Future-Ready Data Center Strategies for Financial Services

IT leaders in financial services are being pulled in two directions. They're under pressure to accelerate AI-driven innovation, while also facing security and regulatory demands and an increasing need to manage risk holistically. As a result, security and innovation compete with one another, while operational costs soar and ROI becomes ever harder to achieve.

It's time for a new approach to the data center, one that delivers performance, security, and compliance without holding back innovation. This requires a consistent hybrid strategy, where security and observability are baked in by design and default, and smarter workload placement enables continuous optimization of performance economics. By standardizing operations across environments, enterprises can reduce complexity and enable AI to scale safely.

A new mandate for financial services organizations

For AI to deliver measurable ROI and become a competitive advantage, it must also become the operating system of the business, not something that's simply bolted on in silos.

The challenge is that modernization needs to happen under strict controls for auditability, governance, resilience and security. However, maturity is often uneven, with significant operational gaps emerging between on-premises and cloud deployments.

Given this inherent complexity, it's difficult for enterprises to maintain control, especially as global governance standards and regulations evolve. This has several architectural implications, including the need to manage third-party risk and ensure data residency and sovereignty where legally mandated.

The only viable long-term strategy is one where platform consistency comes first, with security and compliance embedded by design, followed by performance and cost/power optimization.

One operating model across on-premises and cloud

As regulations and security threats change, organizations need to keep their options open for sensitive workloads, resilience, and sovereignty without compromising performance or hindering cost optimization. However, many existing hybrid approaches involve using two platforms on distinct tech stacks, which leads to greater complexity, higher costs, and increased operational friction in areas like security posture, governance, and compliance.

Red Hat and Intel advocate for a hybrid approach that's treated as an operating strategy rather than a destination. "A lot of our clients are now looking to consolidate and standardize on a single platform that spans both on-premises and the public cloud," says Ryan Goodson, chief architect at Red Hat. "This also makes it easier for them to migrate their applications between the two."

By standardizing server clusters and environments, identity and baseline policies first—with the help of automation—enterprises can enable both portability and repatriation. That's what gives enterprises the ability to optimize performance and cost control while ensuring adaptability to new and diverse regulations.

Having a flexible deployment model built on a single platform also helps assuage sovereignty concerns, which are especially pronounced in highly regulated sectors like finance. "When you look at something like OpenShift—Red Hat's platform for building, managing and deploying containerized applications—it traverses the various 'where am I deploying' questions, which don't really matter anymore," says Phil King, enterprise solution architect at Intel. "Instead, you can link everything together, and you can do so with sovereignty, so that your data is safe."

From checkbox compliance to continuous assurance

Compliance and security go hand in hand, with compliance being demonstrable control, and security being moving-target defense. Essential protections include encrypting data at rest and in transit, role-based access controls, and the principle of least privilege—all of which must be applied consistently regardless of where data and workloads physically reside.

Goodson shares an effective method of embedding security and compliance into the platform itself: "We do it with automation and policy as code so we can provide a comprehensive platform to ensure we're delivering something that's not just secure, but also easy to consume, because businesses can't really take advantage of their investment until that code hits production."

The emphasis on automation aligns with the fact that compliance and governance increasingly assume that controls are provable, repeatable and current, even as workloads move between environments. For example, Europe's GDPR sets conditions for transferring personal data outside the EEA, requiring organizations to demonstrate appropriate safeguards and controls. In the US, the SEC raises the stakes for rapid, defensible actions on incident materiality and disclosure, making evidence, consistency, and repeatability essential.

Policy as code supports this continuous compliance posture by translating requirements into version-controlled, machine-readable rules that can be enforced automatically across hybrid environments. Because those rules live in a system of record, they produce an audit trail detailing policy changes, who approved them, and where they were applied. This reduces configuration drift and manual input—two of the most common sources of compliance gaps.

