Bjumper

Thresholds and Alarms in the Data Center

Thresholds and Alarms in the Data Center

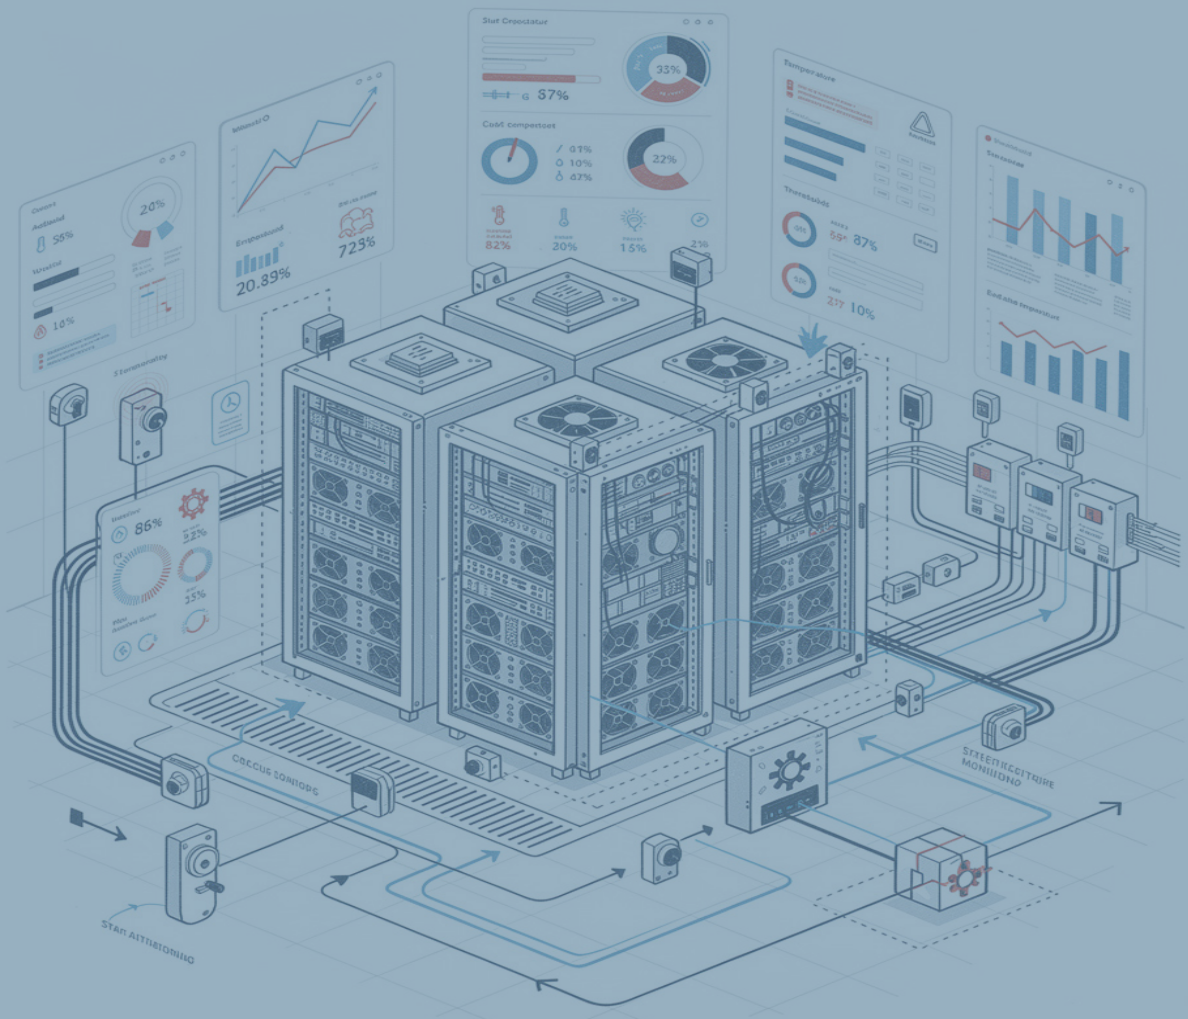

Introduction

A data center operates like a complex organism in which temperature, power, humidity, and connectivity are permanently interconnected. When any of these variables goes out of range, the chain of failures can be rapid and irreversible: from the silent degradation of components to the total loss of service.

