remparo

RECOVERY IN DEPTH™

What Can We Still Trust?

Rethinking Recovery When the Environment May Already Be Compromised

On February 21, 2024, UnitedHealth Group disclosed a cyberattack against its Change Healthcare unit and told regulators it suspected a nation-state actor. On February 27, Optum told the market it was pursuing multiple paths to restoration and would take no shortcuts. Within days, pharmacies and hospitals nationwide were reporting the worst payment disruption the industry had seen. By March 1, the attacker had been correctly identified as the BlackCat ransomware group, not a nation-state. By the end of the first quarter, the company had booked $872 million in direct costs from the incident.

None of that is a story about incompetence. UnitedHealth is one of the best-resourced organizations in American healthcare, and its early public statements were careful and measured. That's exactly why the timeline matters. If a company with that much security investment can't tell, in real time, who is inside its own network or what it can safely stand behind, then no organization's early read on its own breach is something the rest of its ecosystem can safely rely on. The first question every executive team has to answer in a serious cyber event isn't how fast can we recover. It's what can we still trust.

1. The State of DR and BC Today

Paper 1 in this series covered the core problem with Disaster Recovery and Business Continuity Planning: both disciplines assume the recovery environment is clean, an assumption a natural disaster never breaks and a cyberattack breaks by design.

This paper takes up a version of that gap that catches even well-defended organizations off guard. Most enterprises of any size have layered serious security architecture on top of their DR and BCP programs: defense in depth, Zero Trust, network segmentation, endpoint detection and response. That investment is real and it works. It also creates a false sense of coverage, because strong security architecture and a trustworthy recovery environment are not the same thing, and most executive teams have never had to draw that distinction until they needed it mid-breach. This is where healthcare cyber recovery starts to look meaningfully different from a conventional IT recovery program.

2. Why Recovery In Depth™ Is Needed

Two terms get used loosely enough in board conversations that it's worth defining them precisely before critiquing what they don't do. Defense in depth, per NIST, is an information security strategy that integrates people, technology, and operations to establish variable barriers across multiple layers of an organization. Zero Trust, defined in NIST SP 800-207, is a set of concepts for enforcing least-privilege access decisions in the face of a network viewed as already compromised, built on tenets like per-session access and continuous re-verification.

Here's the distinction that matters most: Zero Trust protects access decisions inside an environment. It was never designed to answer whether the rest of the world should trust that environment again after a real breach. An organization can run a mature, well-tuned Zero Trust deployment that correctly blocks an attacker's next move, and still find itself unable to prove to a payer, a partner, a regulator, or its own board that the compromise is actually contained. Blocking the next move and proving the last one is over are two different problems, and only the first one is what most security architecture is built to solve.

Identity verification, the mechanism Zero Trust ultimately depends on to make its per-session decisions, is itself now a named and active attack target. An April 2024 HHS Health Sector Cybersecurity Coordination Center alert describes attackers defeating help-desk identity verification using stolen personal data, and increasingly, AI voice cloning. When the verification layer itself can be spoofed, a security architecture built entirely around verifying identity has a structural weakness sitting at its foundation.

3. Where This Fits in the Overall Model

Recovery In Depth™ is organized around six questions an executive team has to be able to answer during a cyber event:

1. **What can we still trust?**
2. **What must keep operating?**
3. **How do we communicate if our normal channels are suspect?**
4. **How do we keep revenue and cash moving?**
5. **How do we protect customers, patients, employees, and partners?**
6. **How do we rebuild confidence after the technical incident is contained?**

This paper is about the first question, and it's the foundation the rest of the series is built on. Any serious cyber recovery planning has to start here: an organization can't responsibly plan degraded-mode operations, protect revenue, or rebuild stakeholder confidence on an environment it hasn't first established is safe to build on. Papers 3 through 5 all assume the ground this paper covers.

