HBK
HOTTINGER BRÜEL & KJÆR

WHITE PAPER

Ensuring Five-Nines Uptime in the Age of AI

How reliability-centric design and operations are becoming a critical data center advantage.

Adi Dhora
Principal Reliability Consultant, HBK

The rapid growth of AI training workloads is fundamentally reshaping the reliability requirements of modern data centers. While traditional architectures were designed around steady-state, predictable workloads, AI introduces extreme power variability, unprecedented compute density, and training processes that can span weeks or months without interruption. At the same time, the financial and operational consequences of downtime are escalating, with service level agreements increasingly tied to ultra-high availability.

Achieving 99.999% availability ("five nines") is no longer only an operational target - it is a financial and competitive necessity. This whitepaper examines how AI workloads are changing the physics of data center uptime and explores how leading owners and operators are responding. Drawing on HBK's perspective and industry engagement, it outlines how Reliability, Availability, and Maintainability (RAM) digital twins and system-level monitoring are emerging as essential tools for sustaining high uptime in AI-driven data center environments.

The Rising Stakes of Data Center Uptime

AI training workloads have made achieving high-tier uptime significantly more difficult. As a result, attention has returned to the reliability of physical assets, including power infrastructure, cooling systems, compute hardware, and monitoring and control systems.

For colocation providers and hyperscalers such as AWS, Google, and Microsoft, uptime is now a critical priority due to the direct financial exposure tied to downtime. Many service level agreements impose penalties ranging from 25% to 200% of Monthly Base Rent (MBR) if availability falls below 99.999%. Under these conditions, even a single severe incident - such as operating outside a pre-agreed temperature range for several hours or a 26-second power outage - can eliminate a significant portion of rental income.

Beyond SLA penalties, the broader impact of outages continues to grow. Industry surveys show that one in five data center outages costs more than $1 million, accounting for direct costs, lost opportunity, and reputational damage from outage through full recovery.

Reliability as a Winning Data Center Strategy

Reliability is increasingly recognised as a key driver of long-term data center value. Insights from KKR highlight that returns will accrue disproportionately to the most disciplined owners and operators. Those able to overcome structural bottlenecks - such as power availability, land access, grid connectivity, and permitting - while combining strong commercial contracts with high operational availability are positioned to generate higher-quality cash flows and achieve higher asset valuations.

In this environment, uptime is not simply an operational measure. It is a strategic differentiator with direct financial implications.

How AI Is Changing the Physics of Uptime and Reliability

Historically, data centers have been highly successful at managing steady-state, predictable workloads and high-availability storage. Resilience was largely achieved through a software-led response to downtime, where workloads could be shifted to another data center if one site failed. This approach reduced the emphasis on extensive physical fault tolerance and capital investment in redundant systems within individual facilities.

AI training workloads are fundamentally changing this model.

Key characteristics include:

- Multi-megawatt power swings occurring within seconds
- Training sessions that can take weeks or months to complete, where a single power flicker or cooling failure can force the entire process to restart, potentially wasting millions of dollars in compute time
- Extreme compute density, meaning that if cooling fails, equipment can reach critical thermal limits in minutes rather than hours

These factors significantly reduce tolerance for disruption and place renewed emphasis on the reliability of physical infrastructure.

Escalating Compute Demand and Capital Exposure

This shift is only accelerating. Advanced AI training compute demands are now doubling as frequently as every 2 to 3.4 months in some high-performance clusters. At the same time, annualised capital expenditures by large technology firms - who together account for 44% of the total data center market - have doubled compared to 2022 levels.

As compute intensity and capital investment increase, the cost and impact of downtime continue to rise. Sustaining continuous operation under these conditions requires deeper insight into how infrastructure performance, maintenance activities, and operational decisions interact at the system level.

How the Industry Is Responding to the Five-Nines Challenge

At HBK, we are gaining a front-row view of how data center owners and operators are addressing the challenge of delivering ultra-high uptime for AI data centers. Leading organizations are leveraging *HBK's ReliaSoft software to develop Reliability, Availability, and Maintainability* (RAM) digital twins, supported by near real-time monitoring routines for predictive maintenance.

Design and Validation: Embedding Reliability

In the design and validation phase, RAM digital twins are created to simulate and forecast the operational continuity of data center operations. These models encompass the primary causes of outages, including power systems, continuous cooling, compute hardware, and network connectivity.

The analysis methodologies employed in *ReliaSoft's BlockSim software* – such as reliability block diagrams, fault trees, and event trees - have long been used in aerospace and nuclear facility design. They are now becoming increasingly mainstream tools for data center uptime studies.

In-Field Operations: System-Level Risk Awareness

During in-field operations, the *RAM digital twin* is connected to data center information management systems, including CMMS, EAM, BIM, and SCADA platforms. This integration allows the model to incorporate predictive maintenance alerts, findings from preventive maintenance activities, and the current operating status of critical equipment such as power and cooling systems.

Rather than responding to every alert at the component level, operators gain a system-level view of outage risk. This system-level operational approach has proven more effective in helping data center teams prioritise the most important actions required to keep racks, rows, aisles, and entire facilities running continuously.

AI training workloads are redefining the reliability requirements of modern data centers. Multi-megawatt power variability, extreme thermal density, extended training cycles, and rising capital investment mean that even brief disruptions can have disproportionate consequences. At the same time, financial penalties and valuation considerations are intensifying the need for consistently high availability.

In response, leading data center owners and operators are moving beyond reactive maintenance and software-only resilience approaches. By adopting RAM digital twins and system-level monitoring, reliability is being embedded into both the design and operation of AI-ready data centers. As the industry continues to evolve in the age of AI, these practices are becoming central to achieving and sustaining five-nines availability.

Sources

- Data Center Dynamics – Understanding SLA Risk: A Guide for Data Center Investors and Capital Providers https://www.datacenterdynamics.com/en/opinions/understanding-sla-risk-a-guide-for-data-center-investors-and-capital-providers/
- Uptime Institute – 2024 Global Data Center Survey Report https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.GlobalDataCenterSurvey.Report.pdf
- Ingenious – Data Center Tiers Explained https://www.ingenious.build/blog-posts/data-center-tiers-explained
- KKR – AI Infrastructure Insights https://www.kkr.com/content/kkr/sites/global/en/insights/ai-infrastructure.html
- Epoch AI – Compute Trends https://epoch.ai/blog/compute-trends