A data center operates like a complex organism in which temperature, power, humidity, and connectivity are permanently interconnected. When any of these variables goes out of range, the chain of failures can be rapid and irreversible: from the silent degradation of components to the total loss of service.

Economic Context

According to Uptime Institute data, the average cost of an unplanned interruption in a data center exceeds EUR 9,000 per minute. Proper threshold and alarm configuration is, by far, the investment with the highest return in critical operations.

This guide covers three major monitoring domains:

- Environmental: temperature, humidity, and differential pressure.
- Electrical: UPS, PDU, power quality, and energy efficiency.
- Alarm management: key concepts, levels, escalation to the NOC, fatigue, and automation.

Reference Regulatory Framework

Every decision regarding thresholds must be based on recognized standards. The three that most influence daily operations are:

- ASHRAE TC 9.9 (5th ed., 2021): thermal and humidity parameters for data processing equipment.
- Uptime Institute Tier Standard: classifies resilience and determines the rigor of runbooks.
- ISA-18.2: alarm management philosophy for industrial control systems, applicable to BMS and DCiM.
- ITIL 4: IT service management framework. Defines the concepts of event, alarm, and incident, and their operational lifecycle.

Key Concepts: Metrics, Events, Alarms, and Incidents

Before configuring any threshold, it is essential to speak the same language. The industry uses terms that are frequently confused but have precise and differentiated meanings. Their correct distinction is the basis of a coherent monitoring architecture.

Hierarchy of Operational Concepts

Concept	Definition	Managed By
Metric	Continuous measurement of a physical or logical parameter. Does not imply any action	DCiM / EMS (automatic)
Event	State change detected by the system. May be informative or require action	EMS / DCiM (automatic)
Alarm	Event that exceeds a defined threshold and generates active notification. Requires acknowledgment	NOC Operator
Incident	Alarm that is not resolved within the defined time or that causes service impact. Managed through an ITSM ticket	Technical Team + NOC

Metrics: What Is Measured and How Often

A metric is the raw data collected from sensors, devices, or APIs. By itself, it does not indicate whether there is a problem: it is the threshold that turns a metric into a relevant event. Metrics are classified according to their nature:

- **Environmental metrics:** rack inlet/outlet temperature, relative humidity, dew point, differential pressure, airflow, water detection.
- **Electrical metrics:** AC/DC voltage, current per phase, active power (W) and apparent power (kVA), power factor, THD, accumulated energy (kWh), UPS battery autonomy.
- **Efficiency metrics:** PUE, WUE (Water Usage Effectiveness), CUE (Carbon Usage Effectiveness).

Sampling frequency is critical. Temperature sensors usually sample every 30–60 seconds, while electrical current sensors may do so every 5–10 seconds in critical facilities.

Events: The Raw Material of the NOC

An event is any state change that the system detects and records. Not all events are alarms: most are information stored for historical analysis. Event types are:

- **Informational:** the system records the change without notifying anyone. For example, temperature rising from 22°C to 23°C within the normal range.
- **Warning:** the parameter approaches a threshold. The shift operator is notified for follow-up.
- **Critical:** the parameter exceeds the defined threshold. Requires acknowledgment and action within a defined time.

Event management in the NOC relies on correlation tools that group multiple related events into a single incident. If five racks in the same cold aisle simultaneously exceed the temperature threshold, the system should generate a single incident, not five independent alarms.

