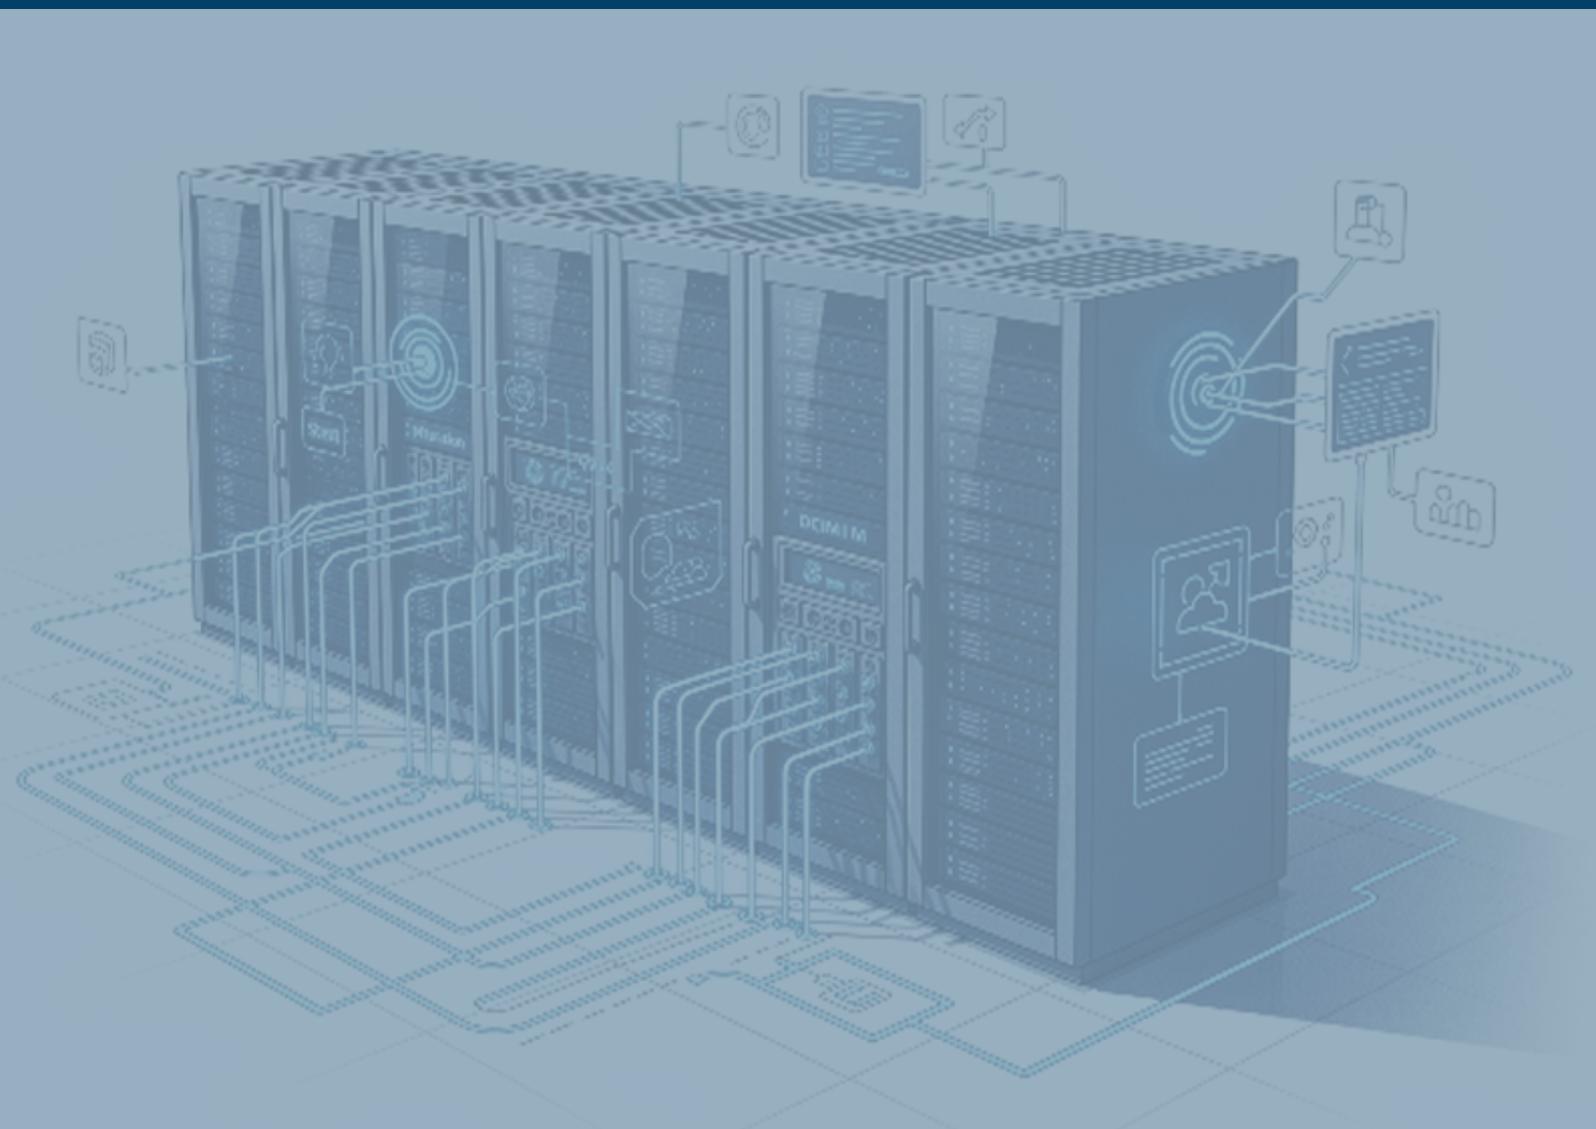


Risks in Implementing a DCIM and Strategies to Mitigate Them

Model, monitor and integrate with criteria



Introduction

Implementing a DCiM is not just a technology project; it is a change in the way infrastructure is understood and managed. As with any significant change, doubts, resistance, and risks naturally arise.

This is normal. What matters is not denying their existence, but identifying them early and addressing them with method and common sense.

With a practical and well-structured approach, those risks stop being a threat and become a natural part of the operational maturity process.

A DCiM delivers value when it achieves three things simultaneously:

1

Describes
your reality
(Modeling)

2

Measures
what happens
(Monitoring)

3

Connects
this data with the rest
of the organization
(Integrations)

It is not about “installing a tool.” It is about **transforming scattered knowledge into controlled operations.**

Project Layers

One way to control risks in a DCiM project is to address its different layers or phases. These do not necessarily need to be sequential, they can run in parallel, but they must always be tracked in parallel.

Let’s review the first ones.



Modeling

Building the “map” of the Data Center

The first questions to ask are, **Do I have the information required for modeling? Where is it located? Who owns it? Is it up to date?**

Required Information

- Structure: site / room / row / rack.
- Electrical chain (UPS, switchboards, PDUs, racks).
- Clear definition of relationships (what depends on what).
- IT inventory (and/or facilities infrastructure depending on scope).
- Naming and classification standards
- Basic connectivity

Does this mean you must have everything before starting a DCiM project? No, much of it will be organized and gathered during the project. But you must be aware beforehand of what you have and what is missing.

Common Risk: Incomplete or outdated modeling turns the DCiM into a “nice picture” that is useless for operations. Do not include data you cannot trust.

Deliverable (Quality Milestone): Coherent Model vs Reality

