HyperNext Data Centers

RESEARCH PAPER

HN-RP-004

# HyperNext BMS

## Supervisory control for Tier IV AI infrastructure

The 3-minute policy. Seven-layer security. What changes when AI workloads run on top.

This is what is Next.


| | |
|---|---|
| Series | HyperNext Research |
| Paper | HN-RP-004 |
| Issued | 20 February 2026 |
| Version | 1.0 |
| Classification | Public release |
| Citation | HyperNext Research, HN-RP-004 |

# HyperNext BMS

This paper is part of the HyperNext Research series. Methodology, assumptions, and source data are stated openly so other operators can reproduce the analysis on their own facilities. Citation as "HyperNext Research, HN-RP-004" is welcome.

## Contents

# 1. What a BMS is for, and what AI workloads change

ABSTRACT

The traditional Building Management System was designed for environments where the workloads it served were unaware of the building infrastructure and largely indifferent to it. AI workloads have changed that. A rack-scale GPU system creates thermal and power dynamics that need sub-second visibility, predictive intervention, and supervisory control reaching across subsystems that used to run independently. This paper describes the HyperNext BMS, the supervisory platform that operates the Phase 1 Hyderabad campus and will scale to the 1.2 GW Kakinada AI Factory. It covers the operating philosophy (3-minute acknowledgement policy. Seven-layer security model. Role-based access). The technical implementation (Modbus, BACnet, SNMP, OPC-UA convergence into a single supervisory data plane). And operational discipline: alarm escalation matrix, tenant SLA handling, predictive maintenance triggers. The paper is intended for facility engineers, BMS integrators, and IT operations teams evaluating their own infrastructure for AI readiness.

A Tier IV data centre running AI workloads cannot be operated like a traditional data centre any more. The BMS has to evolve with it.

## What "supervisory" means

The HyperNext BMS is a supervisory platform, not a control platform. It reads field state from underlying controllers. Modbus PLCs on the mechanical side. BACnet/IP devices on the environmental side. SNMP-capable equipment on the IT-adjacent side. OPC-UA bridges into the SCADA layer. The BMS aggregates this state into a single coherent operational view. Control actions are not initiated from the BMS in normal operating mode. The underlying controllers handle them according to pre-engineered logic. The BMS observes, records, and alerts on results.

This separation is deliberate. Life-safety actuation (fire suppression, breaker tripping under fault, evacuation initiation) belongs to the dedicated panels designed for that purpose, not the supervisory layer. The BMS sees those events and records them. It does not cause them. When the BMS does initiate a control action (an operator requesting a genset changeover from the supervisory HMI, for instance) that action is gated by role-based permission, dual control where appropriate, and a complete audit trail.

The supervisory pattern is what makes the BMS auditable in the regulatory sense. Every action is attributed to a named operator, recorded with timestamp, and replayable from the audit log. The architecture is built for environments where forensic reconstruction of an incident must be possible months after the fact, including the specific question of who authorised which override and on what evidence.

## What AI workloads change

A traditional enterprise data centre running 10 to 20 kW racks operates with thermal time constants measured in minutes. A rack going from idle to full load takes several minutes for the cooling response to stabilise. CRAC units cycle on a slow loop. The supervisory system can poll on a 30-second interval and miss nothing important.

An AI data centre running 600 kW racks operates with thermal time constants measured in seconds. A Vera Rubin Ultra NVL576 cabinet going from inference idle to full inference load draws an additional 380 kW in under five seconds. The cooling response, even with direct-to-chip liquid cooling and aggressive coolant distribution unit pumping, has a measurable lag. If the supervisory system polls on a 30-second interval, it will miss thermal excursions that already started and may have already produced rack-level interventions before the BMS knows the event happened.

The HyperNext BMS therefore runs a sub-second supervisory cycle for thermal-critical points, a one-second cycle for power-system points, and a five-second cycle for environmental and security points. The data ingestion architecture sustains this rate continuously across approximately 12,000 monitored points per phase, with sub-second alarming on threshold violations and 30-day retention at full sample rate.

