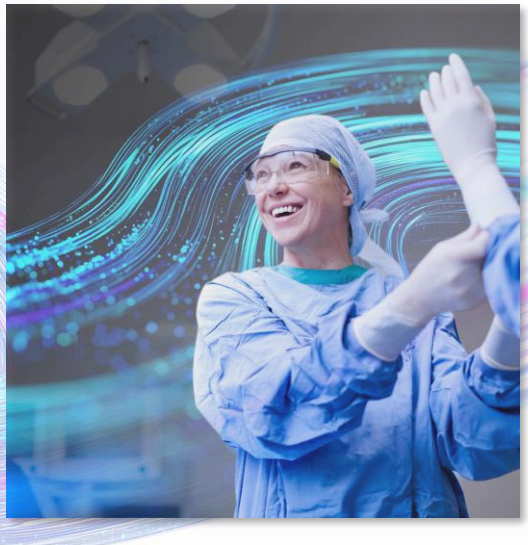
The resiliency gap

Closing the distance between assumed and validated recovery readiness



Explore a practical framework for helping storage and IT teams in healthcare establish recovery readiness.

Introduction

What many healthcare organizations learn too late is that their recovery strategies don't always align with the realities of modern cyber incidents. Traditional disaster recovery is designed for outages and failures. But today, storage and IT teams must be prepared to recover from compromise. Determining which recovery points remain trustworthy may be as important as restoring the data itself.

In healthcare, the challenge is even greater. Patient data is a high-value target, and modern attackers often steal information before encryption begins. A successful recovery may restore clinical systems, but it doesn't undo a HIPAA violation if patient data has already left the building.

For the 13th consecutive year, healthcare experienced the highest average breach costs of any industry at USD 6.64 million.1 As threats become more sophisticated and breaches take longer to identify and contain, IT teams are under increasing pressure to prove that recovery will work—in practice, not just in theory—when patient care depends on it.

This practical eBook—mapped to the National Institute of Standards and Technology (NIST)Cybersecurity Framework—offers a guide to help you demonstrate recovery readiness through testing, validation and continuous assessment.

Together, these practices generate the proof needed to identify trusted recovery points, validate recovery processes and demonstrate that recovery objectives can be achieved under real-world conditions. Teams that use this guide are better positioned to strengthen their organization's cyber resiliency through proven recovery readiness, helping ensure critical healthcare systems can be restored quickly and confidently when a cyber incident strikes.

1. Why validated recovery matters now more than ever

"AI is making attacks faster and cheaper, while breaches keep getting more expensive. When organizations have an extended gap between discovery and remediation, that imbalance shows up directly in breach costs."
Suja Viswesan
Vice President
IBM Security Software2

Healthcare is the biggest target

Healthcare remains one of the most heavily targeted industries for cyberattacks. The reasons are clear: healthcare organizations manage large volumes of sensitive patient information, support critical services that can't tolerate extended downtime, and often operate complex environments where clinical systems, applications and data must remain continuously available.

Attackers are now moving—and causing damage—at machine speed, with AI-driven attacks increasing by 56% since last year.1 In the hands of attackers, AI collapses the time between vulnerability discovery and exploitation. At the same time, ransomware attacks continue to rise. As attacks become faster and more automated, cyber resiliency requires more than traditional backups.

Organizations must be prepared for scenarios that require more than just recovery copies and backup infrastructure.

New and evolving compliance standards

Healthcare organizations face growing pressure to strengthen cyber resilience as HIPAA Security Rule requirements, audit scrutiny and cybersecurity expectations continue to evolve. The HIPAA Security Rule requires covered entities and business associates to protect the confidentiality, integrity and availability of electronic protected health information (ePHI), placing increased emphasis on maintaining resilient systems and safeguarding patient data throughout the recovery process.

Healthcare compliance assessments commonly evaluate:
- Recovery Time Objective (RTO) alignment
- Recovery Point Objective (RPO) alignment
- ePHI confidentiality, integrity and availability
- Recovery testing and validation evidence
- Data integrity verification
- Protection of backup and recovery procedures
- Audit trails and recovery documentation

If an organization can't definitively prove they can restore data and services within mandated timelines, they face financial penalties, immediate operational halts and direct executive liability.