Alarms: From Signal to Action

An alarm is an event that crosses a threshold and requires human intervention. Its effectiveness depends on being actionable: it must tell the operator what happened, where, and what needs to be done. The attributes of a well-designed alarm are:

- **Identification:** device name, physical location (room, row, rack, position), and affected parameter.
- **Current value vs. threshold:** the operator must immediately see how much the limit has been exceeded.
- **Severity:** classified unambiguously (Informational / Warning / Critical / Emergency).
- **Associated runbook:** direct link to the response procedure.
- **Lifetime:** how long the alarm has been active and whether it has been acknowledged.

Incidents and Severitiess

An incident is an alarm that is not resolved within the stipulated time or that impacts service availability. Its management is formalized through a ticket in the ITSM system with an owner, resolution SLA, and full traceability. Severity classification combines two dimensions: impact and urgency.

Severity	Name	Impact	Urgency	Resp. (min)	Resolution
SEV 1 / P1	Critical	Critical service down. Immediate business impact	Maximum	< 5	< 1 h
SEV 2 / P2	High	Significant degradation. Reduced functionality	High	< 15	< 4 h
SEV 3 / P3	Medium	Limited impact. Workaround available	Medium	< 30	< 8 h
SEV 4 / P4	Low	No immediate operational impact. Improvement or prevention	Low	< 120	< 5 days
SEV 5 / P5	Informational	Record without required action	None	_____	_____

Incidents and Severities

Not all alarms should reach the NOC. Prior filtering is essential to avoid operational fatigue. All warning-level or higher events that cannot be automatically resolved are sent to the NOC.

Environmental Domain

- Rack inlet temperature > 28°C (Warning) or > 32°C (Critical).
- Relative humidity outside the 30–70% RH range.
- Temperature or humidity sensor failure (signal loss).
- Water detection at any point on the raised floor or under CRAC/CRAH units.
- Cold aisle differential pressure < 8 Pa (possible containment breach).
- CRAC/CRAH failure or high internal temperature alarm in the cooling unit.

Electrical Domain

- UPS operating on battery mode (utility power failure).
- UPS battery autonomy < 10 minutes.
- UPS load > 85% of nominal capacity.
- Voltage out of range ±10% from nominal.
- Phase imbalance > 15% in any PDU.
- PDU circuit > 90% of nominal capacity (risk of breaker trip).
- Generator failure or inability to start during testing.

IT / Hardware Domain

- Server in thermal throttling state (CPU reduces frequency due to excessive temperature).
- Server shutdown due to thermal protection (iDRAC/iLO action).
- Loss of network connectivity with monitoring equipment.
- Internal server fan failure (hardware alarm via IPMI/Redfish).

Practical Rule for Escalation to the NOC

Every critical or emergency event must reach the NOC without exception. Warning events only reach the NOC if they exceed the configured persistence time (dwell) without being automatically resolved, or if they occur during periods without on-site staff.

Temperature Thresholds

Temperature is the most monitored variable in any server room. Incorrect management generates two equally costly problems: overheating and overcooling, which wastes between 30% and 40% of cooling energy.

ASHRAE TC 9.9 Standard: Recommended and Allowable Ranges

ASHRAE distinguishes between the recommended range, optimal for equipment lifespan, and the allowable range within which the manufacturer guarantees operation, although wear increases.

Class	Recommended Range (C)	Allowable Range (C)	Typical Application
A1	18 - 27	15 - 32	Enterprise servers and critical storage
A2	18 - 27	10 - 35	Volume servers, workstations
A3	18 - 27	5 - 40	PCs, laboratory equipment
A4	18 - 27	5 - 45	Industrial equipment, maximum flexibility
H1 (HPC / IA)	18 - 22	15 - 25	AI systems, GPU clusters, HPC

Measurements are always taken at the server air inlet. Monitoring only the general environment may mask hot spots in specific racks that exceed the recommended limit by 8–10°C.

Practical Configuration of Temperature Alarms

Three-level alarm configuration — warning, critical, and emergency — allows the response to be escalated progressively without overwhelming the operations team.