- Layout: Room sizes and distribution, facilities equipment footprints, racks, electrical distribution (if overhead), airflow distribution (perforated tiles / raised ducts / in-rows / contained aisles).
- Electrical traceability: How energy flows from entry point (utility/generator) to final link (rPDUs/terminals/connectors/AC equipment), considering redundancies.
- Network traceability: How connectivity flows from entry (carriers/offices) to final link (compute/storage).
- IT naming standard: DC/room/row/rack/U position.
- Minimum complete asset data: Location, type, brand, model, IP, labels.

Monitoring

Capturing live and reliable data directly from hardware through standard protocols:

Required Information:

- Reliable capture of electrical consumption, temperature, load, and equipment status.
- Thresholds aligned with operational reality.
- Actionable alarms (not constant noise).
- Correlation between physical and logical events.

Common Risk: Excess data without analysis leads to operational saturation and loss of trust in the tool.

Deliverable (Quality Milestone)

- Correct signals (value/unit/scale).
- Communication stability: Proper network and equipment configuration. Polling interval aligned with number of devices.
- Meaningful alarms: Proper notification policy configuration according to criticality. Avoid overloaded dashboards or inboxes with non-critical alarms (threshold warnings, scheduled disconnections, maintenance entries, etc).
- Minimum complete asset data: Location, type, brand, model, IP, labels.

Integrations

An isolated DCiM loses much of its potential. The real value appears when data flows into other systems. This is when DCiM starts becoming the brain of the Data Center.

BMS to DCiM Integration (unidirectional)

What BMS Provides | Facilities data (climate, energy/status in certain elements, environmental alarms, etc., depending on each BMS).

What DCiM Provides | Operational context:

- “Where is it happening?”
- “Which racks/equipment depend on it?”
- “Which service may be affected?”
- “What capacity/energy is compromised?”

Common Risk | Integration “works” but is not useful (events without context or misinterpreted metrics).

