

AI and the New Edge of Resilience

AI is rewriting the rules of networking, demanding resilience at the edge and beyond. In this new era, out-of-band management is no longer optional—it's the backbone of secure, always-on connectivity.

Report Compiled By: Bill Kleyman

Contributing Editor (Data Center Frontier, Data Center Knowledge, InformationWeek)

Executive Chair, Data Center of Programs, Informa

CEO, Apolo.us



Table of Contents

03	Executive Summary	
04	Introduction	
	Resilience in the Age of AI Infrastructure	
05	Section 1: The Urgency of Resilience in an AI Era	
06	1.1 Network Resilience (Re)Defined	
07	1.2 The Cost of Downtime	
08	1.3 AI Requirements at the Edge	
09	1.4 Security Overlay	
10	1.5 Out-of-Band Management's Role	
	Toward the AI Factory Network	
11	Section 2: The Evolution of Networking:	
	From Legacy to AI Factories at the Edge	
12	2.1 From Static to Modern Infrastructure	
13	2.2 The Shift to Modern Infrastructure	
14	2.3 Challenges in the Modern Network	
15	2.4 Opportunities in the AI Era	
16	Section 3: Out-of-Band Management	
	as the Backbone of AI Resilience	
17	3.1 From Tool to Strategy	
18	3.2 Resilience Through Intelligent Autonomy	
19	3.3 Security at Every Layer	
20	3.4 From Reactive to Proactive: Closing the Resilience Gap	
	Opengear's Differentiator	
21	Section 4: The Partners that Modernize	
	AI Resilience at the Edge	
22	4.1 Keeping AI Operations Afloat: AIDA Cruises	
23	4.2 Scaling Global Network Resilience: Hyve Managed Hosting	
24	Final Thoughts and a Look into the Future	

Executive Summary

Artificial intelligence is driving unprecedented demands on digital infrastructure, and traffic throughput has [doubled](#) since 2020 and is expected to hit [1,088 exabytes per month by 2033](#), as data-intensive AI models scale across cloud and edge environments. In this new era, the cost of downtime is staggering—estimated at an [average](#) of **\$9,000 per minute** for enterprises, and far higher when GPU clusters and AI inference workloads are disrupted.

Traditional network architectures, built for centralized and predictable traffic, are ill-suited for the distributed, latency-sensitive requirements of AI at the edge. Out-of-band management is now essential, providing secure, independent access and rapid recovery when the production network fails.

This paper explores the evolution of networking, the escalating risks of outages in an AI-first world, and how proactive lifecycle resilience strategies can protect organizations from both operational and financial losses.



Introduction

Welcome to the age of artificial intelligence, where every enterprise has become an AI-driven enterprise. Over the past decade, networks evolved to support cloud adoption, mobile connectivity, and digital transformation, but the past two years have accelerated demands beyond anything we've seen before. AI is fueling a surge in data traffic, with utilization climbing nearly 30% year-over-year and projections pointing to [overall growth of 5x to 9x by 2033](#).

At the heart of this growth is the modern data center, but the lifeblood is the network that connects AI workloads, edge devices, and critical applications. This interconnectivity has fundamentally shifted: It's no longer just about faster switches or bigger pipes, but about resilience, security, and out-of-band intelligence to ensure uptime in an always-on AI world. Before diving deeper, it's important to understand how this evolution is reshaping network architecture—and why resilience has become the defining metric of success in the AI era.

Resilience in the Age of AI Infrastructure

According to recent surveys, 84% of CIOs and CSOs have reported an increase in network outages over the past two years, with many organizations losing between \$1 million and \$5 million annually due to downtime.

At the same time, the way organizations think about cloud, colocation, and private infrastructure is shifting. Cloud repatriation is accelerating as enterprises rebalance workloads between hyperscale environments, colocation providers, and private data centers—driven not only by cost but also by the performance and resilience demands of AI workloads.

In fact, nearly **80% of IT leaders expect significant capacity increases due to AI**, and more than **60% are already deploying AI-ready solutions in their facilities**.

The common thread running through all of this is connectivity. AI inference at the edge, distributed workloads, and hybrid cloud models require networks that are not just fast, but resilient and secure.

Interconnectivity modernization is no longer about raw speed alone—it's about designing physical and logical infrastructure that can withstand outages, protect against cyber threats, and enable AI-driven operations without interruption.

