The Storage Leaders Guide

2026 Trends, Threats & Transformations

Julian Topley
Senior Delivery Manager - Storage & Backup
Lloyds Banking Group

Naresh Kattel
VP Storage Management
AllianceBernstein

Kevin Battle
Senior Storage Administrator
Charter Communications

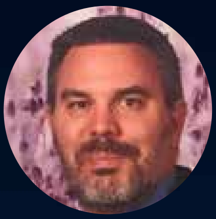

David Brown
Storage Architect
State of Michigan

StorageGuard

Introduction

This **Storage Leaders Guide** brought together four seasoned Storage Execs to explore how enterprises can future-proof storage and backup in 2026.

Across the discussion, one theme came through clearly:
Storage and backup systems are no longer passive infrastructure. They sit at the intersection of security, AI, sustainability, regulation, and business continuity.

The Storage Leaders Guide restructures the December 2025 Virtual Panel into a Q&A-based narrative, preserving the actual viewpoints and quotes from the panelists, organized around four themes: **Industry Trends, Security, Technology & Innovation, and Practical Insights.**

01 Industry Trends

What do you see as the single biggest disruptor in enterprise storage & backup today?

For **David Brown**, Storage Architect at State of Michigan, the biggest disruptor is unambiguous:

> "If you'd asked me this question 20 years ago, I would have said virtualization. Today, the elephant in the room is artificial intelligence. "

David contrasted the early days of SAN:

> "One of the first SANs I ever put in was a whopping 5.5 terabytes. We thought, we'll never outgrow this. Virtualization came along and we were doubling and tripling capacity almost every year. "

Now, AI and machine learning are driving a similar – but even more intense – disruption:

- New workloads (e.g., autonomous vehicles) are discussing yottabytes in capacity planning.
- AI requires instant performance, low latency, and massive parallelism.
- Storage becomes a core dependency for AI's success, not just a downstream consumer.

> "AI and machine learning are a rapid, fundamental force requiring an entirely new approach to enterprise data storage. The higher you predict, the higher your cost – and finance wants to avoid overspending and underdelivering. "

Kevin Battle, Senior Storage Administrator at Charter Communications agreed, zooming in on the hardware and sustainability angle:

> "The new AI workloads require more robust hardware. SSD capacities are increasing and they're getting faster. But that creates a carbon footprint problem – high power demands, high electricity usage. "

He sees a clear direction of travel:

> "The shift I'm seeing is drives becoming more efficient and using less electricity. That's what we need to adopt to."

Takeaways

- AI is the primary disruptor, changing capacity planning, performance expectations, and cost models
- Hardware is evolving quickly, but power and efficiency are now critical constraints
- Predictive capacity for AI workloads is a high-stakes financial decision, not just a technical one

Is sustainability now a real driver in storage decisions, or just a “nice-to-have”?

For **Naresh Kattel**, VP Storage Management at AllianceBernstein, sustainability is absolutely a real driver:

> “It is a driver. When you evaluate storage systems, you have to take footprint, cooling, power, and space into consideration.
> It’s easy to dismiss that, but we take that decision very seriously.„

He emphasized that storage is often the dominant cost center:

> “Storage is, in many environments, about half of the infrastructure cost.
> It’s a big miss for an organization not to analyze that data. It needs to be in the decision every time you go into a refresh or add capacity.„

Naresh also tied sustainability to cloud economics:

> “Even in the cloud, the more efficient a provider is, the better rate they’ll transfer to their clients.
> It’s going to be at the forefront all the time.„

David Brown reinforced that sustainability is no longer just “greenwashing”:

> “Sustainability at one point was considered a green feature, an add-on.
> That’s not the case anymore. It’s intertwined with the financial health and risk management of a modern enterprise.„

He highlighted the operational impact:

> “If an environment can deliver the same capacity and performance while consuming even 50% less power, that’s a massive ongoing OpEx saving.
> Having the ability to do more with less is huge.„

Takeaways

- Sustainability is now tightly linked to cost, risk, and regulatory pressure, not a side benefit
- Storage’s share of power, cooling, and space makes it a prime target for efficiency gains
- Cloud buyers will increasingly favor providers with measurable energy efficiency and cost pass-through

Is virtualization evolving in your storage environments – if so, how?

Julian Topley, Senior Delivery Manager – Storage & Backup at Lloyds Banking Group, described how virtualization has shifted from VM-centric thinking to data-centric and service-centric thinking.