Best Practices |

- Define which signals truly go to operations (the ones actually monitored).
- Normalize units and thresholds from the beginning.
- Pilot: 1 room / 1 set of points --- validate --- scale.

CMDB to DCiM Integration (bidirectional)

What CMDB Provides (logical) |

- Service, application, owner, criticality, IT status.
- Logical relationships (CI to service dependencies).

What DCiM Provides (physical) |

- Location (site/room/rack/U).
- Electrical dependencies, physical ports/patching (if applicable).
- Physical context (environment, energy, capacity).

Real Value | Closing the gap between “what IT thinks it has” and “what physically exists and supports the service.”

Typical Risks |

- “Source of truth” conflicts (who governs which field?).
- Duplicate or overwritten data (naming, states, owners).
- Bidirectional integration without rules --- becomes “data ping-pong.”

Best Practices |

- Define ownership per field (this prevents 70% of problems).
 Typical example:
 Physical location --- DCIM is master.
 Owner/service/criticality---CMDB is master.
- Define a common ID (serial, asset tag, CI ID...) for reconciliation.
- Define synchronization events (add, remove, change, move).
- Acceptance criteria: “Key assets match and fields do not overwrite each other.”
- Design integrations based on real use cases, not technical possibilities.

When properly integrated, DCIM stops being just an infrastructure tool and becomes a decision-support system.

The Most Repeated Blockers

Prerequisites and Access

Before starting with DCIM, the foundation must be ready. Most blockers appear here:

- Architecture and platform ready (servers provisioned, validated versions, correctly sized resources).
- Licenses available and correctly assigned.
- Remote access operational (VPN, jump hosts, credentials).
- Communications enabled between networks (IT, OT, management).

Often these are not complex technical problems, but small dependencies never formally closed.

Rule:

This phase can be quite frustrating. A very effective way to unblock it is to quickly identify the right people within your organization and connect them directly with the vendor contacts who need that validation.

At Bjumper, at least, this is how we work: fewer cross-emails and more direct conversation – and if it can be done face to face, even better.

Data Collection for Modeling

Modeling depends directly on data quality.

- Updated inventory and layout + clarification sessions.
- Documented network connectivity + joint review to resolve inconsistencies.
- Electrical documentation (updated single-line diagrams).
- CAD drawings and physical details for accurate representation.

In practice, discrepancies between documentation and reality are normal. What matters is detecting them before loading the model.

Rule:

A DCiM does not fix bad data, it amplifies it. If the inventory is misaligned, the error does not disappear; on the contrary, it becomes more visible.

Monitoring: Measure Well Before Measuring Everything

A common mistake is integrating all signals from day one. For monitoring to work:

- Standard protocols must be correctly configured (e.g., properly enabled SNMP).
- SNMP tables or MIBs must be identified and documented.
- Credentials and communities validated.
- Critical signals tested before scaling.
- Decide which data truly adds value to the DCiM. Not everything that can be measured should be integrated.

For example:

It makes sense to integrate PDU-level consumption, branch load, and aisle temperature if the goal is to perform energy capacity planning. It does not make sense to start by integrating 200 secondary sensors if we are not yet clear about which KPIs we want to leverage.

Rule:

First “few signals, done well.” Then scale.

**Integrations: The Trick Is to Scope**

Integrating with the BMS, the CMDB, or any other system is not a milestone in itself. Technical integration is not the objective. The real milestone is having an operational use case up and running, with clear acceptance criteria.

For example:

- A change in the CMDB automatically updates the validated physical model.
- A critical BMS alarm generates a correlated and actionable event in the DCiM.
- A load expansion automatically impacts a capacity report.

If there is no concrete use case behind it, integration becomes a project “checkbox” with no real impact.

Rule:

Do not integrate just for the sake of integrating. Integrate to solve something specific.

Turning Risks into Actionable Milestones

A risk in a DCIM project does not disappear just because it is identified in a document. It is reduced when we turn it into something concrete, measurable, and assigned to a responsible owner.

The difference between a project that moves forward and one that gets blocked often lies here – in how we move from “this could be a problem” to “this has an owner and a deadline.” For each milestone, especially in data, monitoring, and integrations, the minimum that must exist is:

Owner:

A specific person, not a department, not “the IT team,” not “Facilities.” A first and last name are required because if something has no owner, in practice it does not exist; and if it has multiple owners, it is unlikely that anyone will prioritize it.

Date:

A realistic and agreed-upon date, not an imposed one. It is better to have a consensual deadline than an optimistic one that gets missed three times. It is also important to differentiate between:

- Technical delivery date.
- Functional validation date.

Because they do not always coincide.

Dependencies:

This is where most surprises appear, because an integration milestone may depend on:

- Firewall openings.
- Credential delivery.
- SNMP activation.
- Validation of a prior inventory.
- Availability of an intermediate environment.