In this paper, we'll explore:

- Why resilience is the **defining requirement** of AI-ready networks
- How the evolution of network architectures is **reshaping** data centers
- The role of out-of-band management in **securing and stabilizing** critical infrastructure
- Real-world use cases of resilient network strategies **powering AI adoption**

Ultimately, this report will connect the dots between digital infrastructure, AI, and interconnectivity modernization—helping organizations rethink how they design for resilience in an era where downtime is no longer an option.



SECTION 1

The Urgency of Resilience in an AI Era

Our connected world isn't experiencing a transformation any longer. With the advent of new AI capabilities, we're going through a renaissance. Artificial intelligence has become the single greatest stress test of modern infrastructure.

Every new AI model, every edge inference node, and every terabyte of training data increases the strain on networks that were never designed for this level of demand. In this environment, resilience isn't just a technical objective—it's a business imperative. The enterprises leading in AI today are those that can keep their networks operational under pressure, recover instantly from disruption, and secure their data at the edge. This section examines why network resilience has become the new measure of readiness in the AI era.

To kick things off, we need to clearly redefine network resilience in this modern age of AI and edge ecosystem.

SECTION 1.1

Network Resilience (Re)Defined

Artificial intelligence has become the defining workload of this decade, transforming how data centers, networks, and enterprises operate. But as AI grows, so too does the pressure on the network—the invisible backbone of every digital process.

Modern network resilience can be defined as:

The ability to provide and maintain an acceptable level of service in the face of faults and challenges to normal operation.

In today's AI-driven ecosystem, resilience means more than simple redundancy. It is the ability of a network to anticipate, withstand, and recover from disruptions—whether caused by cyber incidents, configuration errors, or physical outages—without losing access to critical systems.

According to Opendgear's latest Global Report on Network Management, 84% of CIOs and CSOs report a rise in network outages over the past two years.

Complexity has exploded as AI and edge computing introduce new layers of distributed architecture, automated decisioning, and latency-sensitive workloads. The network is no longer a passive transport layer—it's a dynamic, intelligent substrate that must operate

independently and recover instantly when failures occur. In essence, AI has made resilience the new uptime.

With that in mind, we need to also understand that persistent connectivity and reliance on our networking ecosystem has materially changed the cost of outages. Let's explore what that means in today's world.



↑ 84%

of CIOs and CSOs reported a **rise in outages** over the past two years

SECTION 1.2

The Cost of Downtime

Because of the level of connectivity we require today, the economic and reputational impact of network downtime has never been greater. A recent report titled Measuring the True Cost of Network Outages found that 31% of global enterprises lost more than \$1 million in the past 12 months due to outages, with 17% reporting losses exceeding \$6 million.

Nearly **one in five** organizations experienced **seven or more** outages lasting **30 minutes** or longer within a single year, while **39%** said that outages take **more than a full working day** to resolve.

These figures are even more critical in the AI era. When a high-performance compute cluster sits idle, GPU utilization plummets—resulting in direct financial waste. For organizations running AI training workloads, every hour of downtime can mean tens of thousands in lost compute value. For inference environments at the edge—like autonomous vehicles, hospitals, or financial exchanges—an outage can translate to compromised safety, compliance failures, or brand damage in seconds.

Yet **59%** of organizations lack preventive maintenance programs to minimize downtime, and **41%** cite travel time to remote sites as a top barrier to restoration.

These operational inefficiencies highlight a structural weakness: Many enterprises still rely on traditional, in-band network management models that depend on the very infrastructure that fails during an outage.

The business case is clear. When even brief disruptions can cost millions, **resilience becomes a profit strategy**, not a cost center. Here's the other critical concept to consider: Modern resilience is also evolving our network requirements at the edge. This is compounded by the complexity and importance of modern AI systems, and even more so now that these environments have spread to the edge.





SECTION 1.3

AI Requirements at the Edge

AI workloads are increasingly decentralized. Inference, real-time analytics, and generative models demand processing closer to where data originates—whether in a smart factory, retail store, or healthcare facility.

Global data shows that **28%** of CIOs and CSOs expect **edge computing** to have the single greatest impact on network management within the next five years.

This shift is reshaping how we design connectivity. Low-latency inference requires deterministic networking, consistent throughput, and near-instant failover. Many AI workloads at the edge are highly sensitive to network disruption. When thousands of sensors and edge nodes feed real-time inference pipelines, even minor latency spikes cascade into system-wide slowdowns or failures.