Level	Rack Inlet Threshold (C)	Required Action	Response Time
Normal (OK)	18 - 26	None	____
Warning	26 - 28	Check cooling system status	< 30 min
Critical	28 - 32	Immediate intervention. Activate redundant CRAC	< 10 min
Emergency	> 32 / < 15	Controlled shutdown. Escalate to the NOC as SEV 1	< 5 min

Sensor Placement: Where to Measure

ASHRAE recommends a minimum of 6 sensors per rack: top, middle, and bottom on both the front and rear sides. In practice, many operators start with 3 points per rack (front, top/middle/bottom) and add additional coverage at the ends of each row, which are the points most prone to hot air recirculation.

Placement Rule

Avoid placing sensors in areas with direct airflow, in front of CRAC grilles, near frequently opened doors, or exposed to direct sunlight. These positions generate erroneous readings that trigger false alarms.

Humidity and Dew Point Thresholds

Relative humidity (RH) controls two opposing risks: if it is too low, it promotes electrostatic discharge (ESD); if it is too high, it causes condensation on circuits and metal surfaces.

ASHRAE Reference Ranges

Parameter	Recommended Range	Allowable Range (A1/A2)	Risk Outside Range
Relative humidity	50 - 60 % HR	8 - 80 % HR	ESD (<30 %) / Condensación (>70 %)
Dew point	-9 C a 15 C	-12 C a 17 C	Condensation on components
Rate of change	< 5 C / 20 h	____	Thermal shock on components

Humidity alarm threshold configuration

Level	Relative Humidity	Action
Normal	40 - 60 % HR	None
Low warning	30 - 40 % HR	Check humidifier. Assess ESD risk
High warning	60 - 70 % HR	Check dehumidifier and cooling load
Low critical	< 30 % HR	Activate emergency humidification
High critical	> 70 % HR	Physical inspection. Condensation risk

Configuration of Humidity Alarm Thresholds

Relative humidity is a temperature-relative measurement: the same vapor content in the air produces very different RH readings if the temperature changes. Dew point, however, is absolute: it directly indicates the temperature at which vapor condenses. This value is what truly predicts the risk of condensation when the cold air from cold aisles comes into contact with warmer metal surfaces.

The standard deployment ratio is one humidity sensor for every five racks, since humidity varies more slowly than temperature within the room.

Differential Pressure and Airflow

In facilities with hot/cold aisles and aisle containment systems, differential pressure (ΔP) between the cold aisle and the general room space is the indicator that best reflects air distribution efficiency. A drop in ΔP indicates leakage: cold air escapes through gaps or incorrectly positioned tiles, directly degrading server inlet temperature.

Configuration of Differential Pressure Alarm Thresholds

Situation	Typical Delta P (Pa)	Indication
Optimal (contained aisle)	12 - 25 Pa	Correct distribution, no significant leaks
Warning	< 10 Pa	Possible leaks in seals or tiles. Inspect
Critical	< 5 Pa	Containment failure. Risk of hot air recirculation
Overpressure	> 35 Pa	Check fan/CRAH speed

Airflow Sensors

Airflow sensors are installed at cold air supply points and hot air returns to calculate volumetric distribution. An airflow anomaly combined with an increase in rack inlet temperature confirms a containment problem before servers enter thermal protection.

Reaction Time in HVAC Failure

According to the ASHRAE TC 9.9 specification, when the HVAC system fails, the cold aisle temperature can rise by 30°C in as little as 5 minutes in high-density installations. Ride-through time — the margin before servers shut down due to thermal protection — is the critical parameter that defines escalation urgency.

Electrical Power Supply Thresholds

The power supply chain is the domain that most directly determines availability. Its thresholds not only protect equipment but also allow efficient operation, avoiding over-dimensioning the infrastructure

UPS Monitoring

The following table summarizes reference operating thresholds for generic installations. These are recommended management values, independent of manufacturer-specific thresholds.