If dependencies are not made explicit, the delay seems “unexplainable,” when in reality it was predictable from the beginning. Putting them in writing helps everyone understand what needs to happen before taking the next step.

Acceptance Criteria:

This is the most important point – and the one least often defined – because it is not enough to say “integration completed.” The real question is: what exactly must happen for it to be considered correct?

Examples:

- Are data being received?
- Are they received at the correct frequency?
- Are they properly mapped within the model?
- Do they generate coherent alarms?
- Can the end user effectively use them?

If we do not define this beforehand, validation turns into a subjective discussion.

Delivered ≠ Validated

Delivering something does not mean it has been validated.

- It may be installed, but not tested under real load.
- It may be integrated, but without consistent data.
- It may be modeled, but not reviewed by the room owner.

Validation implies that:

- The responsible person has reviewed it.
- It has been tested under real conditions.
- It meets the agreed acceptance criteria.

Until that happens, the risk is not closed. It has only been “moved.”

Turning risks into actionable milestones is not project bureaucracy – it is the most practical way to avoid long and silent blockages. When there is an owner, a deadline, clear dependencies, and defined acceptance criteria, the project gains traction and the DCIM stops being an ambitious initiative and becomes a controlled implementation.

Checklist

This checklist is not theoretical; it represents the minimum that should be closed before considering each block “complete.” If any of these points are missing, the issue will most likely surface later on.

 **Modelar****Naming and Hierarchy Standard**

Define how rooms, racks, PDUs, equipment, and circuits are named before starting to load information. A clear hierarchy must exist (site > room > row > rack > equipment, for example), along with consistent rules. If each area uses a different naming convention, the model will end up being inconsistent and difficult to maintain.

Complete Inventory/Layout Dataset

Not just an Excel file listing equipment – it must include:

- Exact physical location.
- Nominal power consumption (if applicable).
- Basic connectivity.
- Operational status.

The more structured the dataset is from the beginning, the less rework will be required later.

Minimum Electrical Documentation (Depending on Scope)

The highest level of detail is not always required, but at least the following must be available:

- Updated single-line diagrams.
- Clear identification of lines, protections, and redundancies.
- The relationship between each rack and its actual power supply.

If the scope includes energy capacity management, this point is no longer optional.

Joint Model Validation Session

Before considering the modeling phase closed, there must be a session with the person who truly knows the room. It is not enough for the model to “fit within the tool.” It must align with operational reality, as a joint review prevents many corrections later on.

 **Monitoring****List of Hardware Sources and Priority Signals**

It is not about integrating everything from the start. There must be a clear list defining:

- Which devices are integrated first (PDUs, UPS, cooling systems, etc.).
- Which signals are critical (consumption, load, temperature, status).
- Which signals are secondary and can wait.

This prevents overloading both the system and the team.

Connectivity / Network / Permissions

Many integrations fail not because of configuration issues, but due to incomplete permissions. For this reason, the following must be validated:

- SNMP access or other defined protocols.
- Required firewall openings.
- Confirmed credentials and permissions.
- Network latency and stability.

Validated Pilot (Units / Scale / Stability)

Before scaling, there must be at least one pilot operating correctly:

- Data consistent in units (kW, A, °C...).
- Correct data frequency.
- No intermittent losses.

If the pilot is not stable, scaling will only multiply the problem.

Scaling Plan

Define it in phases:

1. What will be integrated after the pilot.
2. In what order.
3. With which validation criteria.

Scaling without a plan usually generates noise and operational overload.



Integrations

This is where there is the greatest tendency to oversimplify, and where it is most important to be specific.

BMS with DCIM (unidirectional)

The following must be defined:

- Truly priority signals.
- Correct and consistent units.
- Thresholds aligned with real operations (not generic values).
- A prior pilot before mass integration.

Integrating 500 poorly parameterized signals adds no value.

Integrating 20 well-defined ones does.

CMDB with DCIM (bidirectional)

Here, control must be even stricter.

The following must exist:

- Field-level ownership (who is the master of each data field).
- A unique common identifier (shared and stable ID).
- Clear synchronization rules (unidirectional or bidirectional).
- A non-collision test before moving to production..

If it is not clearly defined who governs each field, synchronization will generate silent conflicts that may only be detected months later.

This checklist is not about bureaucracy, it is about avoiding rework, later discussions, and “we thought this was already closed.” When these points are clear, the project flows. When they are not, the DCiM stops being a solution and becomes a new source of incidents.

The objective of this document is to challenge the idea that a DCIM project is inherently complex to implement. If these steps are followed and there is collaboration between the implementer and the system user, the result will be a smooth deployment, delivering a technology that supports your processes and makes operations easier for people.

It will also provide visibility and control over the current and future state of the Data Center. What will be achieved:

1. Minimized risks.
2. Reduced costs.
3. Increased and improved Data Center service levels.