EBOOK

BUILDING CYBER RESILIENCE IN HEALTHCARE

The Role of the Minimum Viable Hospital

rubrik

As healthcare providers rely more and more on technology to efficiently deliver care, they've also become increasingly vulnerable to cyber attacks. According to the FBI, Healthcare received the highest number of cyber threats of any critical infrastructure last year.[1]

During the downtime after a cyberattack, hospitals face significant operational disruptions, potentially compromising patient care and safety. The statistics are alarming: Mortality rates increase 28% for ransomware attacks[2], hospitals face a 30% increase in medical errors[3], and the average recovery time for a ransomware attack is 28 days[4].

Unfortunately, the traditional approaches to business continuity and disaster recovery don't address the unique challenges of recovering from a cyberattack. A paradigm shift is required that acknowledges that prevention alone is insufficient and the measures taken to mitigate cyber attacks require an inversion of disaster recovery (DR) planning assumptions. Healthcare organizations must prepare with the worst in mind: a successful cyber attack that disrupts critical patient care for a period of many weeks, and renders production and DR environments unusable.

THE STATISTICS:

28%

Increase in mortality rates for ransomware attacks

30%

Increase in medical errors within hospitals

28 DAYS

Recovery time for a ransomware attack

[1] FBI Internet Crime Report, 2024

[2] Cyber Insecurity in Healthcare: The Cost and Impact on Patient Safety and Care, 2023

[3] The HIPAA Journal: Healthcare Data Breach Statistics, 2025

[4] Eddy, Nathan. "Ransomware Downtime Costs U.S. Healthcare Organizations $1.9M Daily." Healthcare IT News, 31 Dec. 2024

THE POWER OF PREPARATION

The situation isn't all doom and gloom, however. Despite the pace and impact of the attacks, a body of knowledge exists for healthcare providers to turn a potential crisis into a testament to preparedness. Hospitals can significantly reduce recovery time from the attacks of years past and maintain essential services even during severe cyber disruptions.

It starts with understanding exactly how cyber attacks differ from the disaster recovery events of the past, re-prioritizing the traditional "tier-1" list of applications in this new context, creating a plan that acknowledges the loss of trust inherent in these attacks, and focusing recovery procedures on a core set of truly essential critical facilities, clinical, and business capabilities in the hours and days after an attack.

With the right capabilities and preparations, your teams can move swiftly to execute well-drilled plans—creating some level of predictability in the midst of the crisis and ensuring the organization's limited resources are allocated to the most crucial services.

THE MINIMUM VIABLE HOSPITAL:

When building your plan for cyber resilience, a critical part of this is defining the few core applications most critical to maintaining patient continuity of care. This is your Minimum Viable Hospital.

[5] https://www.hipaajournal.com/cost-healthcare-data-breach-2024/
[6] https://www.comparitech.com/studies/ransomware-studies/ransomware-attacks-hospitals-data/
[7] https://www.hipaajournal.com/cost-non-compliance-hipaa/

BEYOND PATIENT CARE:

UNDERSTANDING FINANCIAL IMPACTS

The financial impact of healthcare cyber disruptions extends far beyond ransom payments. Consider these sobering statistics:

$9.77 MILLION[5]

Average cost of healthcare data breach

$1.9 MILLION[6]

Revenue loss per day of downtime

6-7%[7]

Patient loyalty attrition rate after a publicized breach

But these numbers don't tell the whole story. Cyber disruptions are responsible for so many additional costs:

- Fines associated with regulatory and compliance failures
- Legal action from aggrieved parties
- The overhead cost of remediating your technology and infrastructure after an incident

The bottom line: cyber resilience means patient continuity of care, but is also about protecting your operating margin.

A CASCADING EFFECT

Healthcare systems are sometimes referred to as the most complicated form of human collaboration. So when critical applications fail, the negative effects cascade throughout the organization.