To meet these demands, organizations are adopting hybrid and multicloud architectures, layering local processing nodes on top of centralized compute. The challenge is visibility—operations teams must manage geographically dispersed assets, often with limited on-site expertise. *Smart Out of Band™* (OOB) management provides the necessary control plane for these environments, enabling remote configuration, diagnostics, and recovery even when the production network is unreachable.

But before we dive further into OOB management solutions, which we will cover in later sections, we need to explore a critical component of resilience, AI networks, and modern infrastructure: **security**.

SECTION 1.4

Security Overlay

Resilience and security are now inseparable. Distributed AI systems have dramatically widened the attack surface: Every remote edge device, every API connection, every inference gateway is a potential entry point. Opengear [research](#) found that nearly a quarter (24%) of CIOs and CSOs cite data breaches and API security as top challenges in managing cloud and distributed data environments. Moreover, cyberattacks were among the leading causes of outages for 21% of respondents, underscoring the interplay between downtime and cybersecurity risk.

AI helps defenders detect threats faster—69% of network engineers believe AI integration will improve their ability to respond to cybersecurity incidents—but **it also creates new dependencies that must remain operational under stress.**

When attackers strike during an outage or exploit misconfigured automation pipelines, a resilient management layer is often the only means to isolate and remediate issues securely.

Smart OOB solutions reinforce this critical layer of defense. Because they operate on an independent management plane, they enable secure access and control even when the production network is compromised or unavailable.



69%

of network engineers said the integration of AI will **improve their ability to respond to cybersecurity incidents**

They remain operational and secure during:

DDoS events: Maintain management connectivity when primary links are overwhelmed

Ransomware lockouts: Allow administrators to bypass encrypted systems and begin recovery immediately

WAN or core network failures: Ensure that remediation efforts continue uninterrupted

In a world where Mean Time to Detect (MTTD) and Mean Time to Repair (MTTR) define security success, out-of-band management has become both a resilience mechanism and a cybersecurity control.

With this in mind, let's wrap up this section with a clear understanding of OOB management's role in this new type of networking architecture.

SECTION 1.5

Out-of-Band Management's Role

Traditional network management depends on in-band connectivity—using the production network itself for monitoring and configuration.

But here's the scary part. When that network fails, visibility and control disappear exactly when they're needed most. Out-of-band management provides an independent, always-available access path to critical infrastructure, forming the cornerstone of true resilience.

Losing visibility to critical networking components is a stress-inducing event that no engineer would like to experience. It's a big reason why, according to [research](#), 25% of network engineers plan to increase OOB investments within five years, ranking it alongside 5G and IoT as a top modernization priority.

Smart OOB technologies not only enable remote troubleshooting but also integrate automation and predictive intelligence—key for maintaining uptime in AI environments where outages can cascade across hundreds of nodes.

For instance, Failover to Cellular™ capabilities maintain connectivity when the primary WAN is down, ensuring that AI workloads or edge devices remain online while repairs are underway. Combined with NetOps automation, OOB can automatically detect, diagnose, and remediate issues without human intervention—reducing restoration times from hours to minutes.

This shift aligns with the broader AI-driven transformation of network operations. As 94% of CIOs and CSOs have begun implementing AI for network management, resilience is becoming embedded within the operational fabric.

OOB management serves as the bridge between AI automation and physical network control—providing the **secure, autonomous pathway** that keeps AI systems stable even when the unexpected occurs.

As AI reshapes how data moves, processes, and learns, the very architecture of networking must evolve. Traditional, centralized designs can no longer deliver the agility, visibility, and control that AI workloads demand. The next generation of infrastructure—what many now call the AI factory—requires a fundamental redesign of the network itself. In the following section, we'll explore how this transformation is unfolding, tracing the shift from legacy systems to intelligent, adaptive networks built to power the AI-driven enterprise.

Toward the AI Factory Network

The “AI factory” is more than a metaphor; it's an operational model in which data flows continuously through pipelines of training, inference, and feedback. Like any production line, it depends on precision, visibility, and unbroken continuity. The network becomes the assembly line of intelligence—powering everything from automated design and cybersecurity analytics to autonomous vehicles and real-time retail personalization.

In this new paradigm, downtime is equivalent to production stoppage, and packet loss is a defect in output. The factory floor stretches from hyperscale data centers to 5G towers, micro data centers, and IoT gateways. Each element must communicate flawlessly with the others. That is why the future of networking lies not merely in more bandwidth, but in architectural intelligence—systems that anticipate demand, reroute around failures, and rebuild themselves in real time.



