data foundation. Just as a human engineer requires a reliable source of objective truth before taking action, so does an AI agent – with one critical difference: the agent cannot pause, question, or draw on experience when data is incomplete or wrong.

What Makes Infrastructure Data Different

Infrastructure data is different from business data in at least three important ways.

1. **It is deeply relational.** The value lies not just in individual records but in the chain of relationships among facilities, rooms, racks, devices, ports, circuits, paths, services, and dependencies. A sale creates a record. A customer enrollment creates a record. The data and the event are coupled. Physical infrastructure does not work this way. Unlike customer records or product hierarchies, which can be represented as trees with clean parent-child relationships, infrastructure is a mesh: every element connects to multiple others across multiple layers, and those cross-layer relationships are as operationally significant as the elements themselves.

2. **It is operationally perishable.** Its usefulness depends on whether it reflects the current or near-current state of the environment. A cable gets repatched during an emergency change at 2 a.m. A power circuit gets rerouted to accommodate a new UPS. A fiber splice is added to resolve a fault. An OT device is swapped out by a field engineer. None of these events automatically update a management system unless the organization has disciplined processes and tooling in place to ensure they do.

3. **It requires validation against the physical world.** Knowing where a record came from – its data lineage – is not enough. For high-value use cases, teams need to know whether a record reflects planned state, implemented state, and when it was last verified against physical reality.

The Scale of the Documentation Gap

The gap between the physical state of the infrastructure and its documented state – the "as-built versus as-maintained" gap – is one of the most persistent and costly problems in infrastructure management. A 2026 economic study estimated that lifecycle inefficiencies associated with fragmented infrastructure documentation may exceed $20 billion annually within the United States and potentially more than $300 billion globally across built environments. That figure, drawn from a press-release-backed study rather than a primary regulator or top-tier analyst,

should be treated as directional rather than definitive.[3] The more defensible point is simpler: documentation drift is common, its operational cost is real, and as automation increases, the cost of drift rises.

For human-driven operations, this gap is managed through periodic audits, engineering judgment, and the accumulated knowledge of experienced staff. For AI-driven operations, it is not manageable in the same way. An agent cannot exercise judgment about whether a record might be stale. It reasons from what is there.

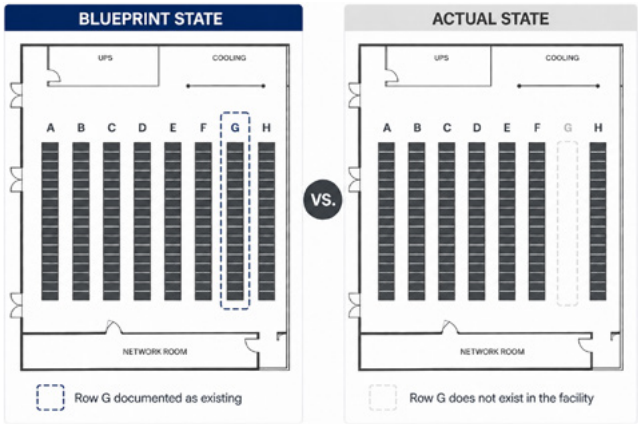

3. Why Data Lakes and Warehouses Are Not the Whole Answer

A reasonable objection at this point is: "We are building a central data platform. All infrastructure data will flow into our data lake / data warehouse / data fabric. AI agents will query that."

This architecture is valuable, and for many use cases it is necessary. It works well for telemetry, monitoring metrics, capacity utilization trends, incident logs, and other machine-generated operational data. AI agents can reason effectively over this data to identify patterns, anomalies, and trends.

But by itself it does not fully solve the infrastructure context problem. The issue is not that data lakes or warehouses are incapable. It is that they are not usually the authoritative operational system for validated physical topology and change-state reconciliation. Several limitations matter in practice:

- **The data model problem.** Data lakes and warehouses are optimized for tabular or semi-structured data – rows, columns, time-series values. Physical infrastructure is fundamentally a relationship graph: cables connecting ports, ports belonging to devices, devices sitting in racks, racks occupying space in rooms, rooms in buildings, buildings in campuses. Infrastructure questions are often graph questions: what depends on this path? Which services share this segment? What is the physical blast radius of this planned change? The topology does not normalize into tabular form without losing the structural information that infrastructure AI actually needs.
- **The provenance and validation problem.** Data governance in a data lake context typically means lineage tracking (where did this data come from?) and schema management (what format is it in?). For infrastructure planning, something different is needed: validation against the physical world. Has this record been confirmed accurate? When was it last verified? Does it reflect a completed change or a planned one? These are infrastructure lifecycle questions, answered by reconciliation processes, field verification workflows, and closed-loop change management, not by data lake governance tooling.
- **The staleness problem.** A data warehouse operates on refresh cycles. Infrastructure physical state changes continuously, and the timing of changes matters acutely for planning. An agent planning a network change needs to know the current connectivity, not the state as of the last batch refresh. A nightly or

periodic ingestion cycle may be acceptable for analytics but is insufficient for high-confidence operational decisions.

- **The relationship traversal problem.** When an AI agent asks, "what services are at risk if this fiber path fails?", it needs to traverse a dependency chain: fiber path → splice points → cable → termination → patch panel → cross-connect → logical circuit → network service → dependent applications → business services. This traversal requires a graph model with validated relationship data at every level. It is not a query that a data lake is designed to answer.

- **The ontology problem.** Perhaps the most underappreciated gap is the absence of domain ontology – the formal model of how infrastructure entities relate to each other and what those relationships mean. An AI agent reasoning about infrastructure needs to understand not just that Device A is connected to Device B, but what kind of connection it is, what it carries, what depends on it, and what the rules are for changing it. This knowledge – the schema of infrastructure relationships – cannot be derived from raw data. It must be encoded in the data model itself. Infrastructure management platforms built over years to model real-world networks carry this ontology as an intrinsic property: the relationships between physical ports, logical circuits, services, and business functions are defined, enforced, and queryable. A data lake ingesting snapshots from multiple systems carries none of this. The ontology must be rebuilt from scratch, a project that takes years and domain expertise most organizations do not have in-house.[15] FNT Command's component library – built over more than thirty years across hundreds of infrastructure deployments in data centers, enterprise networks, and telecommunications environments – carries this ontological knowledge as a product property: more than 75,000 predefined components with their relationships, attributes, and operational rules already encoded. That is not something a data lake ingests. It is something that takes decades of domain experience to build.