The threat inside the backup

Ransomware changed the recovery conversation. Restore from backup is no longer enough because modern attacks can go undetected for days or weeks before encryption begins, meaning backup copies created during that period may contain compromised data or attacker footholds.

Restoring those backups can unintentionally reintroduce the threat and force organizations to repeat the recovery process.

Bad actors also know backups exist and increasingly target them first, creating situations where there may be little or nothing available to restore. And even when recovery copies remain intact, not all recovery solutions are built for speed or recovery at scale. The ability to rapidly access, validate and restore large volumes of data can significantly extend recovery timelines when business services need to be brought back online quickly.

2. What recovery readiness looks like

Recovery readiness is the ability to demonstrate that recovery objectives can be achieved under real-world conditions. The following framework outlines the activities organizations should perform and the recovery evidence they should collect to demonstrate recovery readiness to internal and external stakeholders.

1) Establish recovery assets you can trust

The first step in establishing recovery readiness is identifying which recovery assets exist and whether they're capable of supporting business recovery objectives A completed backup only confirms that data was copied somewhere—but it doesn't show that the copy is isolated from an attacker, usable by the application or protected from unauthorized changes.

Organizations should first identify the business services that represent their minimum viable business, then map those services to critical workloads, application dependencies and recovery objectives. Multiple recovery paths should be maintained, including immutable snapshots, longer-retention backups and immutable recovery copies. The most critical workloads should receive the strongest protection and most frequent validation. At this stage, the goal is to establish which recovery assets can be trusted before an incident occurs.

Recovery evidence produced
- Recovery asset inventory
- Documented RPO and RTO targets
- Recovery copy catalogue with ownership and protection status
- Evidence of immutable and isolated recovery copies

2) Detect threats before recovery options narrow

Early detection is a critical component of recovery readiness. The sooner suspicious activity is identified, the more recovery options organizations have available.

Security tools remain essential, but storage provides another layer of visibility because ransomware changes data as it is written. AI-driven analytics can help identify abnormal behaviour, giving teams more time to investigate threats, preserve trusted recovery points, and reduce recovery impact.

Frequent, accurate telemetry fuels AI-driven detection by continuously analysing entropy, compression, delta compression and access pattern data.

By establishing a baseline of normal activity, the system can identify anomalies earlier, giving teams more time to investigate threats and preserve trusted recovery points before damage spreads. That earlier intervention can reduce the amount of data that must be restored, minimize business disruption and improve recovery confidence through a clearer understanding of when abnormal activity began and what data was affected.

Recovery evidence produced
- Detection reports and storage telemetry records
- Documented anomaly investigations
- Confirmed detection-to-notification timelines
- Evidence of preserved recovery points following suspicious activity

Validate recovery points

Recovery readiness depends on knowing which recovery points can be trusted when an incident occurs. The most recent backup isn't always the safest, as suspicious activity may begin long before it is detected. Teams should evaluate whether recovery points are accessible, intact, relevant to the incident and capable of supporting the workload being restored. Validation should confirm data integrity, identify signs of compromise and establish a known-good recovery point that predates malicious activity.

While recovery copies may be created at different intervals depending on the protection strategy in use, more frequent point-in-time copies can provide organizations with additional recovery options when an incident occurs. A validated recovery point provides greater confidence that restored data can support business operations without reintroducing risk.

Recovery evidence produced
- Verified recovery point with timestamp
- Documented clean-from date
- Recovery-point validation report
- Evidence supporting recovery-point selection

Test Recovery Readiness

Testing is what transforms recovery planning into recovery readiness Organizations should regularly verify that protection mechanisms are operating as expected, validate recovery points and test recovery procedures under realistic conditions. Recovery exercises should restore complete workloads, measure actual RPO and RTO performance, and include the people, processes and dependencies required to support business operations.

Periodic end-to-end simulations should test the full recovery lifecycle, from detection and investigation through restoration and validation. Teams that run tests almost always find something.

That's the point. An issue discovered during testing can be corrected; one found during an incident can significantly extend downtime and increase business impact.