SECTION 2

The Evolution of Networking: From Legacy to AI Factories at the Edge

Networks were once simple highways for information. Today, they are the neural pathways of artificial intelligence—carrying not just data, but decisions, insights, and actions that define entire industries. The legacy architectures of the past can't keep pace with this new rhythm. With AI shifting how and where workloads live, networks must grow from rigid frameworks into agile, self-healing systems capable of supporting intelligence at the edge.

SECTION 2.1

From Static Pipes to Adaptive Infrastructures

For decades, networks were engineered for predictability. And for those decades, that's all we really needed from our networks.

Traditional architectures revolved around centralized data centers, structured traffic flows, and human-driven management cycles. Capacity planning was methodical; routing decisions were deliberate. These designs worked well for conventional IT workloads—email servers, transactional systems, and web applications—but they were never meant to accommodate the unpredictable velocity of AI and edge computing. With GPT-based technologies, new LLMs, and inference at the edge, we're far beyond traditional web apps and database servers.

As Opengear's research notes, the explosion of distributed, data-intensive workloads has exposed the limits of traditional networking. Virtualization, cloud migration, and now AI have layered new dependencies on top of rigid legacy frameworks, making visibility, control, and uptime exponentially harder to maintain. The shift from fixed connections to dynamic, software-defined and automated topologies represents not just a technical evolution, but a philosophical one: Networks are no longer passive conduits—they're living ecosystems that must learn, adapt, and self-heal.

Let's explore that shift a little bit more.



SECTION 2.2

The Shift to Modern Infrastructure

Today's data center is not a single place—it's a distributed organism spanning colocation sites, cloud regions, and thousands of edge locations.

Each node plays a role in what industry leaders now call the **AI factory**: An interconnected fabric of compute, storage, and networking that trains, tunes, and serves machine-learning models in real time.

This new topology relies on pervasive interconnectivity and orchestration. AI workloads generate massive east-west traffic inside data centers and continuous north-south communication with users and devices at the edge. Meanwhile, hybrid and multicloud strategies blur the boundaries between private and public environments.

Here's an important point to consider: **Network management has become an AI problem in itself**—requiring autonomous tools that predict, prioritize, and remediate issues faster than human engineers can react.

At the same time, operational environments are becoming more heterogeneous. Enterprises now run workloads across SD-WANs, 5G links, and cloud interconnects while managing legacy MPLS and fiber circuits. The result is an explosion of complexity—yet the expectation from business leaders remains simple: zero downtime. That's a tough expectation without good systems in place.

Meeting that expectation requires infrastructure that can detect degradation before it becomes failure, isolate faults automatically, and maintain service continuity even when the primary control plane collapses.

Before we go too much further, it'll be important to review the real-world challenges facing modern networks. There is good news with the list below. With challenges come real opportunities to do something great.



Challenges in the Modern Network

Despite advances in software-defined networking and cloud orchestration, three persistent challenges define this new era of connectivity:

01 Rising Density and Data Flow Complexity

AI clusters push bandwidth and power demands beyond conventional limits. A single rack of GPUs can generate 50–100 kW of heat and terabits of data per second. Moving, storing, and securing that data creates immense strain on both physical and logical layers. As traffic scales, so does the potential for congestion, configuration drift, and performance loss across distributed fabrics.

02 Expanding Attack Surfaces at the Edge

The same decentralization that enables AI agility also opens thousands of new entry points for attackers. Each sensor, gateway, or remote server becomes a potential vulnerability.

Resilience studies highlight how cyber incidents and human error remain two of the leading causes of outages, with attackers increasingly exploiting network downtime windows to infiltrate systems.

03 Human Error and Operational Overload

Even as automation rises, people remain a critical weak point. Misconfigurations, missed updates, and reactive maintenance continue to trigger cascading outages. According to field research, organizations with largely manual network operations suffer higher outage frequency and longer recovery times than those that implement automated or out-of-band remediation frameworks.

Each of these factors compounds the urgency for intelligent network design. The modern enterprise can no longer depend solely on centralized NOCs or in-band tools that vanish when connectivity falters. Visibility, resilience, and security must extend all the way to the edge.



SECTION 2.4

Opportunities in the AI Era

As mentioned earlier, while the above challenges can seem daunting, they also create a new set of opportunities for innovation and differentiation.