> "I met the founders of VMware years ago when I was at Goldman Sachs and thought it was a novel invention.
> Then multi-core CPUs arrived and suddenly virtualization became a really big thing because we had something that could leverage it. „

Today, he sees virtualization evolving into services defined as code on resilient storage:

> "The real change is how we look at applications.
> What's important is the data, not the virtual machines. „

He argued that:

- VMs and containers should increasingly be treated as stateless infrastructure-as-code.
- The persistent data becomes the true unit of value that must be protected, governed, and made portable.

Julian outlined three guiding principles:

01 Data becomes the unit of value
"Once you do that, you can move compute wherever you like – different hypervisors, containers, cloud – while still meeting resilience and compliance obligations."

02 Protection must be layered
"Snapshots and replication are the frontline for fast recovery. Backup and vaulting are the last bastion for recovery and forensics."

03 Storage must get "cyber-smart"
"Storage has to offer immutability, anomaly detection, and clean recovery points.
Storage needs to be smart."

Takeaways

- Virtualization is shifting from "protect the VM" to "protect and mobilize the data"
- Stateless components belong in code and pipelines; state belongs in resilient, governed storage
- Finer-grained virtual storage services help limit blast radius and improve access control

Naresh Kattel added that virtualization is also enabling granularity and reduced blast radius:

> "Virtualization is allowing the ability to granularize services – even in storage.
> Instead of large buckets serving many clients, you can virtualize and focus services to specific clients. "

This has clear security advantages:

> "That gives you a smaller blast radius from a ransomware perspective and more control over how you permission storage systems.
> The more you virtualize – as long as you're not proliferating virtual devices – the better you can protect your environment. "

Takeaways

- Virtualization is shifting from "protect the VM" to "protect and mobilize the data"
- Stateless components belong in code and pipelines; state belongs in resilient, governed storage
- Finer-grained virtual storage services help limit blast radius and improve access control

What are you seeing around Backup-as-a-Service (BaaS) adoption?

For **Naresh**, growing complexity is driving interest in BaaS:

> "As storage evolves and data grows, backup is becoming more and more challenging.
> You really need a core set of competencies to manage backup properly. "

He emphasized that this is not about primary/DR copies, but tertiary, last-resort protection:

> "I'm not talking about production or DR. I'm talking about tertiary backups – protecting against the disaster you hope you never have to recover from. "

Why BaaS is accelerating:

> "Backup-as-a-Service is going to continue to grow. You don't want to waste time and cycles backing up.
> You want to give it to the experts – people doing this day in, day out. "

However, **Naresh** stressed the importance of due diligence:

01 **Service Levels**
"You need to understand the SLAs these providers give.
How quickly can you get a restore done when you actually need it?"

02 **Independence and locality**
"Are they creating truly independent backup silos from primary data?
Locality is going to be important."

03 **Economies of scale**
"These players have scale. They can probably give you better cost than you'd get on your own because of the size they get from hyperscalers."

His conclusion:

> "We're definitely looking at BaaS as one of our prongs to provide adequate service. It's an exploding field."

Takeaways

- BaaS helps offset skills gaps and operational burden, but does not remove the need for strong internal governance
- SLAs, independence of backup silos, and recovery performance at scale are critical evaluation criteria
- BaaS is best seen as one component of a multi-layered protection strategy, not a full handoff of responsibility

02 Security

How are storage systems adapting to the rise in ransomware attacks?

Julian framed the shift starkly:

> “We’ve invested for years in perimeter security. But the real vulnerability is the destination – the data. That’s where ransomware now goes, and increasingly backups as well. ”

He noted the mindset change:

> “The assumption is moving away from if and more to when.
> Don’t rest on your laurels thinking you’ve built the most secure perimeter. You need to expect something will happen – and be prepared. ”

Modern storage and backup systems must treat data access as a security signal:

> “Platforms are starting to learn behaviors and infer information back to SIEMs.
> Using ML and AI, they can react in real time – not human time – to protect data. ”

Examples include:

- Triggering immutable snapshots when abnormal behavior is detected
- Cutting off suspect clients automatically
- Feeding rich context into the NIST incident response cycle (prepare, detect, analyze, contain, recover, improve)

Julian’s “exam question”:

> “How can I narrow the gap between something looks odd and we’ve got a clean recovery copy to below human reaction time?”

David described how ransomware has elevated storage from “just capacity” to a core pillar of cyber resilience:

> “Enterprise storage used to be viewed as a data repository – JBODs, then SAN, NAS.
> Now, it’s a component of a company’s cyber resilience strategy. ”

Key adaptations he highlighted:

01 Real-time anomaly detection

“If the system sees a sudden spike in encryption, deletions, or file changes – that’s not normal. An alert should trigger immediately.”