For these reasons, the stronger architecture is usually not "data lake or digital twin," but "data platform plus digital twin." The analytical estate remains important. The digital twin becomes the operational context layer for infrastructure decisions that depend on validated topology and dependencies. A well-maintained infrastructure digital twin is purpose-built to answer exactly these questions.[16]

	Data Lake / Warehouse	Infrastructure Digital Twin
Data Model	⚠ Tabular / time-series data	✓ Relationship-based infrastructure model
Validation	⚠ Provenance tracking only	✓ Validated real-world state
Freshness	✗ Works on refresh cycles	✓ Current operational view
Dependencies	✗ Limited dependency tracing	✓ Deep dependency traversal
Ontology	✗ No built-in infrastructure ontology	✓ Built-in infrastructure ontology

4. The Seven-Layer Digital Twin: The AI Context Layer for Infrastructure

FNT's infrastructure management approach is built on a seven-layer model that spans from physical facility through to business services:

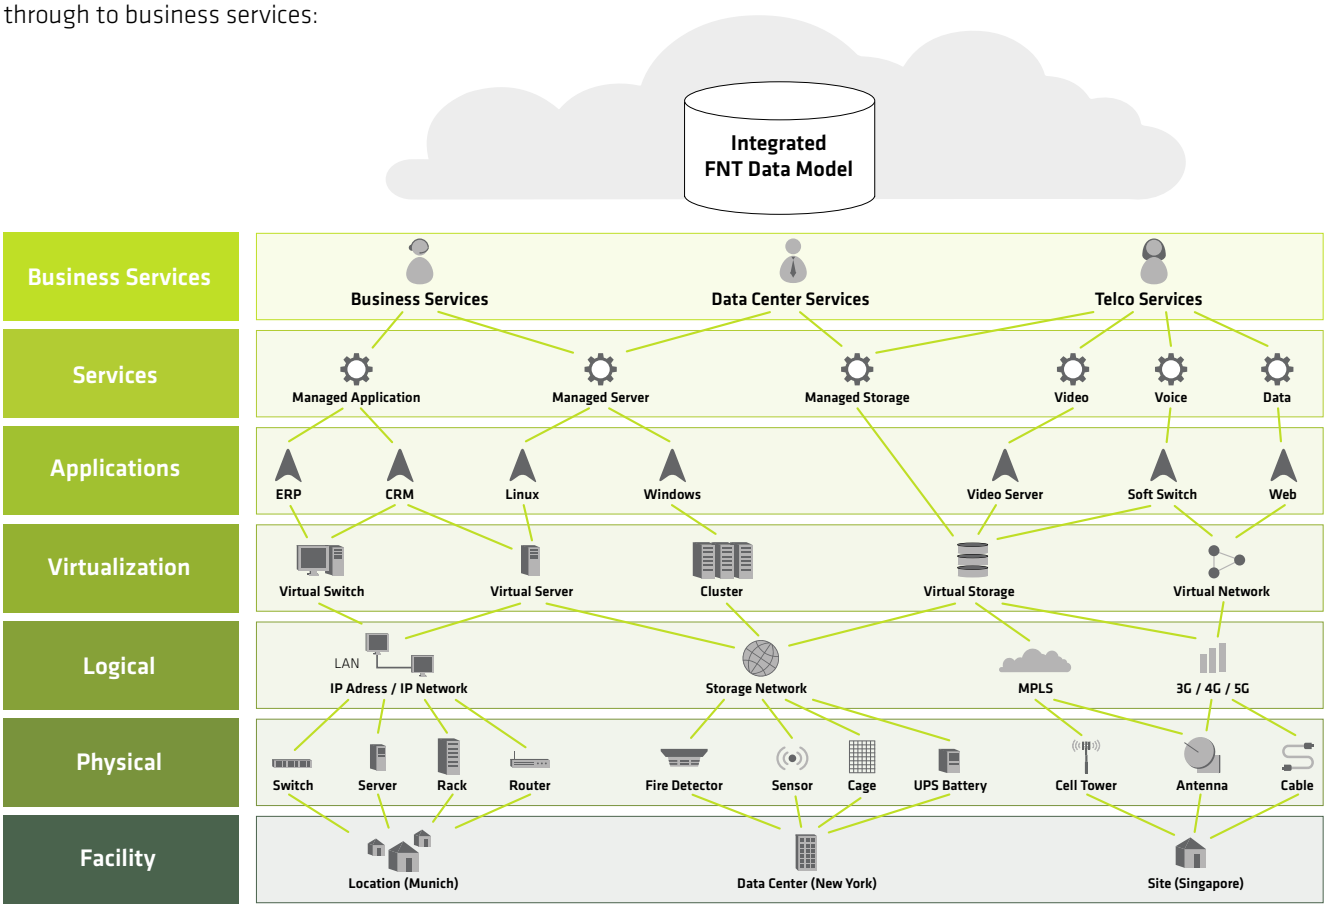

The importance of this model is the maintained relationship structure across the layers. Every logical circuit traces back to physical cable segments. Every application maps to its hosting infrastructure. Every business service carries a complete physical and logical dependency chain. This is the architecture of a true digital twin: not a collection of asset inventories in separate tools, but a unified, traversable model of the infrastructure in its entirety, with relationship-mapped data that can be consumed by AI tools and integration protocols, from graph databases and vector stores to context interfaces such as MCP, without requiring a costly rebuild of the underlying domain model.

For an AI agent operating in the infrastructure domain, this model is the context layer – the structured, relationship-mapped, validated representation of the physical world that enables it to reason correctly.

Critically, that context layer must span time, not just the present moment. Like a navigator who needs charts, a current position fix, and a planned route – not just a GPS reading – an infrastructure AI agent needs three temporal views simultaneously: the historical record of how the infrastructure was configured and changed (past), the current validated state of every asset and connection (present), and the planned and committed changes not yet implemented (future). A monitoring system provides the present. Only a maintained infrastructure digital twin provides all three. Without change history, an agent cannot distinguish an anomaly from a deliberate recent

modification. Without planned state, a capacity agent will commit resources already allocated to a pending project. The temporal completeness of the digital twin is what separates genuine infrastructure intelligence from fast guesswork.[15] This is precisely what FNT's closed-loop model delivers in practice: the Plan stage documents intended changes before execution (future state), the Verify and Update stages ensure the digital twin reflects completed changes (past state), and the current validated record is always queryable between them (present state). No monitoring or network management system maintains all three. No data lake preserves the distinction between planned and implemented state.