SIX SYSTEM FAILURES THAT DIRECTLY IMPACT PATIENT CARE:

Electronic health records become inaccessible, introducing friction into every aspect of the treatment, documentation, coding, and billing process

Laboratory and imaging results cannot be processed or distributed, which has resulted in documented deaths during cyber attacks[8]

Medication dispensing systems fail, creating much higher risk of medical mistakes and greater personal liability for clinicians

Patient monitoring systems become inoperable, taking away the eyes and ears of the nursing staff

Scheduling and registration systems go offline, resulting in a higher rate of patient no-shows and fewer appointments booked

Environmental systems (like positive pressure in operating rooms) and physical security (like access controls) can be affected, slowing medical staff

The uncertainty surrounding which systems can be safely brought back online without fear of spreading malware—and in what order they need to be recovered—can result in extensive delays to recovery efforts. In the pages that follow, we'll explain why and how organizations can minimize this period of uncertainty.

[8] https://www.hipaajournal.com/patient-death-linked-to-ransomware-attack/

THE UNCERTAINTY CHALLENGE

In the wake of a cyber attack, the security and reliability of existing computing environments are usually left suspect. This loss of trust, has only one remedy: to ascertain what the attackers installed, created, or modified via a slow, deliberate forensic process.

This uncertainty represents one of the most challenging aspects of cyber recovery.

Piecing together the forensic picture of what happened shouldn't be rushed and can unfortunately last far longer than organizations (or patients) can afford. In the interim, healthcare providers suffer the extensive consequences of the outage and operational disruption.

Without clear answers, the natural result is analysis paralysis, and an extended period of downtime that predictably results in degraded patient care and unnecessary financial impact to the health system.

IT TEAMS FACE CRITICAL QUESTIONS WITH INCOMPLETE INFORMATION:

WHICH SYSTEMS REMAIN TRUSTWORTHY?

WHERE DID THE ATTACKERS GAIN ACCESS?

WHAT DATA HAS BEEN COMPROMISED?

WHICH BACKUPS ARE FREE FROM MALWARE AND RECOVERABLE?

WHY AREN'T DISASTER RECOVERY PLANS EFFECTIVE AGAINST CYBER ATTACKS?

Traditional disaster recovery plans address physical disruptions—facility loss, power, or connectivity outages—through geographically distant redundant sites. These sites are usually architected to achieve a fast failover. To that end, they're often tightly integrated with production, including stretched networking, stretched clustering, and shared service accounts (credentials for non-human automation processes).

However, this tight integration required for fast failover during a physical attack exacerbates the loss of trust during cyberattacks. When attackers compromise the primary site they typically gain access to the disaster recovery site simultaneously. So how do you design a recovery strategy that addresses the unique challenges of a cyber attack? This table can help you understand the requirements for both disaster recovery and cyber recovery.

	DISASTER RECOVERY	CYBER RECOVERY
THREATS	Natural disasters, hardware failures, site outages, infrastructure disruptions	Malware and ransomware attacks, data breaches, system and identity compromises
STATE OF BACKUPS	Generally trustworthy unless backups are missing or compromised due to physical impacts or outages	Often the target of attacks, with possible deletion or encryption of backups, destruction of backup infrastructure, and lingering backdoors
RECOVERY LOCATION	Always-on or warm failover site with replicated storage	Clean, isolated environment
RECOVERY DATA AND POINT IN TIME	Most recent available data	Unknown. Forensics are needed to determine blast radius, as well as latest malware-free backups
RECOVERY TIME	Fast (if well-tested) and based on replicated data and well-architected failover automation	Slower, unless forensics to determine clean backup points are continuously being conducted, isolated recovery environment is on standby, and the team is well-drilled
TEAMS INVOLVED	Generally just the infrastructure and application teams	Infrastructure, application, security, legal, PR and compliance teams
LINGERING EFFECTS	Some small number of patients affected by a short outage, some short-term revenue loss with most effects limited to a few months time	Extensive patient care impacts due to a long outage, loss of exfiltrated sensitive patient data, regulatory action, extensive litigation costs, class action lawsuits, with full effects felt over a 5 year span

KEY TAKEAWAY:

Revisit your list of critical applications in light of the new threat—a pervasive loss of trust in your security in the wake of a successful attack. Acknowledge in your planning that production resources will be unavailable for an extended period of time. Assume that as currently architected, this loss of trust will extend to your disaster recovery environment. Don't waste time waiting—build into your plan the necessity of an alternate environment, an *Isolated Recovery Environment* (IRE), and build your resilience planning with these constraints in mind.