## The digital twin layer

Every HyperNext data hall is mirrored by a live digital twin. The twin is a 3D model of the physical hall fed by the same telemetry the BMS ingests: thermal sensors, airflow, power draw, humidity, leak detection. Operators see hotspots forming in the digital twin before the underlying alarm thresholds trip. This is preventive operations, not just reactive monitoring. A rack-level temperature gradient that would have produced an alarm in 8 minutes is visible in the twin at minute zero. The intervention window is wider.

The twin is also the substrate on which what-if simulations run. Before any change to the cooling setpoint, the airflow management, or the rack power profile, the BMS team can simulate the change in the digital twin against the current state and compare expected behaviour with the live system. This catches misconfiguration before it reaches production.

# 2. The 3-minute policy

## Why three minutes

The most consequential operational rule in the HyperNext BMS is the three-minute acknowledgement policy. Every active alarm condition must be acknowledged by an operator within three minutes of being raised. If three minutes pass without acknowledgement, the alarm auto-escalates. The row in the Warning Center turns red and is tagged ESCALATED. A notification fires to the Shift Lead. An audit entry is written for senior attention.

The three-minute window is the outcome of two competing pressures. It must be long enough that an operator handling one alarm has reasonable time to look at it, classify it, run the matching playbook, and click acknowledge. It must be short enough that an alarm cannot sit unattended through anything that matters operationally. Three minutes is the threshold we landed on after analysis of historical incident response times. Below three minutes, false-escalation rates rise sharply because operators are still doing useful triage work when the auto-escalation fires. Above five minutes, the rate of incidents where the alarm sat unattended long enough to matter rises. Three minutes is the elbow of the curve.

## The senior-approval requirement for criticals

A second discipline layered on top of the 3-minute policy is the senior-approval requirement for critical alarms. A warning-severity alarm can be acknowledged by any on-duty operator. A critical-severity alarm cannot. The acknowledgement of a critical opens a senior-approval dialog requiring a Shift Lead, Facilities Manager, or Chief Manager to click Approve override before the alarm leaves the active queue.

The dialog is audit-logged with both the operator identity and the senior identity, with timestamp, and with the override reason written in free text. The audit log is immutable. It cannot be edited after the fact. Every entry has cryptographic chaining so tampering with historical entries is detectable. The combined effect: a critical alarm cannot be hidden, cannot be acknowledged without senior visibility, and cannot have its handling rewritten after the incident is closed.

## The escalation matrix

Alarm severity and the corresponding response time are formalised in a four-tier escalation matrix.

| Severity | Examples | First notify | Within | Escalate to |
|---|---|---|---|---|
| Critical | Fire/VESDA, utility loss with genset fail, UPS depleting, SLA breach, security breach | NOC Operator, then Shift Lead | 3 min | Facilities Manager, then Chief Manager. OEM. Emergency services. |

| | | | | |
|---|---|---|---|---|
| Major | Chiller trip, UPS on battery, day-tank low, device offline affecting redundancy | NOC Operator, then Shift Lead | 15 min | Facilities or Engineering Manager. Vendor. |
| Minor | Single sensor fault, warranty expiring, nuisance alarm | NOC Operator | 1 h or next shift | Shift Lead. Maintenance partner. |
| Info | Scheduled events, PM completions, configuration changes | Logged to audit | n/a | Reviewed at shift handover |

The escalation matrix is not advisory. The BMS itself enforces it. Timers are computed against the alarm raise time. Auto-escalation fires automatically at the threshold. The audit log records every escalation event. The matrix is operational infrastructure, not a documentation artifact.

# 3. The seven-layer security model

## Why seven layers

Physical security in a Tier IV facility runs in concentric layers, each needing independent authentication and producing its own audit trail. The HyperNext model uses seven layers, working outward from the rack to the perimeter.

1. Rack or Cabinet lock. Tenant-controlled mechanical and electronic locks on the rack itself. Biometric or key-card override only by HyperNext personnel under dual-control authorisation.
2. Cage or Cabinet row. Tenant cage locks where tenancy is on a per-cage basis. Electronic access with full audit log.
3. Data hall access. Dual-authentication entry (biometric and card) to each individual hall. Escort policy for non-HyperNext personnel.
4. Data centre floor mantrap. Single-occupancy mantrap with weight-and-card verification. Defeats tailgating.
5. Building lobby and reception. Visitor management with pre-registration, host approval, photo capture, and badge issue.
6. Facility entrance. Turnstile-based access for staff. Gate-controlled vehicle and pedestrian flows.
7. Perimeter fence and gate. ANPR-controlled vehicle gate. Heavy-vehicle radar. Perimeter intrusion detection.