Consider what this concretely enables:

Capacity planning AI. An agent can query available space, power, and cooling headroom with knowledge of not just allocated capacity but committed, pending, and available capacity – including the dependency chains that constrain each. "Rack 14G has 4U of free space, 2.4 kW of available power circuit capacity, and sits on a cooling zone that is currently at 78% of design load" is a response a digital twin can deliver precisely. A data lake cannot.

Incident root cause AI. An agent investigating a service degradation can traverse from the impacted service layer down through logical and physical dependencies to identify the affected component and its physical location, the technicians certified on that equipment, and the most recent change that touched that path – all without human involvement in the traversal.

Change impact AI. Before executing a planned network change, an agent can ask the digital twin: "What other services share this physical path? What is the maintenance window constraint for each? Who are the service owners?" This impact analysis – which today can take hours of manual cross-referencing – is a direct query against a maintained digital twin.

Compliance and audit AI. An agent generating a regulatory compliance report can pull a complete picture of infrastructure state, change history, protection dependencies, and service resilience posture – because the digital twin maintains all this data in a structured, auditable form.

None of these use cases are theoretical. Each is a near-term application of agentic AI to infrastructure operations. Each requires the same thing: a validated, relationship-mapped, continuously maintained digital twin as its data foundation.

5. Why a General-Purpose CMDB Is Not Sufficient

A common first response to the infrastructure knowledge problem is to point to the organization's existing Configuration Management Database. If the CMDB already tracks IT assets and their relationships, why is a purpose-built infrastructure digital twin needed?

The distinction matters – and it is structural, not superficial.

A general-purpose CMDB, in the ITSM tradition, is designed to track configuration items (CIs) and the logical relationships between them: servers, services, applications, and their dependencies as understood by IT operations. This is a valuable capability. But it is designed from the service management layer downward, and it stops well short of the physical infrastructure reality that AI agents need to reason about.

The gap reveals itself in four specific ways:

1. **Physical precision.** A CMDB records that a server is "in Data Center A." A physical infrastructure digital twin records that it is in Row G, Rack 14, U-positions 3–5, connected to patch panel PP-G14-A on ports 1 and 2, drawing power from circuit C-G14-PDU-B, and sitting in a cooling zone currently at 78% of design load. The first record is useful for service management. The second is what an AI agent needs to plan a physical change.

2. **Connectivity at port level.** CMDBs model logical connections – "Server A connects to Application B." They do not, in general, track which physical port on which switch connects to which physical port on which server, through which patch cable, to which panel. That level of physical connectivity modeling – essential for cable management, capacity planning, and root-cause diagnosis – is the native domain of infrastructure management platforms, not CMDBs.

3. **Change-validated accuracy.** CMDB data is notoriously difficult to keep current. Without automated reconciliation against physical discovery and closed-loop change management that updates the CMDB on completion of every physical change, records drift rapidly from reality. Practitioners routinely describe CMDBs that are "more aspirational than factual" – populated in the original implementation and then only partially maintained. For AI agents, a CMDB with 70% accurate records is not a useful data source; it is a source of confident errors at scale.

4. **Cross-layer traversal.** The dependency traversal that infrastructure AI agents need – from a degraded business service, down through logical circuits, through physical connections, to a specific cable segment in a specific physical location – requires a model that spans both the physical and logical layers in a single connected graph. General-purpose CMDBs are not designed for this traversal; purpose-built infrastructure management platforms are.

This is not an argument against CMDBs. In environments where FNT is deployed, integration with ITSM CMDBs is standard practice. FNT provides the physical and logical infrastructure layer, the CMDB provides the service management layer, and the integration creates a complete picture. The argument is that for AI agents operating in the physical infrastructure domain, the CMDB alone – even a well-maintained one – is not a sufficient data foundation. Beyond point-in-time queries, the depth and accuracy of infrastructure documentation directly determines the precision of AI-driven simulation and planning. A sparse model produces approximate suggestions, while a detailed, validated twin enables what-if analyses and predictions the organization can act on with confidence.

Best Practice: CMDB + Digital Twin

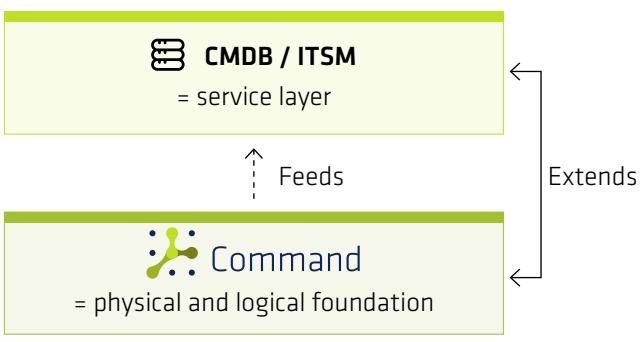

6. The Closed-Loop Imperative: Why Maintenance is the Real AI Enabler

Understanding the value of a digital twin as an AI foundation leads to a reframing of what investment in infrastructure management tooling really means. It is not documentation overhead. It is AI enablement work.

This reframing applies specifically to the closed-loop processes that keep the digital twin accurate: auto-discovery and reconciliation, change management integration, field verification workflows, and the discipline of ensuring that physical changes are reflected in the management system.

Closed-loop process – the AI enablement discipline

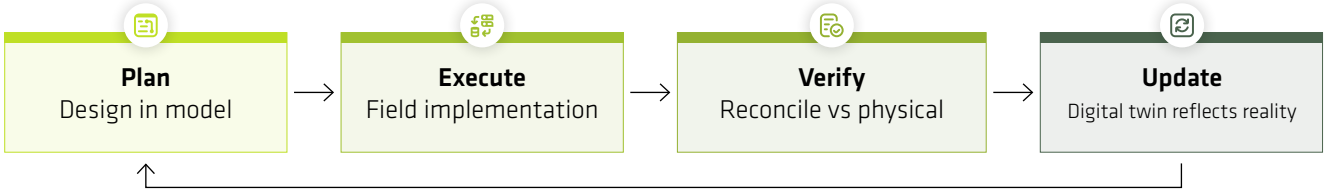

Closed loop – every physical change returns to plan

The closed-loop principle operates in four stages:

1. **Plan –** Changes are designed and documented in the digital twin before execution, including impact assessment against the existing model.
2. **Execute –** Field technicians carry out the change, using work orders generated from the digital twin data.
3. **Verify –** Auto-discovery tools and field confirmation processes validate that the physical change matches what was planned.
4. **Update –** The digital twin is updated to reflect the completed change, closing the loop between planned state, executed state, and documented state.