4. The Recovery Paradox and Air-Gapped Recovery

The Paradox, Stated Plainly

The closer a recovery environment sits to production, the faster it may be able to fail over. The closer it sits to production, the more likely it is to share production's compromise. Most executive teams have never explicitly debated that tradeoff. They inherit it by default, through whatever DR architecture happens to already be in place, and the architecture that makes recovery fastest during a hurricane is frequently the same architecture that makes it least trustworthy after a breach.

Why "Systems Restored" Isn't "Reconnected"

In a highly integrated ecosystem, nothing is fully clean until everything connected to it is clean. Third-party attestation, an incident response firm putting its own name and liability exposure on a letter confirming containment, generally takes weeks and rarely comes in partial form. A breached organization's own self-reported reassurance isn't a substitute for that outside validation, no matter how confidently it's delivered. The Change Healthcare timeline in the opening of this paper shows exactly why: a careful, well-intentioned early statement from the company itself was still wrong about who was inside its network.

Air-Gapped Recovery as the Resolving Architecture

CISA's #StopRansomware guidance defines a genuine air gap precisely: an actual break in the network path that an attacker already inside production cannot traverse. It's not a firewall rule. It's not a separate account on the same network. It's not a replicated environment with tighter permissions. Some practitioners call this a clean room recovery environment, a space built and monitored specifically to stay uncontaminated by whatever happened in production. Ransomware operators actively hunt for and destroy backups they can reach over the network, specifically to remove a victim's ability to recover without paying the ransom. A well-permissioned, monitored, technically "immutable" backup that's still network-reachable is still a target, not a safety net.

It's worth being honest about the tradeoff rather than selling it as costless. Air-gapped recovery has real cost and complexity implications: separate infrastructure, separate identity, and operational discipline to keep it genuinely isolated rather than quietly reconnected for convenience over time. Organizations that adopt it should go in with eyes open about what it takes to maintain.

The Convergence Proof Point

Availity, the largest healthcare information network in the country, launched Rapid Recovery in February 2025, an air-gapped standby environment built explicitly in response to Change Healthcare. Availity described it as a wholly separate and secure backup location, ensuring no contamination from compromised systems, with a five-day recovery target and an 18-hour result in a real-world test. Worth noting here because it's independent validation of the category, not a competitive claim: when organizations serving the same industry, watching the same failure, converge independently on the same architecture, that's evidence the architecture is right, not a coincidence.

Cross-Industry Pairing

- **MGM Resorts, 2023:** reservation, payment, digital key, and customer-service systems were disrupted after attackers reportedly used social engineering against help-desk and identity workflows. It's the non-healthcare proof that identity compromise is business compromise everywhere, not a healthcare-specific problem.

CASE IN POINT: ASCENSION, 2024

Three Clocks, Not One

Ascension detected its ransomware attack on May 8, 2024, engaged Mandiant the same day, and restored core EHR access across its ministries by mid-June. By the metric most dashboards track, that's a five-week recovery. It wasn't the whole story.

Reporting quoting Ascension described negotiations with hundreds of vendors, each with its own requirements, as arduous and time-consuming: one-to-one calls, eight webinars, and ongoing indicator-of-compromise sharing. More than 95% of suppliers were reported as reconnected or reconnecting at the time, but the public record doesn't establish that every supplier required the same written attestation or applied the same acceptance standard. Meanwhile, net days in accounts receivable rose from 46.7 to 78.4 year over year, and Ascension received roughly $874 million in Medicare and commercial-payer advances in fiscal 2024.

Ascension's experience is the clearest public illustration of what "trust" actually costs to rebuild once systems are technically back online: clinical restoration, financial recovery, and partner acceptance all ran on different clocks, and only one of the three showed up on a system-uptime report.

Sources: Ascension Network Interruption Update; Ascension FY2024 MD&A; IT Magazine/WIRED vendor-reconnection reporting; Cybersecurity Dive.