Each layer maintains its own access log. The BMS aggregates them into a single per-visitor or per-staff timeline. A staff member arriving for a 14:00 shift has a complete trace through all seven layers: gate badge-in at 13:42, lobby at 13:46, mantrap at 13:51, hall 2 access at 13:53, cage access at 13:54, and so on. Anomalies (someone present in a hall who has not passed through the mantrap, or a card used at two layers simultaneously) generate immediate security alarms.

## AI camera analytics

The seven-layer model is reinforced by camera-based AI analytics watching for behavioural anomalies that policy alone cannot catch. Tailgating detection (one card swipe followed by two people through the door) runs continuously on every controlled portal. Loitering detection runs at gates and lobby zones. PPE compliance detection runs in the data halls where appropriate. Heavy-vehicle movement is tracked through ANPR cross-referenced against the scheduled delivery list.

The analytics produce alarms for human supervisor review, not for automatic action. A tailgating detection does not lock a door. It raises an alarm that the SOC operator investigates within thirty seconds. This is deliberate. False positives in AI-based security detection are common, and an automated response model would generate too many false-positive lockdowns to be operationally viable. The supervisor-in-the-loop

pattern is the right architecture for current camera-analytics maturity. We expect to revisit it as the underlying models improve.

# 4. SLA handling

## The SLA scenario workflow

The most consequential daily output of the BMS, for the business, is the SLA performance of the platform across its tenants. The platform supports approximately 10 tenants in Phase 1 Hyderabad, with contracted service tiers from 99.9% to 99.999% availability, with associated credit schedules for breaches.

The SLA workflow runs in seven steps from breach detection through credit issue and prevention. Each step is supported by a specific BMS screen and produces specific audit records.

1. Detect. The SLA Tenants screen shows a tenant transitioning to AT RISK state, or an alarm flags a tenant hall. The BMS computes the running uptime against the contracted tier in real time.
2. Classify. The operator identifies which tenants are affected and what contracted tier applies. The incident severity is set accordingly.
3. Root cause. The tenant drill-down shows the implicated asset (chiller B4-CH1 if it was a cooling event, for example). The BMS surfaces historical reliability of that asset.
4. Contain and recover. The matching playbook runs (switch to the standby chilled-water train, for instance). Each step is checkboxed in the incident record.
5. Notify the tenant. A breach notice goes out via the Tenant Alerts workflow, stating the cause and the service-credit position. Tenant acknowledgement is captured.
6. Credits and RCA. Service credits are applied per the agreement. The Root Cause Analysis is completed within the contractual window (typically 5 business days).
7. Prevent. The reliability suggestion is actioned: add N+1, shorten PM interval, refurbish the implicated asset. This step prevents recurrence.

## A worked example

Quantum Bank, a 99.99% tier tenant in B4 hall 1, shows AT RISK status after a 17-minute cooling excursion. The BMS drill-down shows the implicated asset is B4-CH1, a Trane chiller that has been involved in three SLA-impacting events over the past nine months. The operator switches to the standby chilled-water train, confirms hall inlet temperatures recover in the 3D twin, raises an incident with all containment steps logged, sends Quantum a breach notice with the cause and the credit position, and marks Reliability for follow-up.

The Reliability follow-up in this case is to plan an additional chiller (N+2 instead of N+1) for the B4 hall during the next capital cycle, and to shorten the PM interval on the existing chillers from 6 months to 4 months in the interim. Historical reliability data on B4-CH1 supports both interventions on cost-benefit grounds.

# 5. Predictive maintenance

## From schedule-driven to condition-driven

Traditional data centre maintenance is schedule-driven. A chiller gets a major PM every 12 months. A UPS gets a battery test every 6 months. A genset gets a full inspection every 24 months. The schedules come from OEM specifications and are tuned by industry experience over decades. They work. They are also crude. Some assets need attention more often than the schedule implies. Others would run safely for longer.