02 Automated responses

"Upon detection, the system should freeze activity and trigger a snapshot: 'Something's wrong, I'm protecting my data.' That takes the human out of the initial response."

03 Clean-room recovery

"You can't just restore right back into the same environment. You need a clean room to test and verify you're not reintroducing malware."

He connected this directly to business impact:

> "Yes, these systems are expensive, but what's the cost to your business if you don't have them? If you're down for 24 hours, you could be talking millions or billions depending on your business."

Takeaways

- Storage has become a frontline security control, not just a backend system
- Ransomware strategy must include behavior-based detection, automated protection actions, and clean-room recovery
- Speed matters: the goal is to react and protect faster than attackers can do damage

RANSOMWARE ATTACK
You only have 3 days to submit the payment, or your files will be los
Time Left
02:23:59:06

Are immutability and air-gapped backups becoming table stakes?

David's answer was unequivocal:

> "Both immutable backups and air-gapped backups have become table stakes for a true cyber-resilient strategy.
> They used to be advanced features. Today they're mandatory requirements. „

He explained why traditional backup rules are no longer enough:

> "The old 3-2-1 rule – three copies, two media, one offsite – isn't enough. Attackers can still corrupt those locations. „

His team is moving toward an enhanced rule:

> "We're looking at a 3-2-1-1-0 rule: three copies of data, two different media types, one copy stored offline, one copy immutable or vaulted, and zero errors after verification when you restore. „

On cloud and air gapping:

> "In my opinion, cloud-written backups do not constitute an air gap.
> There's still access to the cloud, and as we've seen, the cloud is not a guaranteed safe harbor. „

He used a clear analogy:

> "Think of an air gap as a drawbridge on a castle.
> You raise it up and there's no way to get in or out until you explicitly lower it again. Cloud doesn't behave like that. „

Takeaways

- Immutability and air-gapped or vaulted copies are now baseline requirements, not premium extras
- Backup strategy should evolve from 3-2-1 to 3-2-1-1-0
- Cloud backups are valuable, but do not replace the need for truly offline or logically isolated copies

What challenges do you foresee with evolving data sovereignty laws?

Naresh pointed out that global enterprises face a patchwork of regulations:

> "Every large enterprise operates in different markets, and every market has a regulator. These regulators are unleashed on you to make sure you understand the rules."

He outlined three practical needs:

01 **Know your inventory**
"You need to understand your data inventory – what you have and where it resides."

02 **Understand regulatory impact**
"You need to understand what those laws are and how they impact each region – localization, data movement, all those nuances."

03 **Partner with legal and compliance**
"You have to partner with the appropriate counsel so you adequately understand the requirements."

His bottom line:

> "Complexity keeps increasing – volumes, ransomware, now regulators. It has become a real challenge."

Takeaways

- Data sovereignty is not just "legal's problem" – storage leaders must understand where data lives and how it moves
- Close partnership with legal and compliance is essential
- Architectures must support regionalization, localization, and controlled mobility of data

