


EBOOK

A DCIM Playbook to Close the AI Readiness Gap

Across Enterprise, Neocloud,
and Sovereign Data Centers



Sunbird®

DCIM that's easy, fast, and complete

Table of Contents

- Introduction.....3
- What AI Readiness Means Across Enterprise, Neocloud, and Sovereign Data Centers.....4
- Maintaining Accurate Inventory Across the Asset Lifecycle.....5
- Planning High-Density AI Deployments Before Hardware Arrives.....8
- Supporting Liquid Cooling with Accurate Modeling.....12
- Managing Power Capacity From Utility Feed to GPU.....14
- Integrating Tools to Create a Single Pane of Glass.....16
- Conclusion.....19
- Take the Next Step with Sunbird.....20

Introduction

AI workloads are reshaping data center infrastructure at a pace not seen before. This introduces new operational complexities, particularly for organizations that were not purpose-built for AI.

Yet, with the dynamic environments of today's data centers, documentation can quickly fall out of sync with the real world. When data becomes inaccurate, everything from planning to troubleshooting to reporting can be impacted.

GPU clusters consume 10 or more times the power of general-purpose servers. Liquid cooling is moving from niche to necessary. Procurement cycles are getting shorter, deployment timelines are getting faster, and the cost of getting infrastructure planning wrong is measured in downtime, stranded capacity, and lost revenue.

Today, if you're not prepared for AI, you're already behind.

Closing the AI-readiness gap requires accurate infrastructure data, complete visibility across the physical and virtual environment, and the ability to plan, deploy, and monitor high-density compute at scale. Data Center Infrastructure Management (DCIM) software makes that possible.

This eBook covers the core areas where DCIM software makes an impact: asset lifecycle management, high-density capacity planning and power management, liquid cooling infrastructure modeling, and system integration.



Image Credit: NVIDIA

What AI readiness means across enterprise, neocloud, and sovereign data centers

AI readiness means having the visibility, capacity, and processes in place to plan, deploy, and manage AI infrastructure without introducing unnecessary risk. For most organizations, the readiness gap tends to be rooted in process and data quality such as having incomplete asset records, siloed data, capacity assumptions that are not validated in real time, and manual, time-consuming processes.

The specific gaps, however, vary by environment:

- **Enterprise data centers** are typically retrofitting existing facilities to support AI workloads alongside traditional IT. Common issues are that power and cooling infrastructure were not designed for increased rack densities, space is limited, and tools and data are siloed. Before a single GPU server is deployed, teams need to know whether their existing capacity at the rack, row, and room level can handle the load, and what retrofits, such as liquid cooling infrastructure, are required.
- **Neocloud AI factories** are purpose-built for AI and GPU-intensive workloads, but scale, speed, and multi-tenant complexity create their own management challenges. With rapid customer onboarding, frequent deployments, and large hardware volumes moving through the facility, accurate asset tracking, parts/spares management, and lifecycle management are essential to avoid misplaced equipment, inventory gaps, and deployment errors.
- **Sovereign data centers** must maintain traceability of their physical hardware. Frameworks like NIS2 in Europe, SecNumCloud in France, and the BSI C5 in Germany mandate that infrastructure managers maintain a verifiable inventory to mitigate supply chain risks. Similarly, U.S. federal mandates (such as OMB M-25-03) now emphasize the use of tools to ensure the integrity and availability of government-operated infrastructure. In these environments, a single source of truth is needed to produce reports about the hardware that is powering sovereign workloads.

In all cases, the starting point is the same: accurate, complete, real-time visibility into infrastructure. DCIM software has been providing data center teams with that foundation well before AI entered the picture and its importance has only grown since.



Image Credit: NVIDIA

Maintaining accurate inventory across the asset lifecycle

AI hardware is expensive. For example, [an NVIDIA H100 server can cost \\$200,000 or more](#). A single untracked, misplaced, or incorrectly staged asset represents a significant operational and financial risk. At the scale many AI deployments demand, with hundreds or thousands of GPU servers arriving in compressed timeframes, manual tracking processes can break down quickly.

Many organizations still rely on spreadsheets, siloed systems, and manual processes to track equipment receipt and staging. When assets are not inventoried at the dock, problems can accumulate like unused assets going untracked and purchase orders not reconciled with what's onsite.

The scale and value of AI hardware make these problems even more consequential than they were in traditional data center environments. An undocumented GPU system can impact capacity planning and stall deployments.



Maintaining accurate inventory across the asset lifecycle

Asset lifecycle management from PO to dock to rack

DCIM software supports the full equipment lifecycle, starting at the initial planning phase of modeling complete GPU racks (or complete BOM) on the data center floor. When plans are approved, orders are placed and a shipment arrives on the loading dock, teams can use barcode scanning to log the serial numbers. This automatically creates an item record and populates the serial number and name fields. The purchase order number and other known details (e.g., make, model, storage location) can be added at the same time, giving finance and procurement teams immediate confirmation of receipt. That asset is in the system before it leaves the dock.

When an installation request is created, the same scan reopens the existing record. Teams can update the location, cabinet, and U position, set the status to installed, and apply a physical barcode or QR code asset tag to the device. Future audits can then be completed by scanning those tags directly, with built-in logic and voice-response guiding technicians through the process and an exception report identifying any discrepancies.

Parts and spares management

Managing AI infrastructure means managing more than the primary compute assets. GPU servers require a range of supporting components (e.g., memory modules, storage drives, power supplies, network cards, optics/transceivers, and patch cables) and running short of a critical spare can halt a deployment or extend a maintenance window.

DCIM software supports centralized management of parts and spares inventory. Teams can track quantities, locations, and component details for spare parts across all sites. Email alerts can be configured when stock drops below a user-defined threshold, prompting replenishment before the shortage becomes a problem. For high-value GPU infrastructure where spares are costly and lead times can be long, this visibility is important.




Image Credit: NVIDIA

Maintaining accurate inventory across the asset lifecycle

Asset audits and inventory accuracy at scale


Regular audits are essential to maintaining accurate inventory, especially in AI environments where large volumes of high-value hardware are received and deployed in rapid succession. DCIM software that supports barcode and QR code scanning reduces the time and error rate associated with manual audits. A single technician with a handheld scanner can work through an audit cabinet by cabinet, scanning the barcode or QR code tag on each device. The system provides a voice response confirming each asset and flags any discrepancy between what's scanned and what's in the database in real time. Audits can be completed by one person in less time with a higher accuracy rate and a documented exception report at the end.

Customers using this workflow have reported significant efficiency gains. Delta Dental completed a full audit of 18 racks—a process that would have taken a day and a half manually—in three hours. [Read the Delta Dental case study.](#)

Automated audit trails

In high-stakes AI environments knowing where an asset is isn't enough. Organizations must also know where it has been. Manual spreadsheets fail to provide a temporal record of changes.

DCIM software can generate a comprehensive audit trail for every asset in the facility. Every time a GPU server is moved, a cable is swapped, or a component is serviced, the system logs who did what and when. Instead of scrambling to reconstruct hardware history from emails, infrastructure managers can provide a complete audit trail in seconds.



**Get the full eBook
by clicking the link
below.**

[Download My Free eBook](#)