The HyperNext BMS supplements schedule-driven PM with a condition-driven approach for the major mechanical and electrical assets. For each asset class, the BMS computes a health index from operational telemetry (run hours, load profile, ambient conditions), failure-precursor signals (vibration, temperature, current signature, oil quality where applicable), and historical comparison against similar assets in the fleet. The health index is updated continuously and drives maintenance dispatch decisions.

The output is a Remaining Useful Life (RUL) estimate per asset, expressed in operating hours, and a prioritised maintenance queue. Assets with low RUL are scheduled for refurbishment or replacement ahead of the OEM schedule. Assets with high RUL have their PM intervals extended. The aggregate effect, validated against industry benchmarks, is roughly 15 to 20 percent reduction in maintenance hours per megawatt-year of operation while simultaneously reducing the unplanned-failure rate.

## What we still do on a schedule

Not all maintenance can be condition-driven. The dominant exceptions are safety-critical tests that have to happen regardless of condition (fire suppression discharge testing, UPS load-bank testing) and regulatory inspections that are calendar-bound (statutory pressure-vessel inspections, electrical safety audits). These continue on their fixed schedules with no condition-based modification. The BMS schedules them and triggers the related tenant PM notices but does not attempt to optimise their interval.

HEADLINES

- The HyperNext BMS is a supervisory platform that reads underlying controllers and produces a single coherent operational view, with control actions gated by RBAC and dual-control where appropriate.
- The 3-minute acknowledgement policy with senior-approval for criticals enforces a discipline that scales beyond the cognitive capacity of any individual operator.
- The seven-layer physical security model with AI-camera reinforcement provides the audit-grade access control that Tier IV operations need.
- SLA handling runs as a formal workflow with seven steps, traceable from breach detection through prevention action, and audit-logged at every stage.
- Predictive maintenance based on continuous health-index computation reduces maintenance hours by 15 to 20 percent while reducing unplanned failures, complementing rather than replacing safety-critical scheduled tests.

This paper is a condensed version of the operational reference. The full Operations Manual (HN-BMS-OPS-MANUAL v3.0) is 46 pages and covers per-screen behaviour, what-if scenario playbooks, the alarm escalation matrix, and the complete SLA handling workflow. The manual is available to BMS customers under NDA.

# 6. Protocol architecture

The HyperNext BMS converges four field protocols into a single supervisory data plane. The convergence is not an off-the-shelf integration. Each protocol has its quirks. The data semantics differ. The trust model differs. The section below walks through how the integration is done and why each protocol is used where it is.

## Modbus TCP

Modbus TCP is the primary protocol for mechanical and electrical equipment. The HyperNext deployment uses approximately 6,400 Modbus TCP devices across the Phase 1 Hyderabad campus, including all chillers, all CRAHs, all UPS systems, all PDUs, all switchgear, all transformers, and all gensets.

The strengths of Modbus are simplicity (16-bit register read and write, no schema beyond device documentation), wide vendor support, and predictable timing. The weaknesses are the lack of any built-in security (Modbus TCP has no authentication, no encryption, no integrity check) and the absence of standardised data semantics (the same chiller register can mean different things on different vendor implementations).

The HyperNext approach handles the security gap by network isolation. Modbus TCP traffic is confined to a dedicated VLAN (VLAN 401 in the Phase 1 deployment) that is not reachable from any IT or tenant network. The supervisory data plane sits on a separate VLAN with a controlled inter-VLAN gateway that proxies and validates traffic.

The semantic gap is handled by device-specific drivers. Each vendor and model has a JSON manifest in the BMS device registry that defines what each register means, what the engineering units are, and what the alarm thresholds should default to. Onboarding a new vendor model is a manifest-writing exercise plus a validation pass.

## BACnet/IP

BACnet/IP is the protocol for environmental sensors and CRAH-level controls. It is the default for building HVAC and the HyperNext deployment uses BACnet/IP for approximately 3,200 devices, primarily temperature sensors, humidity sensors, pressure sensors, and the CRAH-level fan control interface.