01 Building Self-Healing Networks

AI itself is becoming the key to network reliability. Predictive algorithms can now analyze telemetry, identify anomalies, and trigger automated recovery before end users notice a disruption. *Smart* OOB systems complement these capabilities by maintaining an independent control path—allowing remote diagnosis, reboot, or configuration even when the production network fails. Together, these create the foundation for self-healing architectures that drastically reduce MTTR.

02 Leveraging Predictive Analytics and Automation

Opengear's analysis of AI investment trends reveals that 94% of IT leaders are now implementing some form of AI-driven network management, signaling a decisive move toward predictive operations.

Predictive models help organizations understand capacity thresholds, forecast outages, and optimize maintenance windows, converting reactive firefighting into proactive resilience. The integration of NetOps automation tools extends this intelligence by removing repetitive tasks and configuration drift—the same causes of downtime highlighted in earlier research.

03 Creating Hybrid Models that Bridge Legacy and AI-Native Systems

Transformation rarely happens overnight. Most enterprises will operate in hybrid states for years—supporting both traditional applications and AI workloads across shared infrastructure. The challenge lies in synchronizing the two worlds: Ensuring that older network assets can coexist with AI-driven orchestration and monitoring platforms. As you think about all of this, consider an important point: True modernization doesn't mean replacement—it means intelligent integration.

Organizations that can merge legacy reliability with AI agility will lead the next phase of digital evolution.

In this context, OOB management once again proves foundational. It creates a universal control layer capable of spanning legacy routers, next-gen switches, and edge nodes—providing consistent access, telemetry, and automation regardless of topology. It's the connective tissue that allows both worlds to operate as one coherent, resilient system.

As organizations reimagine their networks as intelligent, adaptive systems, one principle becomes clear:

Resilience must be engineered, not assumed.

The next frontier is making that resilience tangible—embedding it into physical and digital design through technologies that operate when everything else stops.

In the next section, we'll explore how out-of-band management and Opengear's innovations form the backbone of AI-ready resilience, enabling enterprises to build networks that think, heal, and recover faster than ever before.



SECTION 3

Out-of-Band Management as the Backbone of AI Resilience

Let's pop the hood a little bit and dive further into OOB management. Artificial intelligence has redefined what it means for a network to be resilient. With thousands of distributed devices operating at the edge, networks must now stay online and secure even when their primary pathways fail. This is where out-of-band management becomes indispensable—not as a backup tool, but as a core architectural strategy for modern infrastructure.

SECTION 3.1


From Tool to Strategy

For many organizations, OOB management was once a safety net—activated only when the production network failed. Today, it has matured into a strategic layer that underpins the entire network lifecycle. Modern systems like Opengear’s *Smart OOB* use an independent management plane to maintain connectivity, control, and recovery capabilities regardless of the primary network’s status.

This independence is critical in AI-driven environments, where traffic spikes, distributed inference, and machine-to-machine communication can create sudden surges in load and complexity. Unlike in-band tools that fail when the network goes dark, *Smart OOB* provides persistent, out-of-band access via dedicated Ethernet, serial, and cellular interfaces.

In Opengear’s architecture, *Smart OOB* integrates seamlessly with automation platforms to allow zero-touch provisioning, remote image recovery, and configuration rollback—capabilities that make it not just a monitoring tool, but an autonomous operational backbone for distributed systems.

Let’s focus on autonomous operations a bit more.



SECTION 3.2

Resilience Through Intelligent Autonomy

Modern OOB frameworks combine Failover to Cellular, *Smart* OOB, and NetOps automation to enable true operational autonomy. This multilayered approach allows for continuous uptime even when physical or logical links fail. For example, if a primary WAN connection drops, Opengear's embedded LTE module automatically provides secure cellular failover, ensuring that AI workloads at the edge remain connected and synchronized.

Beyond failover, NetOps automation extends resilience into the predictive domain. Automated scripts can detect network anomalies, trigger diagnostic tests, reboot devices, or deploy configuration fixes without human intervention. This automation aligns perfectly with AI environments, where self-healing and continuous learning are core design philosophies. The integration of these systems allows enterprises to operate AI factories that don't stall when connectivity falters—they self-correct and continue processing data.

And now, with the introduction of Smart Management Fabric (SMF), Opengear extends OOB capabilities into dynamic IP-routed management networks. Using industry-standard protocols such as Open Shortest Path First (OSPF) and secure VPN tunnels, SMF enables scalable machine-to-machine and user-to-machine connectivity across every Opengear-enrolled device.