This discipline – Plan → Execute → Verify → Update – is the mechanism by which a digital twin stays aligned with physical reality. For human-driven operations, this improves efficiency and reduces errors. For AI-driven operations, it is a prerequisite for safe deployment.

There is an important implication here that runs counter to a common assumption: as AI automation increases, the quality demands on the underlying data increase, not decrease. Manual processes tolerate ambiguity because humans apply judgment. An agent executing a provisioning sequence at machine speed does not tolerate ambiguity. It acts on what is there. Every degree of automation raises the cost of inaccuracy in the source data.

This means organizations that have historically deprioritized infrastructure documentation rigor – accepting the as-built versus as-maintained gap as a cost of doing business – are not in a position to safely deploy agentic infrastructure automation. Closing the gap is not a prerequisite for today's AI investments; it is a prerequisite for the infrastructure AI capabilities that will arrive over the next 18 to 36 months.

That does not mean every organization needs a perfect, enterprise-wide digital twin before starting any AI work. It does mean that selective, domain-by-domain trustworthiness is a more realistic readiness target than broad claims of AI maturity built on weak infrastructure records.

7. Infrastructure Management as an AI Integration Layer

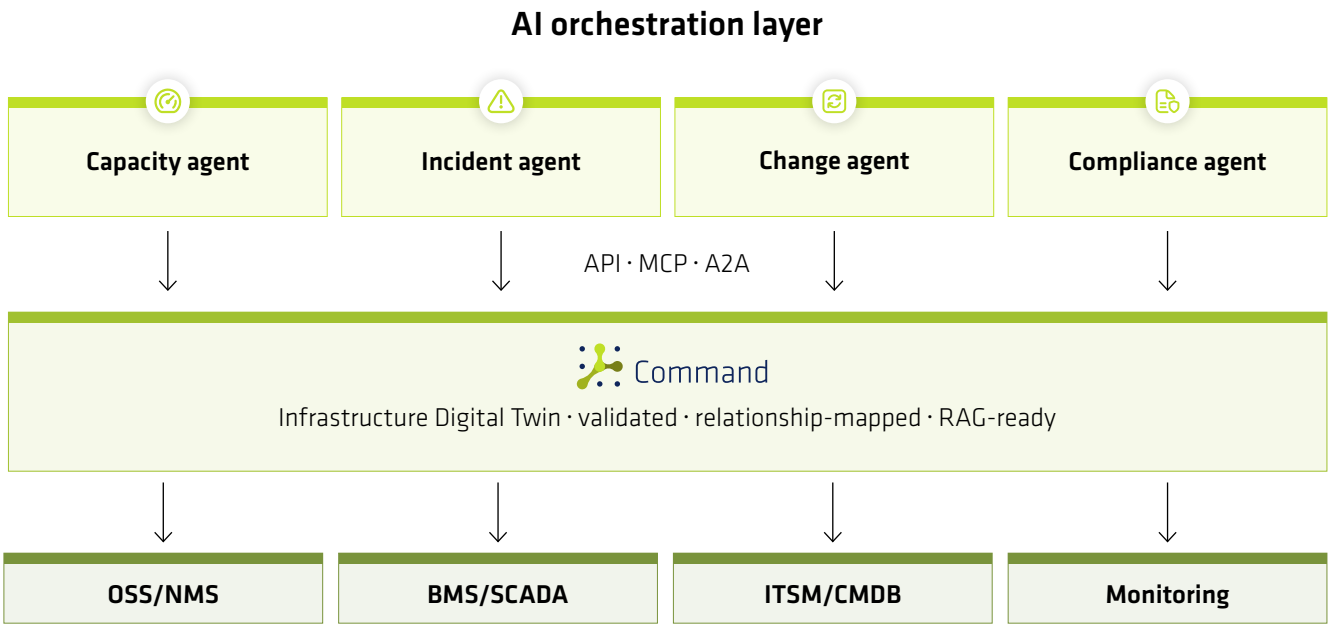

Infrastructure & systems data sources

A digital twin that cannot be queried by AI agents efficiently is not an AI foundation – it is a well-maintained island. The architectural requirement for infrastructure management in the agentic AI era is not just data quality; it is programmatic accessibility.

This means four things in particular:

1. **API-first architecture.** AI agents query infrastructure data through APIs. An infrastructure management platform that requires human interaction for every data retrieval is not agent-accessible. The platform must expose its data model – assets, relationships, dependencies, capacity, change history – through well-documented, performant APIs that agents can call in real time.

2. **Event-driven integration.** Agents operating in infrastructure workflows need to respond to infrastructure events – device state changes, capacity threshold crossings, compliance alerts – not just query on demand. This requires an event-driven integration architecture that publishes infrastructure state changes to downstream systems as they occur.

3. **Support for emerging interoperability standards.** Two developments in the agent ecosystem matter most today: MCP (Model Context Protocol) for tool and context access, and A2A (Agent-to-Agent) protocol for agent-to-agent communication. MCP was introduced by Anthropic as an open standard for connecting AI applications to data sources and tools.[6] A2A was introduced by Google and later advanced under Linux Foundation stewardship as a protocol for interoperable agent communication.[7] ACP was part of the early protocol landscape, but its development has since converged into A2A, so it is better understood as part of that consolidation rather than as a separate enduring standard.[8] Infrastructure management platforms that participate in these integration patterns become first-class components in AI agent workflows – accessible to any orchestration layer that speaks the protocol, without custom integration work for each use case. In practical terms, this is the

infrastructure layer equivalent of what REST APIs did for enterprise software integration a decade ago.

4. **RAG-ready data architecture.** Retrieval-Augmented Generation (RAG) has become the standard pattern for grounding AI model outputs in domain-specific, factual data. Rather than relying on general model training, a RAG-based agent retrieves relevant context from a structured knowledge source at query time – which it then uses to generate accurate, grounded responses and decisions. FNT Command is purpose-built to serve as the RAG retrieval layer for infrastructure AI. Its validated, relationship-mapped data model provides exactly the structured, queryable, domain-specific context that RAG architectures require. When an infrastructure AI agent asks, "what is the current state of the power chain serving this rack?" or "which services share this physical cable route?", the answer is retrieved directly from the FNT digital twin – not generated from training data, not approximated from a data lake – producing outputs that are grounded in the verified state of the actual infrastructure.