The strengths are richer data semantics than Modbus (BACnet has standardised object types, properties, and units) and built-in discovery (devices announce themselves on the network). The weaknesses are higher protocol overhead (BACnet has more on-the-wire structure than Modbus) and a similarly weak security model (BACnet/IP has no authentication by default, with optional BACnet-SC providing TLS but not widely deployed).

The HyperNext BACnet/IP traffic is on VLAN 402 with the same network isolation approach as Modbus. The discovery protocol is limited to the supervisory network: tenant or IT networks cannot discover or address BACnet devices.

## SNMP v3

SNMP is the protocol for IT-adjacent equipment that has historically been managed by IT teams rather than facilities teams. The HyperNext deployment uses SNMPv3 for approximately 800 devices, primarily network switches in the supervisory fabric, UPS systems whose vendor provides SNMP in addition to Modbus, and gateway devices that bridge Modbus/BACnet onto IP networks.

SNMPv3 (unlike v1 and v2c) includes authentication and encryption. The HyperNext deployment uses HMAC-SHA-512 for authentication and AES-256 for encryption. Read community strings are rotated quarterly. Write access via SNMP is disabled across the entire deployment: configuration changes happen through device-specific management interfaces with separate auth, not via SNMP set.

## OPC-UA

OPC-UA is the protocol for the bridge into the broader SCADA layer (the campus-level controls that integrate with the building automation system) and for high-density data interfaces with select OEMs (Schneider EcoStruxure, Siemens Desigo) that ship OPC-UA endpoints natively.

OPC-UA has the strongest security model of the four protocols. Mutual TLS authentication, certificate-based identity, and standardised security policies (Basic256Sha256 minimum in the HyperNext deployment). The data model is also the richest: nested address spaces, type hierarchies, and structured access controls.

The HyperNext deployment uses OPC-UA on a dedicated security zone (VLAN 410) with certificate rotation every 90 days. Approximately 200 endpoints are exposed, primarily aggregation points rather than individual devices. The supervisor layer subscribes to OPC-UA event streams for everything that changes.

## Data plane architecture

The four protocols converge in a Kafka-based ingestion layer. Each protocol has its own gateway service that translates from the wire format into a unified internal event schema. The schema is:

```
INTERNAL EVENT SCHEMA
{
  "ts":             ISO 8601 with millisecond precision and timezone offset,
  "device_id":      HyperNext canonical device identifier (HN-####),
  "facility":       campus identifier (e.g. "hyd-1"),
  "building":       building identifier (e.g. "b2"),
```

```
    "floor":          floor identifier (e.g. "f3"),
    "asset_type":     enumerated set (chiller, ups, crah, transformer, ...),
    "vendor":         vendor name (e.g. "trane", "schneider"),
    "point":          canonical point name (e.g. "chw_supply_temp"),
    "value":          numeric or boolean,
    "engineering_unit": SI unit string,
    "quality":        OPC-UA quality code (good, bad, uncertain, ...),
    "source_protocol": wire protocol (modbus_tcp, bacnet_ip, snmp_v3, opc_ua)
  }
```

Approximately 12,000 monitored points produce events at a peak rate of 4,000 events per second across the Phase 1 deployment. The Kafka cluster handles ingestion. Downstream consumers include the supervisory HMI (real-time view), the historian (time-series database with 30-day full-resolution retention and 7-year decimated retention), the alarm engine (rule evaluation against incoming events), and the analytics layer (machine learning for predictive maintenance).

# 7. Alarm rule schema and the rule library

The 3-minute policy described in Section 2 needs an alarm rule library that defines what counts as an alarm in the first place. The schema for these rules and the categories they cover are below.

## Rule schema

```
ALARM RULE STRUCTURE
{
  "rule_id":          short unique identifier (e.g. "chw_supply_high"),
  "scope":            point matcher (e.g. {"asset_type":"chiller","point":"chw_supply_tem
  "condition":        evaluation expression (e.g. "value > 12.0 for 30 seconds"),
  "severity":         enumerated (info | minor | major | critical),
  "ack_required":     boolean,
  "ack_role":         minimum role for acknowledgement (operator | shift_lead | manager),
  "escalate_after":   timer in seconds (default 180),
  "playbook":         reference to runbook in Part D (e.g. "cooling_excursion"),
  "tenant_impact":    expression evaluating whether tenant SLA is at risk,
  "metadata": {
    "owner":          team responsible (e.g. "mechanical"),
    "review_cadence": how often rule should be reviewed (annual | quarterly),
    "last_reviewed":  date of last engineering review,
    "false_positive_rate": running 90-day moving average
  }
}
```