This means that automation platforms—whether FortiManager, Panorama, or VMware vCenter—can now communicate directly with remote devices over a secure OOB channel, even when the production network is unavailable.

At a technical level, this architecture reduces MTTR dramatically. Engineers can remotely access equipment using console servers—even in isolated edge environments—to reimage routers, restore firmware, or execute diagnostic scripts. Combined with automation frameworks such as Ansible and RESTful APIs, OOB management becomes the control plane for a self-sustaining network ecosystem.

Earlier in the paper, we touched on security. Now, let's go a little bit deeper into security architecture, OOB management, and resilience at the edge.



SECTION 3.3

Security at Every Layer

As networks become more distributed, the attack surface grows exponentially. The same connectivity that empowers AI and edge computing also exposes critical systems to new threats. OOB's architectural independence creates a dedicated, secure management plane that isolates administrative functions from production traffic, reducing exposure to cyberattacks.

Opengear strengthens this model with a multilayered security design that includes:

Encrypted access protocols such as **SSH, TLS, and PKI authentication**, ensuring all management sessions remain private and tamper-proof

Role-Based Access Control (RBAC) and **centralized credential management** to enforce least-privilege access and prevent lateral movement

Audit-ready logging and **compliance monitoring**, providing a full record of access attempts and configuration changes for forensic visibility



This approach ensures that even during an outage or cyber incident, administrators can securely log into devices, contain damage, and restore service without exposing critical assets to public networks.

In AI and edge environments—where automation systems act faster than human operators—this visibility and segmentation are mission-critical. OOB ensures that human operators always retain an independent, secure

command channel, even when the production network or automation framework is compromised.

It's important to quickly touch on one final point: creating resilient network architecture that goes from reactive to proactive management systems.

SECTION 3.4

From Reactive to Proactive: Closing the Resilience Gap

Despite its growing importance, research shows that nearly 80% of enterprises still wait until a major incident before deploying OOB solutions, exposing themselves to unnecessary risk. This reactive posture contrasts sharply with the demands of AI operations, where real-time adaptability and continuous uptime are paramount.

Organizations that treat OOB as part of their core design strategy, rather than an afterthought, achieve faster recoveries, lower operational costs, and greater confidence in their network’s ability to withstand disruption.


Opengear’s Differentiator

Opengear has advanced this model by integrating hardware, software, and lifecycle support into a single, cohesive platform. The company’s latest business model bundles network automation, secure management tools, and technical support into one partnership that scales from data center to edge.

This approach transforms OOB from a simple redundancy measure into a **foundation for resilient AI networking**, aligning perfectly with how enterprises are redesigning their digital infrastructure.

As networks evolve into self-healing, AI-driven ecosystems, the value of out-of-band management becomes most visible in real-world operations. From hyperscale data centers to remote edge sites, enterprises are already using Opengear’s technology to safeguard uptime, automate recovery, and maintain secure connectivity under pressure.

In the next section, we’ll look at how these capabilities come together in practice—real-world examples of *Smart OOB* and the SMF delivering resilience where it matters most.





SECTION 4

The Partners That Modernize AI Resilience at the Edge

Across industries—from maritime operations to global managed-service networks—organizations are applying Opendgear technologies to sustain AI performance, minimize downtime, and protect business continuity. These case studies reveal how *Smart OOB*, Lighthouse®, and the SMF enable AI-era networks to stay intelligent, adaptive, and always on.

The first example we'll explore is AIDA Cruises, a German cruise line operating a fleet of ships that travel to numerous global destinations.

SECTION 4.1

Keeping AI Operations Afloat: AIDA Cruises

At sea, resilience is not optional—it's mission-critical. On each of AIDA Cruises' 11 ships, over 64,000 networked devices support everything from guest Wi-Fi, onboard commerce, and security cameras to navigation and logistics systems. These are basically mission-critical floating data centers. Any connectivity disruption, even for a few minutes, can ripple across these systems, impacting passenger experience, safety, and operations.

To solve this, AIDA deployed Opendgear's *Smart* OOB solution, establishing a dedicated management network spanning two redundant onboard data centers and multiple distribution rooms per vessel. Each Opendgear device acts as a secure, always-available gateway to critical infrastructure components—core switches, WLAN controllers, firewalls, and servers.