FNT's infrastructure management solution is built on an API-first architecture with robust integration capabilities, and the platform roadmap moves toward native support for emerging agentic AI protocols. The objective is to make the FNT infrastructure digital twin – spanning physical assets, logical services, and virtual resources across data centers, enterprise IT, and telecommunications networks – directly queryable by AI orchestration layers; not as an after-thought, but as a designed, first-class integration point.

The significance of this is not technical. It is strategic. An infrastructure management platform that is natively accessible to AI agents becomes the infrastructure context layer for the entire AI automation stack. Capacity planning agents, incident management agents, change orchestration agents, compliance reporting agents – all of them, querying one validated source of infrastructure truth. The right retrieval architecture follows directly: authoritative operational systems first – because they carry validated topology, relationships, and intended state that no downstream copy can fully replicate – analytical platforms second for historical context and trend data, and raw data sources last, only where no authoritative record exists. In infrastructure AI, source selection is not a secondary concern. It is the primary architectural decision.

8. Three Scenarios: What This Looks Like in Practice

Scenario 1:
Automated Capacity Expansion in a Multi-Site Colocation Environment

A hyperscale enterprise customer submits a request to expand their compute footprint at a colocation facility by 20 additional racks, with a target in-service date six weeks out.

In a traditional workflow, this triggers a manual capacity study: available white space, power circuit capacity, cooling zone headroom, cross-connect availability, connectivity to the customer's existing footprint. Depending on staffing and data quality, this study takes two to five days. The project then moves into planning, provisioning, and construction coordination.

In an AI-assisted workflow, a capacity planning agent receives the request and queries the infrastructure digital twin directly. Within minutes, it has identified three candidate locations that meet the space, power, and cooling requirements, ranked by proximity to the customer's existing footprint, cross-connect costs, and cooling zone efficiency. It has automatically generated a capacity impact report for each option, identified the infrastructure changes required, and drafted a delivery timeline. A human operations lead reviews and approves the recommendation. The detailed work order is generated automatically.

This workflow is not speculative – the AI reasoning components exist today. What most colocation operators are missing is the data foundation that makes the query answerable with confidence. A digital twin with accurate, real-time capacity data across all dimensions – space, power, cooling, connectivity – is what transforms a capable AI model into an operational AI agent.

Automated Capacity Expansion: Before vs After

Scenario: Customer requests 20 additional racks with a six-week target in-service date

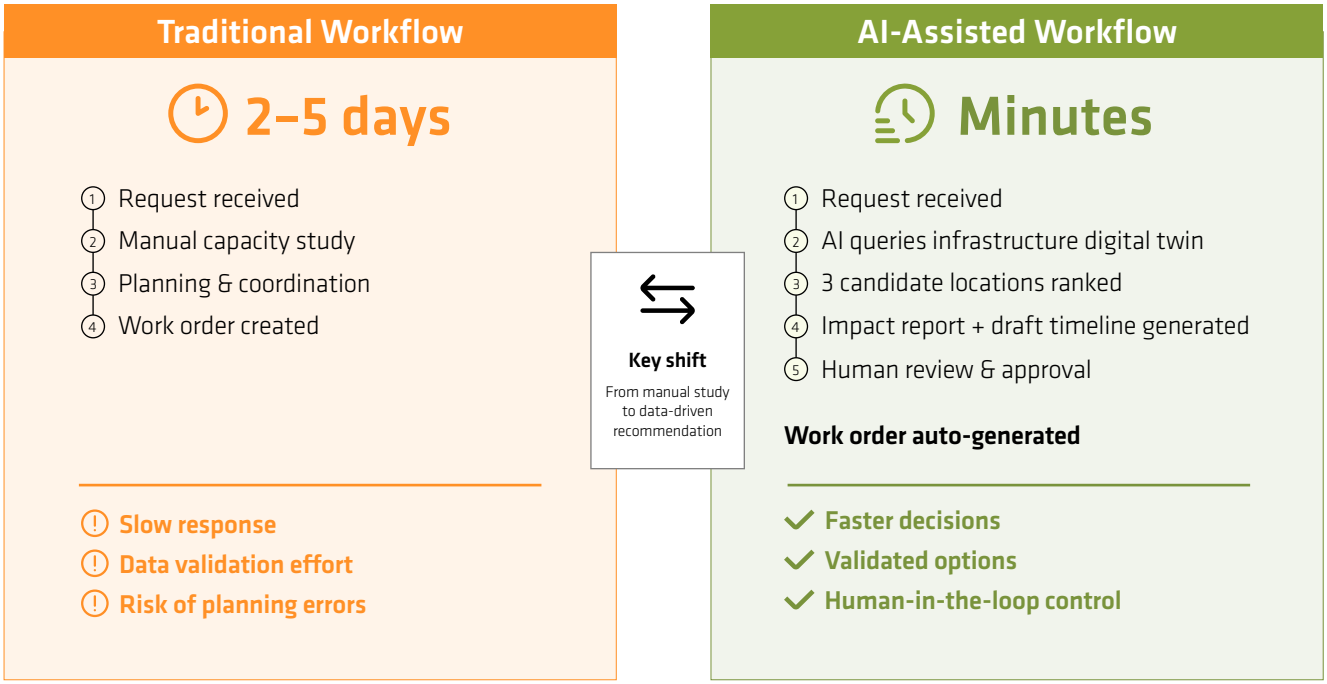

Scenario 2:

Automated Root-Cause Identification in a Telecom Network

A network performance monitoring tool detects degraded performance on a regional circuit serving 47 enterprise customers. In a traditional workflow, a network engineer begins manually tracing the affected path through multiple systems – documentation, topology, monitoring – to identify the physical segment responsible. On a complex network with siloed documentation, this can take hours.

With an AI diagnostic agent operating against a maintained network inventory that spans logical circuits, physical cables, and outside plant topology in a single model, the trace is automated. The agent traverses the dependency chain from affected services to the physical layer, identifies the cable segment showing anomalous behavior in monitoring data, correlates it with a recent maintenance activity in the field records, and produces a root-cause hypothesis with the physical location, the affected customers, and the relevant field contact – in under five minutes. Beyond reactive fault isolation, the same unified documentation layer enables proactive simulation – the ultimate Telco digital twin use case. An agent can verify circuit redundancy without triggering an actual network outage, identify single points of failure across the topology before they become incidents, and correlate historical performance data with infrastructure configuration to predict where degradation is most likely to occur next. Real-world impact is already measurable: in deployments where infrastructure documentation spans both logical and physical layers with full dependency mapping, impact assessment cycles that previously took three days have been completed in under three hours.

