remparo

RECOVERY IN DEPTH™

Staying Open for Business

Operating and Communicating When Normal Systems Are Suspect

A nurse charting on paper. A registrar checking eligibility over the phone instead of a portal. A runner carrying lab results down a hallway because the system that used to send them electronically is offline. None of that looks like a security incident from the outside. It looks like a hospital in 1985. It's also exactly what "recovery" looks like for the days or weeks between the moment systems go down and the moment they're fully back, and almost no recovery plan actually plans for it.

The U.S. Navy has a doctrine for this problem, though it wasn't written with cyberattacks in mind. "Every sailor is a firefighter" is more than a slogan. Every recruit, regardless of specialty, gets base damage-control training and a casualty-response station layered on top of their primary job, so the capability to save the ship doesn't sit with one specialist team. It's distributed across the whole crew. Most organizations plan recovery the opposite way: concentrated in IT, with everyone else expected to simply wait.

1. The State of DR and BC Today

Standard Business Continuity practice already defines a Recovery Time Objective and Recovery Point Objective for each critical activity, per ISO 22301. What it doesn't define well is the space in between. An activity either meets its objective or it doesn't. There's no standard vocabulary for a registration desk running at half speed on a paper backup process, or a lab running behind but still running. That's a real capability, and most BCP frameworks have no way to describe it, let alone plan for it.

The deeper issue is scope. Most recovery plans stop at whether IT can restore the system. Very few plans go the next step and ask whether the people and workflows around that system can keep functioning while it's down. A restored server with no one trained to work around its absence in the meantime is only half a recovery plan, and half a plan is what passes for hospital business continuity at most organizations today.

2. Why Recovery In Depth™ Is Needed

Paper 1 in this series named five gaps in standard BCP practice. Two of them are the ones this paper is built around. Training tends to be role-specific: people are trained for their own job in their own plan, not for someone else's role during a crisis, which is exactly the flexibility a breach

demands. And risk assessment tends to run department by department, which misses the fast, low-cost, adjacent workaround that doesn't replace a critical system but meaningfully offsets the gap while that system recovers.

This paper adds a third gap that's specific to communication rather than operations. Most organizations plan to communicate about a breach using the same corporate email system a breach is likely to compromise. If the incident that took down the EHR also touched the email server, the plan for telling employees, patients, and partners what's happening runs through the exact channel that may no longer be trustworthy. That's not a hypothetical edge case. It's the default architecture at most organizations, and it needs to be planned around before an incident, not discovered during one.

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. This paper takes up the second and third together:

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?**

Operating and communicating are treated as one connected problem here because in practice they are. Neither can be planned responsibly without knowing which environment can be trusted, which is the ground Paper 2 covered. This paper builds the operating layer on top of it: once you know what you can trust, what has to keep running on it, and how do you tell people what's happening while it does.

4. Degraded-Mode Operations and Independent Communications

Degraded-Mode Workflows

An electronic health record can't be replaced by paper charting alone, and no one is proposing that it should be. But staff who retain access to a separate, internet-facing eligibility and benefits-check system can keep that one specific function moving while the EHR goes through incident response, without waiting on the EHR's own recovery timeline at all. An eligibility verification outage doesn't have to mean a registration outage if that separation was planned in advance. That's the pattern worth generalizing: identify the narrow, specific functions that can run independently of the system that's down, and make sure staff already know how to reach them before the day they need to.

Pre-Trained Redeployment

A sales team that has never handled a support call is not a redeployment plan. It's a name on a spreadsheet. The clinical parallel is direct: cross-trained surgical staff can be redeployed to a strained emergency department if a breach forces a surgical unit to close, but only if that cross-training happened before the incident. Discovering during a crisis that a plan exists on paper but was never actually rehearsed is its own kind of failure, and a common one.

The Distributed-Capability Framing

The Navy analogy in the opening of this paper connects to a real body of organizational research: Karl Weick and Kathleen Sutcliffe's work on High Reliability Organizations, and specifically their principle of "commitment to resilience," built on distributed, redundant capability across a whole team rather than concentration in one specialist group. Worth being precise about what this is and isn't: the Navy pairing is original framing, drawn from a well-documented public training model, not an established piece of business-continuity doctrine in its own right. It's offered as a useful way to think about the problem, not as peer-reviewed proof that it works in a hospital setting.

There's a bottom-up complement worth introducing briefly: individuals documenting what they personally need to keep their own function running, feeding that into the company plan rather than sitting apart from it. Worth stating honestly rather than overselling: this requires real training investment, it creates duplicate content, and it doesn't remove the need for someone to curate what actually gets kept and trusted. It changes that job. It doesn't eliminate it.

Independent Communications

An organization mid-incident still needs to reach employees, patients or customers, vendors, regulators, insurers, outside counsel, and its own board, through channels that don't depend on the same corporate email system a breach is likely to compromise. That means a pre-established alternate channel, tested in advance rather than improvised during the event: a separate domain, a phone tree, a status page hosted outside the compromised environment. None of this is exotic. All of it has to exist before the incident to be useful during one.