## Rule categories

| Category | Example rules | Approx count | Owner |
|---|---|---|---|
| Mechanical thermal | chw_supply_high, hall_inlet_high, condenser_temp_high | 340 | Mechanical |
| Mechanical flow | pump_no_flow, cdu_low_dp, fan_no_rotation | 180 | Mechanical |
| Electrical voltage | ups_voltage_low, bus_undervoltage, dc_voltage_excursion | 240 | Electrical |
| Electrical current | pdu_overload, breaker_imminent_trip, ground_fault_warning | 210 | Electrical |
| UPS and battery | ups_on_battery, battery_string_imbalance, runtime_short | 120 | Electrical |

| | | | |
|---|---|---|---|
| Genset | genset_fail_to_start, day_tank_low, coolant_temp_high | 180 | Electrical |
| Environmental | hall_humidity_high, room_co_elevated, intake_pm25_high | 140 | Environmental |
| Water and leak | leak_zone_active, sump_high_with_no_pump, fire_water_low_pressure | 90 | Mechanical |
| Life safety | vesda_prealarm, novec_pressure_low, fire_panel_fault | 110 | Life Safety |
| Security | perimeter_breach, mantrap_tailgating, after_hours_access | 200 | Security |
| Tenant SLA | tenant_uptime_at_risk, tenant_power_at_threshold, sla_breach | 80 | Operations |
| Asset health | vibration_high, lubricant_degraded, runtime_above_threshold | 110 | Reliability |

Approximately 2,000 active alarm rules across the Phase 1 deployment. Each rule is reviewed at least annually. Critical and major rules are reviewed quarterly. Rules with false-positive rates above 5 percent in the trailing 90 days are flagged for tuning.

## Rule lifecycle

New rules are proposed by the responsible engineering team, reviewed by Operations, validated against historical data (does the rule fire on the historical incidents it should have fired on, and not fire on the non-incidents it should have stayed quiet on), and approved by the Chief Manager before being activated. Changes to existing rules follow the same process.

Rule disabling is a deliberate operational action requiring senior approval. A rule cannot be disabled to suppress nuisance alarms. The correct response to nuisance alarms is to tune the threshold, not to disable the rule. The audit log records every rule change with the engineer who made it and the approval chain.

# 8. Three incident case studies

The case studies below are anonymised but representative of incidents handled at the HyperNext Phase 1 deployment since commissioning. Each illustrates a different aspect of the BMS operational discipline.

## Case Study 1: Cooling excursion on B4-CH1, May 2026

| Stage | Time | Event |
|---|---|---|
| Detection | T+0 | BMS detects CHW supply temperature on B4-CH1 rising at 0.4 deg C per minute. Rule chw_supply_drift fires at WARN severity. |
| Operator response | T+0:48 | NOC operator acknowledges the alarm. Begins investigation per the cooling_drift playbook. |
| Diagnosis | T+2:14 | Compressor 2 on B4-CH1 has tripped on high discharge pressure. Compressor 1 is carrying full load. |
| Containment | T+3:30 | Operator initiates changeover to standby chilled-water train. B4-CHWP3 ramps up. CHW supply temperature recovery begins. |
| Tenant impact | T+17:00 | Hall inlet temperatures recover to design point. Quantum Bank (tenant in B4-H1) is at 99.97% trailing 30-day uptime, contracted at 99.99%. AT-RISK status set. |
| Notification | T+22:00 | SLA breach notice sent to Quantum Bank via Tenant Alerts. Cause stated as "Cooling excursion, compressor trip on B4-CH1". Service credit position computed. |
| Recovery | T+8 hours | B4-CH1 returned to service after compressor restart and pressure check. Standby train remains primary until next maintenance window. |
| Root cause | T+4 days | RCA completed. Compressor 2 trip caused by intermittent fault on discharge pressure transducer reading falsely high. Transducer replaced. |
| Prevention | T+30 days | Reliability suggestion implemented: pressure transducer redundancy (dual transducers with vote-2-of-2 logic) added to all B4 chillers during next capital cycle. |

What the case study illustrates: the cascade from detection to prevention runs through the BMS at every stage. No step is verbal-only. Every decision is logged. The audit trail is reconstructible months later.