Data-in-use is the final frontier in security and compliance

The financial services sector has been reluctant to share and collaborate over data. While protocols for protecting data at rest or in transit are well-established, data must typically be decrypted in memory during processing. This creates a window for exposure, especially in shared computing environments. "Security threats have grown, and now there are in-memory exploits to worry about, but technology like confidential computing can go a long way to helping secure against those threats," Goodson says.

AI can amplify data-in-use exploits due to the sheer volume of data loaded into memory or who might have access to the infrastructure. For example, generative AI-related policy violations are rising due to uncontrolled data sharing with third-party tools, potentially leading to sensitive data being used for model training or exposed through third-party APIs.

"Protecting data in use is really the last frontier of data protection," King says. "We've got it at rest, we've got it in transit, and if we can protect it in use, we've really encompassed all the key areas." To address this need, King recommends confidential computing technologies like Intel® Software Guard Extensions (Intel® SGX) and Intel® Trust Domain Extensions (Intel® TDX) Intel's SGX and TDX, which are designed to help protect applications, containers, and virtual machines on supported platforms that are designed to protect applications, containers and virtual machines. When technologies are directly integrated into enterprise platforms as hardware-based security capabilities and comprehensive provisioning software, enterprises can protect data where work actually happens.

Maximizing performance efficiency with a unified view

Security and compliance need to be strong and operationally sustainable, so they can't come at the expense of performance, availability, or cost control. In other words, performance is part of the risk posture, since latency and instability result in outages which lead to missed SLAs, delayed fraud detection and blind spots in incident response. "Not only do we want to make it safe; we also need to make it fast so that we can deliver on those business objectives," says Goodson.

Observability is a fundamental enabler of performance, cost management, and security. After all, you can't optimize—or secure—what you can't see, even more so across hybrid dependencies like apps in the cloud and data on-premises. Having a platform that bridges hybrid complexity with a unified view allows teams to make decisions that address all three aforementioned areas for any workload.

As an operating strategy, a unified hybrid cloud supports versatile workload placement and improved cost-to-performance efficiency. With data center energy consumption expected to double by 2030 (IEA), power-aware design decisions that reduce the need for overprovisioning will become more important. "Performance is somewhat subjective," says King. "A lot of people chase the highest possible performance, but in many cases, enterprises should instead focus on efficient performance to deliver performance where they need to and perhaps be a little more conservative when they don't."

The operationalization of AI in financial services is a prime example. Many enterprises have yet to see any real ROI from AI projects due to factors like not using the right accelerators. One common mistake is overprovisioning GPU instances for inference. "If we can use CPUs instances instead of the expensive GPUs for something like inferencing, that goes a long way to balancing power budgets and being able to devote those GPUs to where they are really needed," says Goodson. When enterprises are constrained by data center power budgets, that balance becomes critical for realizing ROI.

What to do next

As financial services organizations operationalize AI and modernize core workloads, the hybrid cloud has become the go-to operating model for the future-ready data center—one that meets security, compliance, performance, and cost-control requirements simultaneously.

Here's a rundown of how to get there:

- Assess workloads for adherence to security, compliance and performance policies
- Standardize orchestration on one platform and operating model across environments
- Automate deployment with policy-as-code and repeatable deployment workflows
- Secure your hybrid cloud with zero trust and confidential computing for data in use
- Observe data and workloads with end-to-end visibility across risk, performance and cost
- Optimize workload placement with the right accelerators and power-aware sizing

Intel and Red Hat have been innovating together for over 25 years to bring secure, high-performance computing infrastructure to enterprises around the world.

© 2026 Red Hat, LLC. Red Hat and the Red Hat logo are trademarks or registered trademarks of Red Hat, LLC or its subsidiaries in the United States and other countries. Intel, the Intel logo, Intel SGX, and Intel TDX are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Studio
by informa techtarget

Expert led.
Impact driven.

Studio is Informa TechTarget's global content studio offering brands an ROI rich tool kit: Deep industry expertise, first-party audience insights, an editorial approach to brand storytelling, and targeted distribution capabilities. Our trusted in-house content marketers help brands power insights-fueled content programs that nurture prospects and customers from discovery through to purchase, connecting brand to demand.

Learn more