The unified network documentation is not incidental to this outcome. It is the mechanism that makes the dependency traversal possible. Without the maintained relationship model from logical service to physical infrastructure, the agent has no path to traverse.

Topology Traversal for Faster Root-Cause Diagnosis

From affected service to physical fault segment

Traditional Workflow	**Network Engineer** • Reactive fault awareness	**Network Engineer** • Multiple documentation tools • Manual topology traversals and fault location	**Network Engineer** • Manual solution design and execution based on engineering guides and SOPs
AI-Assisted	**Fault Agent** • Proactive awareness and prediction • Correlation	**Fault Agent** • Intelligent impact and fault analysis based on dependencies • Automatic topology traversals	**Fault Copilot + Network Engineer** • Automatic solution design • Remediation plan suggestion by Copilot • Human in-the-loop to accept the plan

Scenario 3:

Continuous Compliance Monitoring in an Industrial OT Environment

A manufacturing company operating under NIS2 and ISO 27001 obligations needs to demonstrate continuous awareness of its OT network topology and the protection status of its critical control systems. The traditional approach is periodic audit, which is expensive, point-in-time, and increasingly insufficient under regulatory frameworks that demand demonstrable, continuous resilience.

With an AI compliance agent operating against a maintained OT infrastructure digital twin – devices, network segments, protection dependencies, change history – continuous monitoring becomes automated. The agent detects when an unauthorized device is added to a protected network segment, when a configuration change alters the protection classification of a critical asset, or when a maintenance activity creates a temporary gap in a protection chain. It generates the compliance alert, documents the event, and updates the risk posture report – automatically and in real time.

The prerequisite is not the AI agent. It is the infrastructure digital twin that reflects the OT network accurately enough that the agent can detect meaningful deviations. Organizations that have invested in OT asset documentation and network inventory are positioned to deploy this capability in the near term. Those that have not are not.

9. Infrastructure AI Readiness: An Assessment Guide

Understanding where an organization stands relative to this challenge requires an honest assessment of its infrastructure documentation posture. The following framework describes four levels of infrastructure AI readiness, based on the maturity of the digital twin and the processes that maintain it. This is an FNT-defined readiness assessment guide to help organizations in understanding how prepared their infrastructure and documentation are. It is different from TM Forum's Autonomous Networks maturity model – which uses a different level structure specifically for network autonomy – though both frameworks rest on the same foundational principle: that higher levels of operational autonomy require progressively more complete, accurate, and governed infrastructure data.

Stage	State	Documentation Posture	Likely AI Posture
Fragmented	No reliable system of record across infrastructure domains	Records are incomplete, local, or stale; as-built vs. as-maintained gap is large and unmanaged	AI use should remain exploratory and advisory only; autonomous workflows will produce unreliable outputs
Consolidated	A central platform exists, but coverage and process discipline vary	Some domains are usable; others drift quickly; no consistent closed-loop process	Human-in-the-loop assistance is feasible, but agents require human validation of every output
Maintained	A digital twin exists for priority domains with closed-loop updates and active reconciliation	Accuracy is tracked and continuously improved in operationally important areas	Selective agentic automation becomes feasible with clear governance boundaries
Optimized	The model is broadly maintained, event-aware, and programmatically accessible across domains	Variance from physical reality is detected and corrected quickly; accuracy sufficient for autonomous decision-making	More autonomous workflows viable in carefully governed use cases; agents can plan, execute, verify, and report

The key transition is usually from Consolidated to Maintained. That is the point at which infrastructure data becomes sufficiently reliable for higher-trust operational use. The shift requires primarily process discipline and tooling investment, not a technology breakthrough. The tooling exists. The processes are wellunderstood. The gap is organizational commitment.

The shift from Maintained to Optimized requires, additionally, the API and protocol integrations described in Section 7 – making the digital twin natively queryable by AI orchestration layers in real time.

In real life, this transition is not only a technical shift, it is a trust transition. Moving from AI that advises to AI that acts – from "humans in the loop" to "humans out of the loop" – is the step that most operators find hardest to take.

Moving from disparate domains towards one consolidated knowledge source is the definitive event which unlocks transitioning from AI that advises to AI that acts, from humans in the loop to humans out of the loop, from TM Forum level 3 to 4. Many large enterprise organizations, colocation operators, and telecom networks today operate somewhere between Levels 1 and 2 in at least part of their

estate. The determining factor for progression along the maturity scale is whether priority domains are trustworthy enough for the AI-enabled workflows the organization wants to deploy.

A practical starting point for self-assessment: What percentage of infrastructure records in your most critical facility, network segment, or OT environment has been verified against physical reality in the last 12 months? The honest answer to that question is often more revealing than any dashboard summary – and it tells you which stage you are at, regardless of what your CMDB reports.

Infrastructure AI readiness – maturity levels

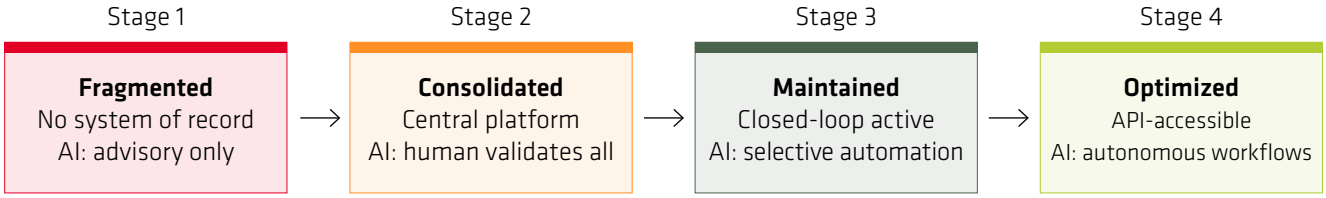

Key transition:
Stage 2 to 3 requires process discipline, not a technology breakthrough
Stage 3 to 4 additionally requires API/protocol integration

10. The Regulatory Imperative: Compliance as a Second Driver for Digital Twin Investment

The operational AI readiness argument above is the primary case for investing in an accurate infrastructure digital twin. But an increasingly powerful second argument is emerging from the regulatory environment. The relevant point is not that regulations require a digital twin – they do not. The point is that they increase the importance of accurate asset records, dependency visibility, change traceability, resilient documentation, and the ability to explain or evidence operational decisions. For many organizations, a maintained digital twin is one of the most scalable ways to satisfy those obligations.