UPS Metric	Normal	Warning	Critical
Output voltage (VAC)	220 - 230 V	210 - 219 V / 231 - 240 V	< 210 V o > 240 V
Load (%)	0 - 75 %	75 - 85 %	> 85 %
Battery charge (%)	90 - 100 %	50 - 90 %	< 50 %
Battery autonomy	> 15 min	5 - 15 min	< 5 min
Battery temperature (C)	20 - 25 C	25 - 30 C	> 30 C / < 15 C

Smart PDU Monitoring

Smart PDUs (iPDUs) are the most granular measurement point in the electrical chain. They allow threshold configuration at circuit (infeed), phase, and even individual outlet level. The values in the following table are generic reference operating thresholds. The metrics that must be monitored are:

PDU Metric	Normal	Warning	Critical
Current per phase (A)	< 80 % of nominal	80 - 90 % of nominal	> 90 % of nominal
Phase balance (%)	< 10 %	10 - 20 %	> 20 %
Input voltage (V)	220 - 230 V	210 - 219 V	< 210 V o > 240 V
Power factor	0,95 - 1,0	0,85 - 0,95	< 0,85
THD (Harmonic distortion %)	< 5 %	5 - 10 %	> 10 %

Exceeding 80% of a circuit's capacity is the industry-standard warning signal because it guarantees margin for transient peaks and prevents breaker trips, which generate immediate and uncontrolled interruptions.

Water and Leak Detection

Water is one of the most underestimated risks in the data center. Leaks in water-cooling systems can cause short circuits and irreversible damage within minutes..

Sensor Types and Placement

- Point sensors: detect the presence of water at a specific location. Installed at the lowest points of the raised floor, under cooling units, and near condensation trays.
- Linear detection cables: continuously cover large surfaces. Ideal for room perimeters and underfloor aisles. A single cable can cover dozens of meters.

Response Protocol

Water detection has no intermediate levels: any positive signal is a critical event (SEV 1) requiring immediate notification through all channels and automatic ticket creation with the attached response protocol.

Leak Response Protocol

1. Identify the source (visually or by the position of the activated sensor).
2. Shut off the water supply to the affected circuit if possible.
3. Assess whether there is equipment at immediate risk.
4. Never shut down equipment due to proximity to water without confirming actual contact.

Alarm Management KPIs

The maturity of the alarm system is measured with specific indicators. The most commonly used in the sector are:

KPI	Definition	Recommended Objective
MTTD	Mean Time to Detect: average time to detect the event	< 2 min
MTTA	Mean Time to Acknowledge: time until acknowledgment of critical alarm	< 5 min
MTTR	Mean Time to Resolve: average incident resolution time	< 60 min (SEV 1
False alarm ratio	Percentage of alarms that do not correspond to real events	< 5 %
Alarms per shift (8 h)	Number of active alarms per operational shift	< 30 per shift

Alarm Management Architecture and NOC

Having well-configured sensors is not enough if the alarm architecture is not designed so that the right information reaches the right person at the right time. Poor alarm management leads to operational fatigue: the team begins to ignore notifications because there are too many or they are too insignificant.

Severity Levels and Color Coding

Level	Color	Operational Definition	Response Time
Informational	Blue	Event recorded without operational impact. No active notification	——
Warning	Yellow	Parameter approaching the limit. Requires scheduled review	< 30 min
Critical	Orange	Parameter outside safe operating range. Intervention required	< 10 min
Emergency	Red	Imminent risk of service loss or physical damage	< 5 min

Escalation Protocol to the NOC

Automatic escalation guarantees that no critical alarm goes unanswered. Typical escalation logic works in three stages:

- **Level 1** (0–5 min): notification to the on-duty operator via email and SMS. If there is no acknowledgment within 5 minutes, it escalates automatically.
- **Level 2** (5–10 min): notification to the technical shift supervisor and automatic ITSM ticket creation with event data and response runbook.
- **Level 3** (> 10 min without resolution): notification to the facilities manager or technical management. In emergency events, escalation to level 3 is immediate.

Maintenance and Suppression Windows

During planned work, DCIM systems allow maintenance mode to be activated to suppress alarms on specific devices without disabling global monitoring. This avoids noise during maintenance windows and the risk of overlooking real events in the rest of the infrastructure.