CASE IN POINT: SCRIPPS HEALTH, 2021

Thirty Letters, One Reconnection

Scripps Health operated under downtime procedures for roughly three and a half to four weeks after a May 2021 ransomware attack, restoring Epic and its other major systems by late May. The technical recovery was, by most measures, complete within a month.

Reconnecting to its partner ecosystem took longer and required something DR plans rarely account for: formal proof. CIO Shane Thielman later said Scripps produced roughly 30 vendor assurance letters, and that some vendors initially rejected them and asked for more technical documentation before they'd trust the connection again. What those vendors wanted wasn't a status update. They wanted confirmation that independent outside experts had been involved and that the malware had actually been contained.

Scripps has said these requests didn't ultimately delay its restoration access, which is itself the point worth taking from this case: when an organization has a real attestation package

ready, being asked to prove trustworthiness doesn't have to become a bottleneck. Not having one is what turns the question into a delay.

Sources: AHA Scripps leadership retrospective; Fierce Healthcare; Becker's Hospital Review; IT Magazine/WIRED assurance-letter reporting.

Ascension and Scripps are three years apart, in different-sized systems, with different technical footprints. What they share is the pattern this paper is about: technical restoration and ecosystem trust are two separate recoveries, running on two separate clocks, and only one of them is visible on a system-uptime dashboard.

5. Conclusion

Strong security architecture answers a narrower question than most executive teams assume it does. Defense in depth and Zero Trust tell you whether you can detect and contain an active threat. They don't tell you whether you can be trusted again, by a payer, a partner, a regulator, or your own board, once containment is declared. Those are different questions, and an organization that only prepares to answer the first one will find itself improvising an answer to the second in the middle of a crisis.

Air-gapped, independently administered recovery is the architecture that answers the second question, and Ascension and Scripps both show what it costs an organization to answer it without one: months of parallel effort spent proving what a pre-built attestation process could have made routine. That's the case for trusted recovery environments in one sentence: they turn a proof problem into a planning problem, and planning problems are solvable before the incident starts.

Once an organization knows what it can still trust, the next question follows immediately: what has to keep running on it. People, workflows, and the ability to communicate, even in degraded mode, while full recovery is still underway. That's where this series goes next.

> *"The question is no longer whether you can fail over. The question is whether you can fail over into something safe."*

Coming next: "Staying Open for Business." It covers what it takes to keep operating in degraded mode and to communicate credibly when an organization's normal channels are suspect.

Sources

1. UnitedHealth Group SEC Form 8-K filing, February 21, 2024, and UnitedHealth Group First Quarter 2024 Results.

2. Optum status update, February 27, 2024, on restoration approach; MSSP Alert, "UnitedHealth Group: A Cyberattack Timeline."

3. The Record / Recorded Future News and The Register, reporting on BlackCat/ALPHV attribution for the Change Healthcare attack.

4. CBS News, "UnitedHealth says Change Healthcare cyberattack cost it $872 million."

5. NIST SP 800-53 Rev. 5 and NIST CSRC glossary (defense-in-depth definition).

6. NIST SP 800-207, Zero Trust Architecture.

7. HHS Health Sector Cybersecurity Coordination Center (HC3) Sector Alert, April 3, 2024.

8. CISA joint #StopRansomware guidance (air gap definition and backup-targeting guidance).

9. Availity press release, "Availity Sets New Standard in Healthcare Cybersecurity With Launch of Rapid Recovery," February 2025.

10. CNBC and TechCrunch reporting on the MGM Resorts (2023) cyberattack.

11. Remparo Recovery In Depth™ case research, Ascension (2024): Ascension Network Interruption Update; Ascension FY2024 MD&A; IT Magazine/WIRED; Cybersecurity Dive.

12. Remparo Recovery In Depth™ case research, Scripps Health (2021): AHA leadership retrospective; Fierce Healthcare; Becker's Hospital Review; IT Magazine/WIRED; HIPAA Journal.