## Case Study 2: Genset failure to start, March 2026

Test condition: scheduled monthly black-start test on B2 genset group. Genset DG07 fails to reach rated speed within the 12-second design window. Genset DG08 (BACKUP in the N+N pair) starts normally.

The BMS detects DG07 fail-to-start at T+13 seconds. CRITICAL alarm raised. Senior approval required for acknowledgement. Shift Lead approves with reason "test condition, not live utility loss". Genset DG07 is taken out of service. The N+1 capacity of the B2 genset bank is maintained by promoting DG09 from standby spare to BACKUP in the pair.

Root cause investigation finds that the starter motor solenoid on DG07 has developed an intermittent fault. Replacement scheduled. The fault would have caused a genuine outage if it had occurred during a real utility-loss event with DG08 simultaneously unavailable. The monthly black-start test caught it; the test discipline justified its own cost.

## Case Study 3: Loss of monitoring on AHU-12, January 2026

The BMS detects that AHU-12 (B1 mechanical room) has stopped reporting on Modbus TCP. Device registry shows OFFLINE status. The mechanical equipment itself continues to operate (verified by visual inspection). The issue is in the monitoring path, not the plant.

This is a MAJOR severity rather than CRITICAL because the redundancy of the plant is not affected; only the visibility is. The investigation finds a failed Modbus-to-IP gateway in the B1 mechanical room. The gateway is replaced. Monitoring restored within 90 minutes.

The case illustrates an important operational pattern: monitoring failures are escalated as MAJOR, not as warnings, because operating blind on a critical asset is a risk equivalent to operating with reduced redundancy. The BMS does not treat "the sensor stopped reporting" as a less serious condition than "the sensor reported a problem".

# 9. References and standards

The HyperNext BMS implementation draws on the following standards and references.

## Protocol standards

- Modbus Application Protocol Specification V1.1b3 (2012) and Modbus Messaging on TCP/IP Implementation Guide V1.0b (2006), Modbus Organization.
- ANSI/ASHRAE Standard 135-2024, BACnet, A Data Communication Protocol for Building Automation and Control Networks.
- IETF RFC 3411 to 3418, the SNMPv3 specifications.
- IEC 62541 (multiple parts), OPC Unified Architecture.

## Data centre standards

- Uptime Institute Tier Standard: Topology, current version. The Tier IV requirements for fault tolerance and concurrent maintainability.
- Uptime Institute Tier Standard: Operational Sustainability. The operating discipline that underpins the BMS design philosophy.
- ANSI/TIA-942-C, Telecommunications Infrastructure Standard for Data Centers.
- ASHRAE TC 9.9, Thermal Guidelines for Data Processing Environments, 5th edition.

## Security standards

- IEC 62443 (multiple parts), Industrial communication networks. Network and system security. The reference for OT security architecture.
- NIST SP 800-82 Rev 3, Guide to Operational Technology Security.
- ISO/IEC 27001:2022 Information security management.

## Adjacent HyperNext publications

- HN-BMS-OPS-MANUAL v3.0 (03 June 2026). The full 46-page Operations Manual referenced throughout this paper. Available under NDA to BMS customers.
- HyperNext Research HN-RP-002, "800VDC Power Architecture for the AI Rack Era". September 2025.
- HyperNext Research HN-RP-006, "Liquid Cooling at 660 kW per Rack". June 2026.

HyperNext Data Centers

## HyperNext Research

We publish engineering and policy papers because the Indian conversation about AI infrastructure needs more substance than marketing material provides. The papers state methodology openly so other operators can run the same analysis on their own facilities. They report findings that may not flatter the HyperNext commercial position. They get review from the engineering team and the communications partners.

Correspondence on methods, figures, and conclusions: **hello@hypernxt.com**. We read every email.

**HN-RP-004** · HyperNext BMS
20 February 2026 · v1.0

www.hypernxt.com/research
hello@hypernxt.com · +91 99784 23333