remparo

RECOVERY IN DEPTH™

The Recovery Illusion

Why Disaster Recovery and Business Continuity Aren't Built for a Breach

A hospital that loses a data center to a power outage can often fail over to a backup site within the hour. A hospital hit by ransomware may need weeks before any outside partner will trust its network again. Same organization. Same word: recovery. Two completely different timelines, because the second kind of disruption doesn't just knock out infrastructure. It knocks out trust.

Natural disasters destroy infrastructure. Cyberattacks destroy trust. Most disaster recovery and business continuity programs were built to solve the first problem. They were never built to solve the second. That gap, the space between disaster recovery and true healthcare cyber recovery, is what this series is about.

1. The State of DR and BC Today

Disaster Recovery, as most organizations practice it, was built to answer a narrower question than the one healthcare and other regulated industries face today. Traditional healthcare disaster recovery plans were designed around fire, flood, hardware failure, power outages, and regional service interruptions. Those disruptions damage infrastructure without touching an organization's ability to trust what its own systems are telling it. In every one of those scenarios, the recovery environment is assumed to be clean. The disaster took out a data center. It did not take out the organization's credibility.

Hospital business continuity planning covers more ground, at least on paper. ISO 22301 defines a full Business Continuity Management System built around a Business Impact Analysis that sets a Recovery Time Objective, a Recovery Point Objective, and a Maximum Tolerable Period of Disruption for each critical activity. DRI International's Professional Practices, also recognized by the Business Continuity Institute, lay out ten practice areas covering everything from risk assessment to crisis communications. NIST SP 800-34 covers the more technical, systems-focused half of the same problem. These are mature, well-tested bodies of practice.

In most organizations, though, BCP lives inside InfoSec or compliance. It gets tested once a year in a tabletop exercise, then filed away rather than operating as a discipline executives outside those functions think about day to day. Both DR and BCP share a structural assumption that doesn't survive a cyberattack: that recovery runs on a single clock, and that an activity's recovery either meets its Recovery Time Objective or it doesn't. Neither discipline has a native

way to describe a function running at reduced but valuable capacity while full recovery is still underway.

That gap matters more than it used to, because the “sophisticated attacker” story most organizations tell themselves about why breaches happen is largely a myth. Verizon's 2026 Data Breach Investigations Report found the human element present in 62% of breaches across all industries, up from 60% the year before. The share of breaches that trace back to a person, not a purely technical failure, is not improving. Healthcare breaches sit undetected for an average of 279 days before containment: the longest of any industry, and five weeks longer than the global average.

The Change Healthcare breach is the case that makes this concrete. Attackers got in through a Citrix remote-access portal that had no multi-factor authentication on it, a control that violated the company's own policy according to UnitedHealth CEO Andrew Witty's May 2024 testimony before the U.S. Senate. This was not a sophisticated attack against a hardened target. It was one unenforced basic control in a large, complex environment, and it triggered what the American Hospital Association would later call the most significant and consequential cyberattack against U.S. healthcare in history.

None of this is unique to healthcare. Manufacturing, retail, and financial services all run DR and BCP programs built on the same physical-disruption assumptions healthcare does. Take the 2024 cyberattack on CDK Global, the software platform thousands of U.S. auto dealerships depend on for sales, financing, and service. It forced dealerships nationwide back onto paper forms and manual reconciliation for weeks. Different industry, same illusion. Hospital ransomware recovery follows the identical pattern: the industry keeps measuring itself against a disaster recovery standard built for a different kind of disruption.

2. Why Recovery In Depth™ Is Needed

DR and BCP don't fall short because they were built badly. They fall short because they were built to answer a different question than the one a cyberattack actually asks. That's the real difference between disaster recovery and cyber recovery: one assumes the recovery environment is clean, and the other can't.

A cyberattack doesn't just damage a server. It can compromise identity systems, administrative privileges, endpoint management tools, cloud consoles, backup repositories, monitoring platforms, and the communication systems an organization would normally use to coordinate its own response. The organization may not know when the attacker got in, what was touched, or whether the environment it's failing over to inherited the same compromise. This isn't a rare pattern: 96% of ransomware attacks now target backup repositories directly, and 76% of those attempts succeed.

