SENSAKA

sensaka.com

BMC Monitoring Scope

The hardware-layer telemetry your observability stack can't see

What BMC monitoring ingests

Power telemetry

Input/output power per supply, efficiency curves, voltage drift

Thermal data

Per-fan RPM, inlet/outlet temperatures, CPU thermal margin, GPU envelope

Memory health

ECC corrected and uncorrected error counts, DIMM-level diagnostics

System Event Log (SEL)

The canonical hardware event record, vendor-normalized

Firmware inventory

BIOS and BMC versions, security posture, compliance state

Asset metadata

Service tags, serial numbers, models, installed components

Pre-boot & post-failure state

POST errors, IPMI bus health, last-power-down reason

Protocols

Redfish (dominant), IPMI (legacy baseline), iDRAC, iLO, iBMC, XClarity

Why it matters in 2026

1. **AI workload density changed the failure surface**
   Modern AI servers draw 50–100+ kW per rack; early thermal signals — fan curve drift, inlet gradients, thermal margin — only exist at BMC resolution.
2. **EU EED reporting requires IT-equipment efficiency**
   Annual reporting on server utilization and IT efficiency for data centers above 500 kW. Facility-only DCIM misses exactly this IT layer.
3. **Observability platforms hit a structural ceiling**
   OS-agent tools stop reporting when the OS hangs, the host is unreachable, or the failure happens before boot. BMC monitoring continues regardless.

What it does NOT include

BMC monitoring does not replace application observability: no traces, OS process metrics, container telemetry, or Kubernetes state. It is the layer beneath those platforms, not a substitute — each sees what the other cannot. And a BMS (Building Management System) watches UPS, CRAC units, and room sensors; it does not see inside servers.

Public knowledge from Sensaka · free to share with attribution
Source: sensaka.com/glossary/what-is-bmc-monitoring