NOC Escalation Protocol

Alarm fatigue occurs when the volume of notifications exceeds the processing capacity of the operations team. The result is that alerts begin to be ignored, including critical ones with serious consequences for availability.
The most common causes are:

- Thresholds that are too narrow, generating alarms due to normal fluctuations.
- Absence of persistence time (dwell), triggering isolated events without real relevance.
- No deduplication: multiple sensors alerting independently about the same event.
- Informational alarms sent through the same channels as critical alarms.

Industry rule: more than 1 alarm per minute during a normal operational shift indicates that the alarm architecture requires review.

DCIM / BMS / ITSM Integration

- **DCIM**: consolidates environmental and power data, applies threshold logic, and generates standardized alarms.
- **BMS:** manages mechanical systems (HVAC, pressure groups, fire detection) and shares data with the DCiM.
- **ITSM**: receives alarms from the DCiM and generates tickets with owner, runbook, and resolution SLA.

Monitoring Architecture: Deployment Levels

The monitoring architecture is not unique: each installation requires a model adapted to its size, criticality, and resources. There are three reference models that scale progressively.

Autonomous Model (Single Room)

Suitable for small installations, fewer than 20 racks, or remote telecommunications rooms without permanent staff. A local controller collects data from temperature, humidity, and current sensors and sends alerts via email or SMS. There is no integration with higher-level platforms.

- Advantages: consolidated view, alarm deduplication, centralized reporting.
- Limitations: requires stable connectivity between sites and greater configuration complexity.

Network Model (Multiple Rooms or Buildings)

Local gateways normalize and forward telemetry to a central platform. This model enables unified visibility, alarm correlation between locations, and centralized QA/QC (quality control of received data).

- Advantages: consolidated view, alarm deduplication, centralized reporting.
- Limitations: requires stable connectivity between sites and greater configuration complexity.

Fully Integrated Model

The DCIM platform is integrated with the building's BMS and the corporate ITSM. Alarms automatically generate tickets with the included runbook, dashboards show real-time status to operators, facilities managers, and management, and periodic reports document deviations from agreed SLAs.

- Advantages: maximum traceability, automated response, audit support.
- Limitations: higher initial investment, requires data integration between heterogeneous systems.

Operational Best Practices

Sensor Calibration and Maintenance

A miscalibrated sensor is worse than having no sensor at all: it generates false confidence. Temperature and humidity sensors must be calibrated at least once a year or whenever data shows systematic deviations from reference measurements. Current sensors require semiannual verification in critical facilities.

Periodic Threshold Review

Thresholds configured during commissioning may become inadequate as the data center load changes. Quarterly threshold review and semiannual validation through simulation exercises are practices recommended by the Uptime Institute.

Operational Runbooks

Without a runbook, the response depends on the individual knowledge of whoever is on duty, which introduces variability and lengthens resolution time. This is one of the factors that feeds the operational fatigue described in chapter 10.3: if each incident is handled differently, the team spends more time deciding what to do than actually doing it. A well-maintained catalog of runbooks is the most effective antidote.

Alarm Management KPIs

The key performance indicators for alarm management (MTTD, MTTA, MTTR, false alarm ratio, and alarms per shift) are defined and tabulated in chapter 10.5, together with the recommended objectives for each. It is recommended to consult that section as the central operational reference.

Sources and References:

- ASHRAE TC 9.9: Thermal Guidelines for Data Processing Environments, 5.a ed. (2021). ashrae.org
- Uptime Institute: Global Data Center Survey 2024. journal.uptimeinstitute.com
- MFE-IS: Data Center Environmental Monitoring: An In-Depth Guide. mfe-is.com
- INOC: 5 Key Processes for Network Operations Centers. inoc.com envigilance.com: ASHRAE TC 9.9 - Data Center Thermal Guide (2026).
- U.S. Department of Energy: Data Center Metering and Resource Guide. datacenters.lbl.gov