Recovery In Depth™ starts from an uncomfortable assumption that DR and BCP don't: in a serious breach, the production environment may no longer be trustworthy. Instead of asking

only “how fast can we recover,” the framework asks a second question first: how fast can we recover into an environment we know is safe?

To be clear about what that means for existing investment: this is not an argument against DR and BCP. Recovery In Depth™ doesn't replace either discipline. It surrounds both, using the recovery objectives BCP already defines and the technical recovery DR already provides, and adding the operational connective tissue that neither discipline requires by default.

The same logic applies to prevention. Nothing in this series argues for spending less on stopping breaches in the first place. Prevention is the cheapest layer of defense there is, and it should stay funded accordingly. Prevention and recovery are sequential defenses against the same threat, not competing line items in the same budget. An organization that treats recovery planning as optional because its prevention program is strong is making the same mistake as one that skips its fire evacuation plan because the sprinkler system is new.

3. Where This Fits: The Six Executive Questions

Recovery In Depth™ is easiest to understand as a set of six questions an executive team has to be able to answer during a cyber event. They sit above the technical detail of any single system's recovery:

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 makes the case that these are the right six questions to be asking, and shows where the standard DR/BCP playbook goes quiet on each of them. The pieces that follow in this series take them up one at a time: what an organization can still trust; what has to keep operating and how to communicate while it does; what a breach costs in revenue and cash; and how to protect the people affected and rebuild the confidence of everyone the organization depends on.

Positioned against the two disciplines already in place, DR is the technology recovery engine and BCP is the operational recovery playbook. Recovery In Depth™ is the executive discipline that runs technical, operational, financial, and trust recovery in parallel, rather than waiting for one to finish before starting the next.

4. The Illusion Itself: Where DR and BCP Quietly Fail

The Real-Time Replication Problem

Here is the design choice that makes most DR architecture fast and fragile at the same time. Aggressive Recovery Point and Recovery Time Objectives require a live, or near-real-time, connection between production and the recovery environment: continuous replication, shared credentials, synced backup agents, often the same administrative trust. That connection is exactly what an attacker inside production can use to travel into recovery. The same design decision that makes DR fast during a hurricane is what makes it untrustworthy after a breach.

Five Gaps Standard BCP Doesn't Cover

Underneath that architectural problem sit five gaps in how standard Business Continuity practice actually gets applied. They show up consistently across breach after breach, regardless of industry:

- **Binary recovery.** An activity either meets its Recovery Time Objective or it doesn't. There's no standard model for a function running at reduced but valuable capacity in between.
- **Flat recovery targets.** RTOs are set per activity, not nested from company to department to individual technology, so leadership can't see which single technology delay is gating which department-level outcome.
- **Role-specific training only.** People are trained for their own role in their own plan, not for someone else's role during a crisis, which is exactly the capability a breach demands.
- **Siloed risk assessment.** Run department by department, it misses the fast, low-cost, adjacent workaround that doesn't replace a critical system but meaningfully offsets the gap while that system recovers.
- **Financial impact as input, not output.** Standard practice prices out financial impact in advance to justify a security budget, then hands the actual tracking to insurance and finance retrospectively, once incident response has already closed.

What the Illusion Looks Like in Practice

Two recent incidents show what happens when these gaps meet a real attack.

At Universal Health Services, a September 2020 cyberattack forced the company to suspend user access to IT applications across its U.S. operations. UHS didn't stop delivering care. It delivered care using established backup processes, including offline documentation, while it worked to restore electronic medical records, laboratory, and pharmacy systems. Those backup processes weren't an improvisation invented in the moment. They existed because someone, at some point, had planned for a day when the primary systems wouldn't be there. That's proof the fifth gap above is fixable, and that organizations who close it in advance perform better under pressure than organizations that discover the gap during the incident.

Not every organization has that kind of plan in place. And even organizations that recover their core systems relatively fast can be caught by a different version of the same illusion: the assumption that IT actually knows what it's recovering.

CASE IN POINT: ARDENT HEALTH SERVICES, 2023

Faster Recovery, Same Blind Spot

When Ardent Health Services detected ransomware across its hospital system on November 23, 2023, its technical recovery was comparatively fast: core systems, including Epic, were restored within about two weeks, and all 25 of its emergency departments were again accepting ambulances within a week of the attack.