Recovery evidence produced
- Recovery test reports
- Validated RPO and measured RTO results
- Recovery simulation reports
- Gap analysis and remediation plans
- Historical recovery performance records

Maintain recovery readiness

Recovery evidence loses value if it's not maintained over time. Applications, infrastructure, data and business requirements continually change, which means recovery readiness must be reassessed on an ongoing basis. Organizations should review recovery capabilities whenever critical workloads move, protection policies change new dependencies are introduced or operational responsibilities shift.

Recovery test results should be tracked over time to identify trends, emerging risks and areas for improvement. The goal isn't to create more documentation, but to maintain an accurate understanding of current recovery readiness so stakeholders can make informed decisions when an incident occurs.

Recovery evidence produced
- Recovery readiness review records
- Updated recovery documentation
- Change logs linked to recovery validation activities
- Risk and remediation tracking reports
- Current recovery evidence demonstrating ongoing readiness and remaining risk

3. Recovery readiness requires shared ownership

No single team owns recovery. Infrastructure teams manage recovery assets, security teams investigate and contain threats, and clinical, business and IT leaders establish priorities and acceptable levels of risk. Effective recovery requires all three groups to work from a shared understanding of what must be recovered, how recovery success will be measured and who is responsible for making key decisions during an incident.

Align leadership around recovery risk

Executive stakeholders are accountable for business risk, so recovery discussions should focus on operational and financial impact rather than technical capabilities alone. Some workloads may tolerate hours of data loss, while others require recovery points measured in minutes. Start with the potential consequences of downtime, including disruption to patient care, clinical system availability and exposure of ePHI. Then assess current readiness honestly: When was the last full recovery test performed, and did it achieve the required recovery objectives? If the answer is unclear, that uncertainty is itself a risk.

Rather than seeking approval for a broad resilience initiative, consider proposing a realistic recovery exercise that actively involves executive stakeholders. That allows leadership to observe current capabilities, validate assumptions and identify areas for improvement. The results can provide a stronger foundation for future investment and planning decisions.

Connect IT and security around response and recovery

Storage, infrastructure and security teams each play a role in determining whether recovery efforts succeed. Storage teams may identify unusual data activity, while security teams investigate and respond to potential threats. Without a defined process for sharing information and coordinating decisions, valuable signals may not reach the right people before trusted recovery options narrow.

Organizations should establish clear escalation paths, agree on indicators of suspicious activity and define how recovery points will be evaluated and approved for restoration. Recovery exercises should test not only technical recovery procedures, but also how information moves between teams and who is responsible for key recovery decisions. Effective coordination helps organizations detect, contain and recover from incidents with greater confidence and consistency.

4. Conclusion

Recovery readiness isn't a destination but an ongoing discipline. As applications, threats and clinical systems and business requirements evolve, organizations must continue to test assumptions, validate recovery capabilities and address gaps as they emerge. The framework in this guide is designed to support that process.

The next step is to choose one area, run one test, answer one unanswered recovery question and keep building momentum from there. Every step reduces uncertainty and helps close the resiliency gap between assumed readiness and validated readiness.

Learn how IBM & SCC can help you

IBM helps healthcare organizations move from assumed resilience to validated recovery readiness through a combination of resilient infrastructure, cyber resilience capabilities and recovery-focused innovation.

IBM FlashSystem® data storage arrays thread several parts of the recovery story together. Safeguarded Copy helps preserve protected recovery points that can be evaluated and validated before restoration. Ransomware Threat Detection inspects every IO at the storage layer for data corruption in progress. An additional Safeguarded Copy can help preserve an option when a threat is detected. And, in conjunction with IBM Storage Defender Sentinel, it can add recovery point validation and recovery support

The differentiator is the combination of resilient storage, immutable protection, AI-driven detection, automated recovery point preservation, recovery point validation and integration with wider operations.

Also, IBM's use of computational storage in the FlashCore® Module to handle processor-intensive functions close to the drive helps free array resources for other work, strengthen cyber resilience, accelerate recovery and optimize the storage infrastructure with less manual intervention.

For more information visit https://www.scc.com/insights/resources/scc-healthcare-storage-solutions-with-ibm-flashsystem/

Sources
1. Cost of a Data Breach Report, IBM, 2026.
2. https://newsroom.ibm.com/2026-07-29-ibm-study-one-in-four-malicious-breaches-are-ai-enabled,-costing-companies-6-million-on-average