Through Opendgear's Lighthouse software, engineers at AIDA's onshore network operations center can access any onboard system securely, even if the ship's production or satellite network is offline.

The results have been transformational:

90% reduction in MTTR for network incidents

Remote firmware updates for 64,000+ devices performed securely from shore, with full rollback and logging

ROI achieved in under nine months, saving thousands of dollars in technician travel and downtime costs

End-to-end visibility and compliance across a global, mobile fleet

AIDA's deployment demonstrates how *Smart* OOB and Lighthouse deliver autonomous resilience—empowering the network to recover on its own even when isolated and enabling AI-driven monitoring to continue operating across highly mobile, bandwidth-constrained environments. A self-healing network at sea is an extraordinary tool to have.

While AIDA's floating data centers prove that *Smart* OOB management can maintain AI connectivity in the harshest and most remote conditions, the same architecture is transforming scale and agility on land. For Hyve Managed Hosting, the challenge isn't mobility—it's managing thousands of devices across 35 global data centers while maintaining 24/7 availability for enterprise clients running AI workloads.



"Console access via out-of-band management is our insurance that we can reach every critical device at any time. Without it, remote management would simply not be possible at sea."

Diogo Almeida
Head of Network & Connectivity
AIDA Cruises

SECTION 4.2

Scaling Global Network Resilience: Hyve Managed Hosting

Hyve Managed Hosting, a UK-based managed service provider with more than 35 global data centers, delivers secure cloud and colocation services to enterprises that require constant uptime. As its footprint expanded, Hyve faced growing challenges managing distributed infrastructure and supporting AI-driven workloads that demand low latency and uninterrupted connectivity.

To overcome these pressures, Hyve standardized on Opengear's **Smart OOB** architecture, **Lighthouse Automation Edition**, and **SMF**. Together, these solutions form a centralized, independent management network that operates even when the primary network fails. Each device ships preconfigured from the UK, connects via embedded LTE for Failover to Cellular, and automatically registers with Lighthouse to enable secure, zero-touch provisioning.

The approach delivers both technical and operational gains:

Rapid site rollout through **automated configuration templates** and **Ansible/Python integration**

Layer 3 OOB access to physical and virtual assets without relying on serial ports

Centralized visibility and **telemetry** for proactive monitoring and predictive maintenance


The impact has been immediate. Hyve **reduced provisioning times by more than half, cut travel expenses by 70%, and strengthened SLA compliance across continents**. Out-of-band automation now allows engineers to remediate or recover entire sites remotely, without interrupting AI workloads or waiting for on-site intervention.

For Hyve, Opengear transformed reactive troubleshooting into proactive resilience—a foundation that keeps AI environments secure, scalable, and continuously available across its global infrastructure.

"With Opengear, we can manage any device anywhere in the world—no matter what's happening on the primary network. The ability to remotely troubleshoot, reboot, or recover an entire site without dispatching engineers has completely changed the way we operate."

Jon Lucas
Co-Founder
Hyve Managed Hosting





Final Thoughts and a Look into the Future

Artificial intelligence has transformed the network from a passive transport layer into a living system that must anticipate, adapt, and recover in real time. As data moves closer to the edge, the cost of disruption grows—not only in lost revenue, but in lost trust, insight, and opportunity.

The stories of organizations like AIDA Cruises, operating resilient networks across the open sea, and Hyve Managed Hosting, scaling secure operations across continents, prove that intelligent network design can keep even the most complex systems continuously online.

Looking ahead, the future of connectivity will hinge on one defining principle: resilience engineered through intelligence. Networks must evolve beyond redundancy toward autonomy—able to detect, defend, and self-heal faster than human intervention. AI will not only drive the workloads we deliver; it will shape how our networks think, learn, and protect themselves.

The challenge for every enterprise is clear: Evaluate your digital foundation, strengthen your edge, and build systems that can thrive under pressure.

Because in the era of AI, resilience isn't a goal—it's the minimum requirement for what comes next. Here's the most important point to consider as you take this all in:

AI will redefine what networks can do.
Resilience will define which ones survive.



Opengear delivers secure, resilient access and automation to critical IT infrastructure, even when the network is down. Provisioning, orchestration, and remote management of network devices, through innovative software and appliances, enables technical staff to manage their data centers and remote network locations reliably and efficiently. Opengear's business continuity solutions are trusted by global organizations across financial, digital communications, retail, and manufacturing industries.