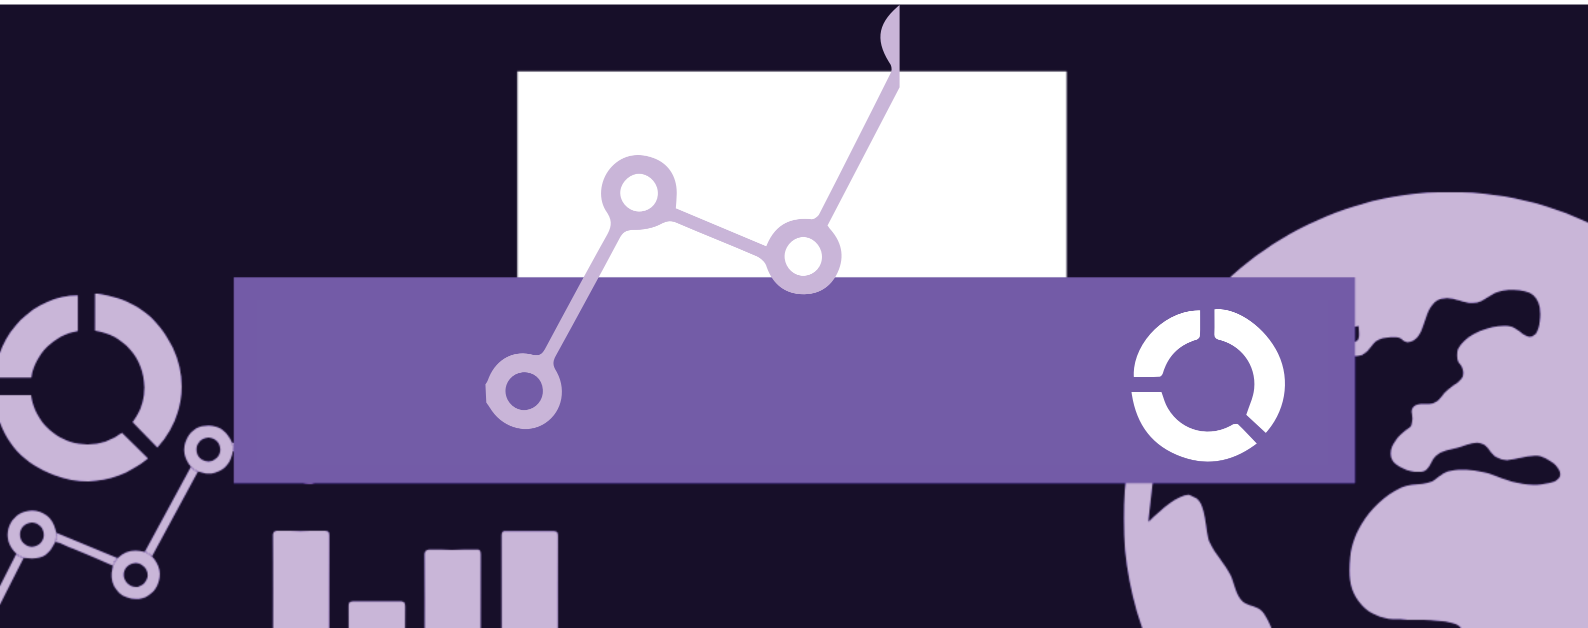

Should InfoSec teams build deeper storage-security expertise?

Kevin strongly believes in deeper cross-team understanding:

> "Absolutely. Right now everything is segmented – separate security teams, separate backup teams, separate storage teams.
> But with new AI-aware software, everyone has to coordinate and be aware. "

He gave a simple but powerful example:

> "Organizations often put decommissions to the side.
> Systems still connected to the network, still in DNS – these are risk points that people overlook. "

Traditional tools aren't designed for this:

> "Many tools we've tested didn't see into the perimeter, into the SAN and storage.
> Rightfully so – they weren't designed for that. "

He highlighted the role of storage-security tooling:

> "That's what I appreciate about solutions like StorageGuard, which I've used for three years.
> It gets inside the perimeter to scan storage and SAN. Those are areas other tools just don't touch. "

Julian agreed on shared responsibility – but warned against blurring it:

> "The security and storage teams must share the problem, but the key is not to blur responsibility. "

He framed it in the following way:

01 **Security team:** owns threat, policy, escalation

02 **Storage & backup team:** own configuration, change, fabrics, protocols

The challenge is context and volume:

> "No single team can hold all the context – CVEs, NIST, ISO, internal standards – in their heads.
> We need ways to bring this knowledge into a central point where teams can work from it. "

Julian explains how 3rd party solutions can help:

> “Tools like StorageGuard really help here – they continuously discover storage and backup systems and add that contextual security layer, mapping configurations and firmware against live vulnerabilities and good practice. „

He also stressed that this cannot be a one-off project:

> “Today, we often treat this as a project – get a PDF full of findings and then ‘get on with it.’ But it’s an ongoing task. Someone has to own threat scanning and interpreting outputs. „

Takeaways

- InfoSec does not need to become expert in FC zoning or array internals – but it must understand storage risk context
- Storage teams must own configuration and remediation, guided by risk prioritization from security
- A sustainable model requires continuous discovery, shared context, and clear ownership, not one-off assessments

03 Technology & Innovation

How is AI changing storage management – is it hype or reality?

David was clear:

> "It's definitely reality.
> But 'AI' is such a broad term that you need to shrink the scope and talk about specific areas."

He sees AI most tangibly in AIOps for storage:

> "We're using AI for IT operations – supporting customer support, security, backup and recovery.
> For storage, it turns systems into active participants in maintaining uptime and data security."

Examples he gave:

01 Predictive maintenance
"Models continuously analyze logs, sensor readings, and performance to predict failures before they impact production."

02 Intelligent tiering & triage
"We used to write scripts to move data between tiers based on age and activity.
AI can now constantly analyze data and predict when it might be needed, staging it preemptively."