European Union: NIS2, DORA, and the EU AI Act

NIS2 (Network and Information Security Directive 2) required Member States to transpose the directive and expanded cybersecurity obligations across 18 critical sectors including energy, transport, telecommunications, digital infrastructure, healthcare, water, and public administration.[9] For affected organizations, this raises the value of maintained records of systems, dependencies, and security-relevant operational context. Demonstrating compliance at audit requires precisely the kind of structured, current infrastructure documentation that a digital twin provides. Without it, compliance becomes a point-in-time exercise rather than a continuous operational posture.

DORA (Digital Operational Resilience Act) requires financial entities and their critical ICT service providers to strengthen ICT risk management, operational resilience, incident handling, and related governance.[10] DORA requires comprehensive ICT risk management frameworks that include detailed documentation of ICT assets, their dependencies, and their resilience posture – and financial entities must be able to demonstrate this to supervisors on demand. The asset and dependency mapping that FNT's infrastructure digital twin provides is precisely the data layer DORA auditors are looking for.

The EU AI Act creates increasing scrutiny around governance, traceability, and explainability for AI systems classified as high-risk, which includes AI used in critical infrastructure management. A maintained, auditable infrastructure context layer directly supports those expectations: an AI agent that makes an autonomous infrastructure decision must be able to reference what data it acted on and why, which is only possible when that data is validated and traceable.

EU regulatory timeline – infrastructure documentation obligations

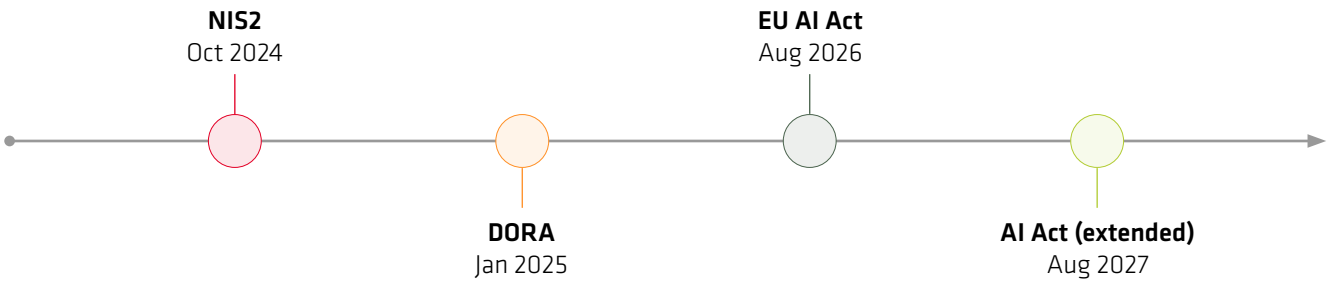

North America: NERC CIP and the Federal Continuous-Monitoring Model

NERC CIP (North American Electric Reliability Corporation Critical Infrastructure Protection) is the mandatory cybersecurity framework for entities that own, operate, or use assets connected to North America's Bulk Electric System. Standard CIP-002 requires responsible entities to identify and categorize BES Cyber Systems and to review and update those identifications at least once every 15 calendar months.[12] This is a concrete example of the continuing compliance importance of maintained asset categorization and evidence. An infrastructure digital twin that tracks OT asset inventory, network topology, and configuration baselines provides the continuous monitoring that makes audit evidence available on demand rather than requiring a manual collection exercise.

The broader US federal security model under FISMA and NIST emphasizes hardware and software asset inventories and continuous monitoring as core elements of information security programs. NIST SP 800–53 Rev. 5 includes controls for inventory of system components (CM-8), and NIST SP 800–137 establishes the Information Security Continuous Monitoring model for federal agencies.[13, 14] This is not equivalent to a mandate for a specific platform, but it reinforces the operational and compliance value of current, queryable infrastructure records across federal and contractor environments.

APAC: Cybersecurity Act

In the Asia-Pacific region, Singapore's **Cybersecurity Act** establishes obligations for Critical Information Infrastructure (CII) owners across eleven critical sectors including

energy, transport, and telecommunications. CII owners must identify designated systems, report incidents, and comply with codes of practice that include asset identification and risk management requirements. As Asia-Pacific frameworks continue to align with international standards such as NIST CSF and ISO 27001, the underlying obligation is consistent with those of NIS2 and NERC CIP: organizations must demonstrate current, accurate knowledge of their infrastructure and its interdependencies. A maintained infrastructure digital twin serves this compliance need as directly in Singapore or Australia as it does in Frankfurt or Chicago.

The Unified Argument

Taken together, these frameworks support a practical conclusion: the same infrastructure data discipline that improves AI readiness also improves compliance readiness. A maintained digital twin is therefore not a niche documentation investment. It is a multi-purpose operational capability that simultaneously supports:

- higher-confidence AI-assisted and increasingly autonomous infrastructure workflows,
- better evidence for asset, dependency, and resilience questions under NIS2, DORA, and NERC CIP,
- faster response to audits and regulatory reviews, and
- clearer traceability when AI-driven decisions depend on infrastructure state under the EU AI Act.

The infrastructure digital twin is not three separate investments for three separate purposes. It is one investment that serves all three simultaneously. This changes the ROI calculation materially – and it changes the urgency of the investment for any organization subject to these frameworks.

Regulatory Readiness: Where Infrastructure Documentation Matters

A digital twin is not mandated – but accurate infrastructure records support compliance readiness

Framework	EU	North America	APAC	Infrastructure documentation relevance
NIS2	✓			Current asset records, dependency visibility, and evidence for audit readiness.
DORA	✓			ICT asset and dependency mapping, resilience documentation, and supervisory evidence.
EU AI Act	✓			Traceability, explainability, and auditable data context for AI-driven decisions.
NERC CIP		✓		Maintained asset categorization, topology visibility, and on-demand audit evidence.
FISMA / NIST		✓		Inventory, continuous monitoring, and queryable infrastructure records.
Cybersecurity Act			✓	Compliance and risk management requires accurate knowledge of infrastructure.

11. The Repositioning: From Documentation Tool to AI Foundation

Traditionally, DCIM, network/resource inventory, ITAM, and cable management have been justified through operational efficiency, planning accuracy, and compliance support. Those arguments still matter. But in the AI era, they are no longer the whole story.

The stronger framing is this: a validated, continuously maintained infrastructure digital twin is becoming a foundational data asset for infrastructure automation. It is not a nice-to-have for AI maturity. It is a prerequisite. AI is only as reliable as the infrastructure knowledge behind it. FNT does not claim to be the AI. It provides something more fundamental: the authoritative knowledge that makes AI reliable. Not the AI. The reason it works.