But the recovery process surfaced a different problem. According to a later retrospective account from Ardent's own IT leadership, the investigation uncovered approximately 520 applications running across the environment, against an internal inventory that had expected roughly 390. An organization moving fast to recover a system it thought it understood found, in the middle of the crisis, that it didn't fully know what it was recovering.

In miniature, that's the whole illusion: a DR plan is only as good as the inventory it's built against, and most organizations have never stress-tested whether that inventory is accurate.

Sources: Ardent Health Services official incident disclosure; HealthSystemCIO retrospective; Fierce Healthcare; Healthcare Dive.

CASE IN POINT: COMMONSPIRIT HEALTH, 2022

"Restored" Is Not "Recovered"

CommonSpirit's ransomware attack shows how many different clocks a single incident actually starts. The company contained the attack on October 2, 2022, and by November 9, about five weeks later, most of its markets had regained EHR access. To a dashboard tracking system uptime, that looks like recovery.

The forensic investigation that determined what attackers actually did told a different story. An external firm ultimately found that data had been copied from two file-share servers rather than from the EHR itself, and the detailed file review wasn't complete until February 21, 2023. Individual notifications to affected patients weren't finished until April 2023. The technical outage lasted weeks. The full recovery, by CommonSpirit's own record, lasted closer to seven months.

Sources: CommonSpirit Health official update; Massachusetts filed data-event notice; CommonSpirit 2023 Annual Report; HIPAA Journal.

"Systems restored" and "fully recovered" are not the same event. Any recovery plan that measures itself only against the first is measuring the wrong thing.

This is also where the paradox at the center of the next paper in this series starts to come into view. 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've inherited it by default, through whatever DR architecture happens to already be in place.

5. Conclusion

Disaster Recovery and Business Continuity Planning are not obsolete, and nothing in this series argues that they should be replaced. They're necessary. They are also, on their own, insufficient for a cyber event. Both were built on the assumption that trust survives the disruption, and a cyberattack's entire objective is to make sure it doesn't.

This is the gap Recovery In Depth™ is built to close. Not a replacement for DR and BCP, but the executive discipline that runs alongside them, treating technical recovery, operational continuity, financial protection, and stakeholder trust as four clocks that all have to be managed at once, not four problems solved one after another once the last one is finished.

The rest of this series takes up the Six Executive Questions one at a time, starting with the one every executive team has to answer first, before any of the others matter: not how fast can we recover, but what can we still trust.

> *"Natural disasters destroy infrastructure. Cyberattacks destroy trust."*

Coming next: "What Can We Still Trust?" It covers the Recovery Paradox, why Zero Trust and defense in depth don't answer the reconnection problem, and air-gapped recovery as the architecture that does.

Recovery In Depth™. Rebuilding trust.

Sources

1. Verizon, 2026 Data Breach Investigations Report.

2. IBM, Cost of a Data Breach Report 2025 (healthcare breach identification/containment timeline).

3. TechCrunch and Cybersecurity Dive coverage of UnitedHealth CEO Andrew Witty's May 2024 U.S. Senate testimony.

4. American Hospital Association statements on the Change Healthcare cyberattack.

5. ISO 22301:2019, Business Continuity Management Systems — Requirements.

6. NIST SP 800-34 Rev. 1, Contingency Planning Guide for Federal Information Systems.

7. DRI International / Business Continuity Institute, Professional Practices for Business Continuity Management.

8. Sophos, State of Ransomware 2025 (backup-targeting and compromise-success rates).

9. Public reporting on the CDK Global (2024) cyberattack and its impact on U.S. auto dealerships.

10. Universal Health Services public statement; Healthcare IT News and HIPAA Journal reporting on the September 2020 UHS incident.

11. Remparo Recovery In Depth™ case research, Ardent Health Services (2023): Ardent official incident disclosure; Fierce Healthcare; Healthcare Dive; HealthSystemCIO retrospective; Advisory Board Radio Advisory.

12. Remparo Recovery In Depth™ case research, CommonSpirit Health (2022): CommonSpirit official update; Massachusetts filed data-event notice; CommonSpirit 2023 Annual Report; HIPAA Journal.