He used a healthcare example:

> "If a procedure is scheduled two weeks out, AI can start pulling the right information together and presenting it to the health environment.
> There's no guesswork – it knows what's related to that patient."

His conclusion:

> "AI is taking away time-consuming triage and tiering work.
> It still needs people to validate it, but as machine learning grows, so does the benefit to storage."

Julian framed AI as a data control plane:

> "AI is real in storage when it behaves like a data control plane.
> It can give data a purpose, elevate what drives insight, and help stop powering data obesity."

Concrete roles he sees for AI:

- Locating and classifying data across fragmented estates
- Separating wheat from chaff for retention, backup, and analytics
- Enforcing role-based access and views in lakehouse-style infrastructures
- Operating at real-time scale where humans cannot keep up

> “The way we see storage is going to completely change.
> It will also change how applications behave – we’re going to have to rethink a lot of this. „

Takeaways

- AI in storage is already real, especially in AIOps, anomaly detection, and tiering
- The highest value comes when AI acts as a data control plane, not just a performance optimizer
- Governance and human validation remain essential, but manual triage and scripting are increasingly automatable

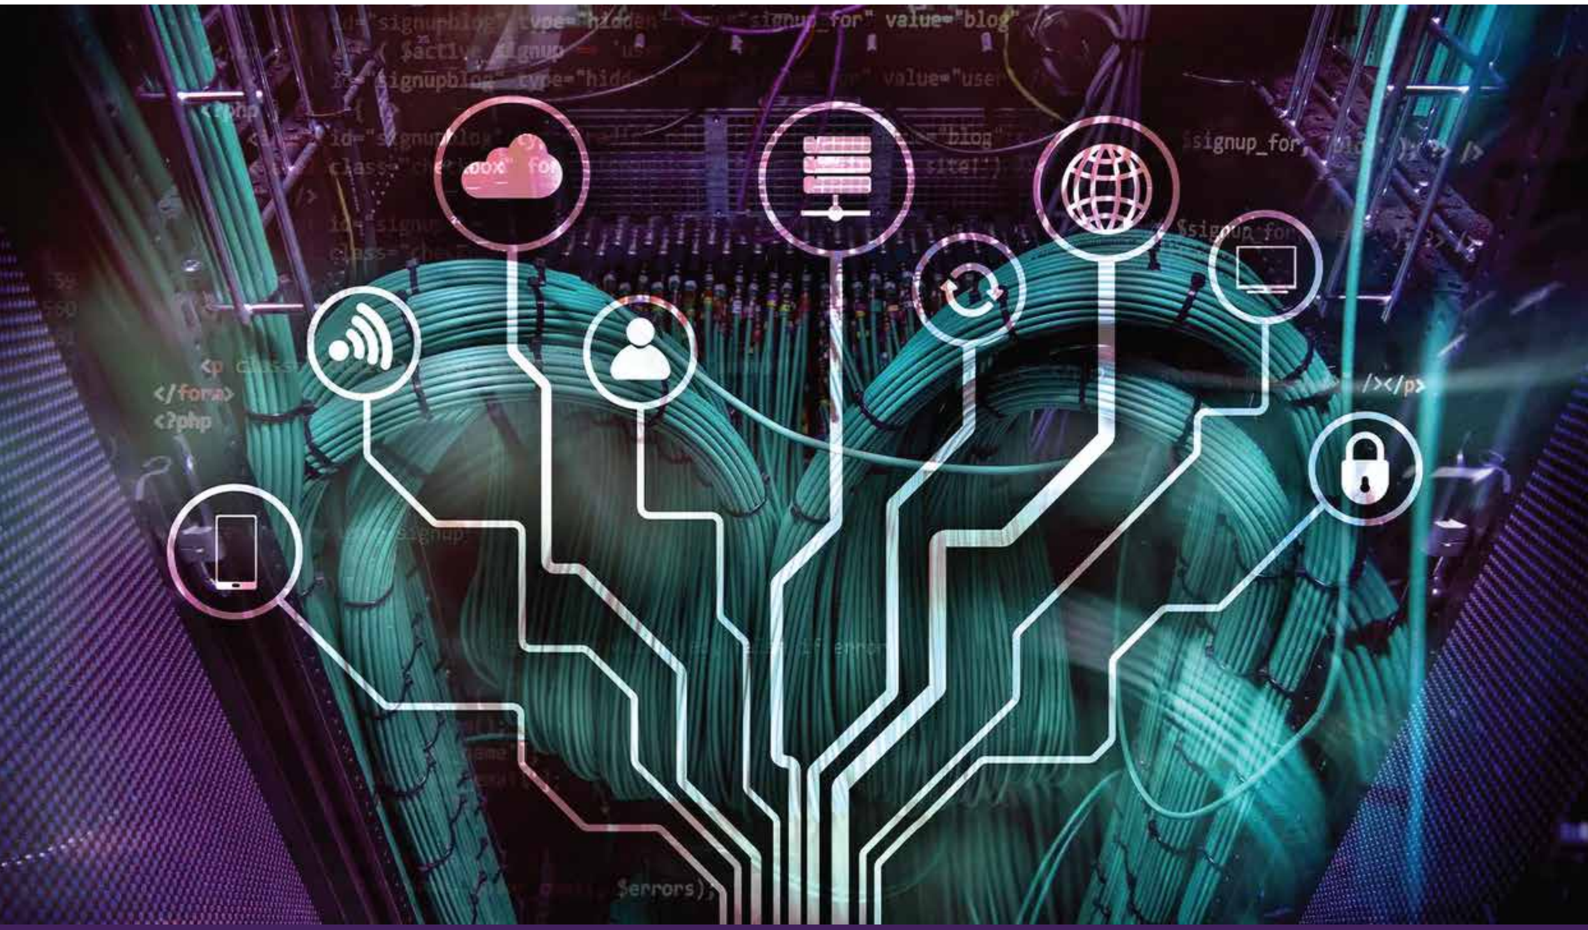

04 Practical Insights

If you had to give one piece of advice to enterprises preparing for the next phase of storage, what would it be?

Naresh Kattel

Use AI and immutability to offset constraints

> "AI is a great place to start. Embrace what it can do for you, because we're going to be constrained budget-wise in every way possible. "

He recommends focusing AI on **metadata and operations:**

> "I'm not talking about proprietary data – I'm talking about metadata of how we manage, provision, and consume storage.
> Some of this you can build in-house and really use to your advantage. "

On security:

> "We touched a lot on ransomware. Immutability of your data – on primary where you can, and in your vault – will go a long way to protecting against the threats we see. "

Kevin Battle

Plan for data growth and hire AI-aware talent

> "Since data is doubling and tripling every two or three years, companies need to really consider the infrastructure – and that's going to incur costs. "