That shift changes three things:

1. **The ROI calculation expands.** The value is not limited to documentation efficiency. A DCIM or NRI platform that enables automated capacity planning, intelligent incident response, and continuous compliance monitoring has a return profile that far exceeds its cost as a documentation tool. The investment should be evaluated against the value of the AI capabilities it enables. For organizations subject to NIS2, DORA, or NERC CIP, the compliance value adds a further dimension to this calculation.

2. **The timing becomes more strategic.** Organizations that deprioritized infrastructure documentation rigor in favor of other IT investments have a closing window to build the data foundation before their AI roadmaps arrive at infrastructure use cases and find nothing reliable to build on. Eighteen months of closed-loop process discipline now is worth far more than the same effort applied reactively after a failed AI deployment.

3. **The ownership model rises.** Infrastructure context stops being a back-office record-keeping issue and becomes part of enterprise data, resilience, and governance strategy. Under DORA, it is also a regulatory compliance asset with board-level accountability. This means the conversation about infrastructure documentation investment no longer belongs only in IT operations – it belongs in the boardroom.

Conclusion: Authoritative Data is Not Optional

The infrastructure AI era is not hypothetical. It is here, in early deployment, and accelerating. The organizations building it are discovering, consistently, that model capability is not the constraint. Data quality, data completeness, and data accessibility are the constraints.

For the physical infrastructure domain, the only adequate response to this constraint is a maintained, validated digital twin that covers the full stack from facility and cable through logical services – one that is API-accessible, event-driven, and increasingly natively integrated with agentic AI protocols.

The argument that DCIM, NRI, and cable management are legacy overhead in an AI-driven future is precisely backwards. They are the precondition for an AI-driven future that is reliable, safe, and scalable. Every rack, cable, circuit, and device that is accurately modeled today is one less potential source of error when an autonomous agent acts on that data tomorrow.

Where to Start

The right starting point is not a grand transformation program. It is an honest assessment of the highest-value domains:

- Where are records trustworthy, where are they drifting, and which closed-loop processes would raise confidence fastest?
- What percentage of physical infrastructure records in your most critical facility or highest-risk network segment have been verified against physical reality within the last 12 months?
- Where are the largest gaps – in capacity-constrained facilities, multi-domain network segments, OT environments with active NIS2, or NERC CIP exposure?

A targeted reconciliation effort in those priority areas – combined with the closed-loop process discipline to maintain accuracy going forward – is the foundation on which everything else is built. It does not require a complete overhaul before AI projects begin. It requires knowing which domains are trustworthy enough to support automation today, and a credible plan to extend that coverage.

The organizations most likely to lead in infrastructure AI will not necessarily be those with the most ambitious AI messaging. They will be the ones that invested early and deliberately in making infrastructure knowledge trustworthy, governable, and usable.

You cannot automate what you cannot verify. The verified knowledge of your infrastructure lives in its digital twin. Building and maintaining that twin is the most consequential AI readiness investment an infrastructure organization can make right now.

Endnotes

1. Bain & Company, "Why AI Stumbles Without a Solid Data Strategy," 2025.
2. Bain & Company, "Building the Foundation for Agentic AI," Technology Report 2025.
3. GlobeNewswire / UMIP, "New Economic Study Identifies a $300 Billion Infrastructure Identity Gap in the Built Environment," March 2026. Used here as a directional estimate rather than a definitive market benchmark.
4. MIT Technology Review Insights, "Building a strong data infrastructure for AI agent success," March 2026.
5. ISACA, "Resilience and Security in Critical Sectors: Navigating NIS2 and DORA Requirements," White Paper 2025.
6. Anthropic, "Introducing the Model Context Protocol," November 25, 2024.
7. Google for Developers, "Announcing the Agent2Agent Protocol (A2A)," April 9, 2025; Linux Foundation, "Launches the Agent2Agent Protocol Project," June 23, 2025.
8. IBM Research / IBM Think, ACP materials noting ACP's merger into A2A under the Linux Foundation umbrella, 2025.
9. European Commission, "NIS2 Directive: securing network and information systems," confirming the 17 October 2024 transposition deadline.
10. EUR-Lex, Regulation (EU) 2022/2554 (DORA), confirming application from 17 January 2025.
11. European Commission, "AI Act | Shaping Europe's digital future," including staged applicability and the 2 August 2027 extension for some high-risk AI systems embedded in regulated products.
12. NERC, "CIP-002-5.1a – Cyber Security – BES Cyber System Categorization," Requirement R2, review and approval at least once every 15 calendar months.
13. NIST SP 800-53 Rev. 5, "Security and Privacy Controls for Information Systems and Organizations," including CM-8 (System Component Inventory).
14. NIST SP 800-137, "Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations."
15. Appledore Research / FNT Software, "The Secret Ingredient for Network Digital Twins: Why Unified Inventory Is the Foundation for Network Autonomy and Business Agility," FNT White Paper 2026. Available at fntsoftware.com
16. Dell Technologies / Forrester, "The Challenge: Disconnected Data, Disconnected Decisions," 2024.

About the Authors

Daria Batrakova is Director Telecom Solutions at FNT Software. She has worked in network operation, OSS integration and solution advisory roles in the telecommunications field for over 20 years

Oliver Lindner is Director of Product Management at FNT Software, where he leads the DCIM and infrastructure management product line. With more than fifteen years of experience in data center infrastructure management, connectivity, and digital twin architecture, he works with global enterprises, colocation operators, and network service providers at the intersection of infrastructure visibility, operational efficiency, and emerging automation capabilities.

About FNT Software

FNT Software, headquartered in Ellwangen (Jagst), Germany, simplifies the management of highly complex digital infrastructures in companies and public authorities with its FNT Command Platform. With the cloud-enabled "Software made in Germany", IT, telecommunications and data center infrastructures can be efficiently recorded as digital twins and documented across all levels from buildings to digital services. The software also offers open interfaces and numerous functions for planning, implementing and automating transformations and changes in an integrated manner. FNT's customers include more than 500 companies and government agencies worldwide, including more than half of the DAX-40 listed corporations. FNT* operates offices in several locations in Germany as well as in New York, Singapore and Timisoara and has an international partner system with market-leading IT service providers and system integrators.

*Refers to FNT Software GmbH, including its subsidiaries, and FNT Services GmbH.

Transparency note:
Written by Experts | Authors: Daria Batrakova, Oliver Lindner | AI-Enhanced