Case Illustrations

- **Cross-industry:** the 2024 CDK Global cyberattack forced auto dealerships nationwide back onto paper forms and manual reconciliation for weeks. Different industry, same operating problem: the software went down, and almost no dealership had a rehearsed plan for running the business without it.
- **A faster comparison:** when Ardent Health Services detected ransomware across its hospital system on November 23, 2023, it moved in phases rather than all at once. All 25 of its emergency departments were again accepting ambulances within a week, even before core systems like Epic were fully restored on December 6. That staged, tracked approach, stabilize the most critical function first while broader restoration continues, is what deliberate degraded-mode planning looks like when it works.

CASE IN POINT: PROSPECT MEDICAL HOLDINGS, 2023

What “Degraded” Actually Looks Like

Prospect Medical Holdings' 16-hospital system identified unauthorized access on August 1, 2023 and didn't have systems fully back online until September 12, about 40 days later. In Connecticut alone, state records documented 29 separate ambulance-diversion notifications over that stretch, including a full 17-day diversion period at Manchester Memorial Hospital.

Facilities ran on paper records and curtailed emergency, outpatient, imaging, blood-draw, and procedure services. Nearly half of elective procedures at the reviewed Connecticut hospitals were postponed. The state ultimately advanced roughly $7.5 million to the hospitals because they couldn't properly file Medicaid claims while their systems were down, a straightforward case of claims filing during ransomware recovery becoming a state-level cash-flow problem rather than just a hospital one.

Prospect's public record supports outside forensic involvement, but not a public attestation process, which means its 40-day technical-recovery window can't be attributed to attestation delay the way some other cases in this series can. What it does show clearly is how many separate things have to keep functioning, clinically and financially, while a full recovery plays out over weeks rather than days.

Sources: Becker's Hospital Review; Connecticut Mirror investigation; Prospect Medical updated notice; Crozer Health substitute HIPAA notice.

CASE IN POINT: SCRIPPS HEALTH, 2021

Running the Hospital Without the Hospital's Systems

Scripps Health's May 2021 ransomware attack is best known in this series for what it took to prove the network was safe again. It's just as instructive for what it took to keep operating while that proof was being assembled. Scripps leadership later described malware propagating laterally across the environment, followed by a deliberate, broad disconnection: sites, data centers, internet access, remote connectivity, phones, and the EHR itself.

What followed was roughly three and a half to four weeks of hospitals running on paper, staff using personal phones to communicate, clinics curtailing activity, and operations relying on runners and local command structures to move information the systems normally moved automatically. Emergency departments used ambulance-diversion procedures in the meantime.

None of that was elegant, and Scripps has never presented it that way. It's a real example of an organization keeping care moving without its usual tools, and it's the clearest illustration in this series of the gap this paper is about: the weeks between systems going down and systems coming back, where the actual work of staying open for business happens.

Sources: AHA Scripps leadership retrospective; Fierce Healthcare financial report; Becker's Hospital Review EHR restoration report.

Prospect and Scripps sit two years apart and tell the same underlying story from different angles: how fast the systems come back is only part of the recovery. What happens to patients, staff, and cash flow while they're down is the other part, and it's the part standard DR and BCP plans are least likely to have rehearsed.

5. Conclusion

Recovery isn't finished when systems are restored. It's finished when people can operate and communicate through the gap that remains until they are. That gap can last days or, as Prospect Medical Holdings shows, well over a month, and an organization's ability to function through it depends entirely on planning done before the incident, not improvisation during it.

Degraded-mode workflows, pre-trained redeployment, and independent communications all cost money to build and money to run through during an incident. None of them are free, and none of them are optional if the goal is staying open rather than simply waiting for IT to finish. This is what healthcare operational resilience actually looks like in practice, not a policy binder but a rehearsed set of workarounds people can execute under pressure. The next paper in this series puts a number on what happens when they aren't there.

> *"Most recovery plans restore systems. Recovery In Depth™ restores the business."*

Coming next: "The Board Case for Recovery in Depth." It puts a number on what a breach actually costs, from financing costs and revenue loss to the insurance gap and the effect on neighboring providers.

Recovery In Depth™. Rebuilding trust.

Sources

1. ISO 22301:2019, Business Continuity Management Systems — Requirements (RTO/RPO definitions).

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

3. DVIDS and Navy.mil coverage of U.S. Navy Surface Force damage-control training (“Every Sailor Is a Firefighter”).

4. Karl E. Weick and Kathleen M. Sutcliffe, Managing the Unexpected: Sustained Performance in a Complex World (commitment-to-resilience principle, High Reliability Organizations).

5. Remparo Recovery In Depth™ case research, Ardent Health Services (2023): Ardent official incident disclosure; Fierce Healthcare recovery report; Healthcare Dive patient-portal update.

6. Remparo Recovery In Depth™ case research, Prospect Medical Holdings (2023): Becker's Hospital Review; Connecticut Mirror investigation; Prospect Medical updated notice; Crozer Health substitute HIPAA notice.

7. Remparo Recovery In Depth™ case research, Scripps Health (2021): AHA leadership retrospective; Fierce Healthcare; Becker's Hospital Review.