His advice:

> "There has to be a balance.
> Forward-thinking hiring – people that are AI-aware at all levels, including security – will be very helpful. "

David Brown

Look in the mirror as much as the crystal ball

> "Crystal ball questions are always fun but difficult to answer.
> I don't try to predict technology because of the rate of change. "

Instead, he uses hindsight:

> "I keep a mirror on my desk. I can't predict what's to come, but I can see what's occurred.
> The change virtualization brought to storage two decades ago is what I foresee AI already doing now. "

His outlook:

> “Data will continue to grow. If you thought storage grew quickly over the last 10–15 years, I’m anxiously curious to see what it looks like 20 years from now.”

Julian Topley

Be greedy for the right things, and don’t build a Frankenstein

Julian’s closing advice was memorable:

> “You need to be greedy – but for the right things.”

Specifically:

01 Greedy for bandwidth and proximity
“Be greedy for bandwidth and proximity design, so data can reach GPUs and CPUs quickly and hot datasets live close to compute.”

02 Greedy for utilization
“Plan architectures where expensive accelerators and licenses aren’t sitting idle.”

03 Greedy for guardrails
“Be greedy for guardrails – catalogs, lineage, access controls, sovereignty rules.
Use this cycle to build a well-governed data fabric you can safely electrify later.”

He closed with a warning:

> “If we just bolt AI and accelerators on and stitch it all together, history tells us we’ll get something powerful but very hard to control – a kind of Frankenstein.
> Use this cycle to build a data fabric that turns watt-hours and terabytes into useful outcomes, not a powerful uncontrollable beast.”

Conclusion

Across all four panelists, a coherent message emerges:

- AI is real in storage and backup – but must be applied with intention and governance.
- Ransomware and regulatory pressure have turned storage into a core pillar of cyber resilience and risk management.
- Sustainability and cost are inseparable from capacity and performance planning.

Enterprises that invest now in:

- Immutable, air-gapped recovery paths
- AI-assisted operations and data control planes
- Well-governed data fabrics and cross-team operating models

will be best placed to face 2026's disruptions – without creating the very Frankenstein they're trying to avoid.

StorageGuard