ZERO HOUR:

CONSIDERING PEOPLE AND PROCEDURES

Even with all the right technological measures in place to rapidly recover applications, there still may be a period of time in which critical systems are completely offline, due to uncertainty and the loss of trust. This means organizations must implement comprehensive non-technical downtime procedures that address patient safety and clinician concerns about personal liability. This conversation must be taken far beyond the usual refrain around how "younger doctors and nurses haven't charted on paper."

Which departments and processes can run without any network access or IT applications for a period of many weeks? The understanding that this won't be a short outage forces planners to make deliberate decisions about resource allocation that prioritize patient safety while the system operates at significantly reduced capacity.

This discussion has to start with patient well-being. Documented procedures should be in place to determine which patients are most vulnerable, what procedures can be safely deferred to free up resources, and which service lines are critical to maintain during a cyber outage. This involves stratifying patients based on acuity and dependency on electronic systems, identifying which surgeries or interventions must continue despite downtime, and determining which outpatient services could temporarily operate with manual processes versus those requiring full system functionality.

The natural tendency is to default to an "all-or-nothing" mindset. Some IT organizations are given wildly unrealistic assignments—to ensure that all services and systems are back online within an impractically short timeframe. Regardless, IT cannot recover everything at once. Senior leaders and clinicians who refuse to engage in prioritizing some services over others must understand that inaction will have an impact on patient safety.

Hopefully, with a clear understanding of patient status, organizations can deploy manual procedures that address their specific needs. That will require knowing how to circumvent some access controls now offline, access to lower tech communication methods, manual double-checks for medication administration, plans for higher staffing levels and additional runners, and surging critical skillsets

throughout the hospital to provide real-time consultation. In some cases these preparations might include renewing the organization's subscriptions to physical reference materials, and finding workarounds for the loss of robotics.

Organizations must also consider the impact on the staff. Anecdotes exist across all industries of staff choosing to retire early or find new employment rather than endure what they've previously experienced elsewhere, which is a testament to the human toll of these events. Employees who have experienced previous cyberattacks may be hesitant to work in what they perceive as an unsafe environment or face grueling recovery hours.

There needs to be clear communication about why standard safety protocols can't be followed during a cyber crisis and staff need to be reassured about their exposure to personal liability—while the organization is running in a degraded state, clinicians will be understandably concerned about the impact on their own licenses and risk of litigation when operating in a higher-risk environment. Additionally, organizations must create sustainable staffing plans that acknowledge staff concerns. They should provide support that helps maintain the larger workforce and increased overtime needed to care for patients during extended downtime periods.

QUESTIONS TO ASK

PATIENT WELL-BEING WHEN SYSTEMS AREN'T OPERATING PROPERLY

01

WHICH PATIENTS ARE THE MOST VULNERABLE?

02

WHAT PROCEDURES CAN BE SAFELY DEFERRED?

03

WHICH SERVICE LINES ARE CRITICAL?

DEFINING THE MINIMUM VIABLE HOSPITAL

The Minimum Viable Hospital concept centers on identifying the absolute minimum technological capabilities required to deliver safe patient care during an interim period measured in weeks. This critical exercise forces healthcare leaders to distinguish between systems that are not immediately urgent and those that are truly essential.

There is a tendency to view all roles and the applications they use as critical, but given the constraints of time, IT personnel available, and compute/storage resources in an IRE, hard decisions will have to be made for an interim period of time while the organization operates in a degraded capacity.

Subsets of the following key systems and applications might be considered:

- Electronic Health Record
- Patient Telemetry
- Pharmacy Systems
- Radiology Information Systems
- Healthcare Information Systems
- Healthcare Coding and Billing
- Communication and Collaboration Tools
- Supply Chain Systems
- Facilities Management
- Employee Management Systems
- Identity and Network Infrastructure

CRITICAL APPLICATIONS TYPICALLY INCLUDE SOME SUBSET OF THESE CATEGORIES:

01

COMMUNICATION AND COLLABORATION TOOLS

> **Priority Functions:** Internal coordination via secure messaging, emergency notification systems, basic virtual care capabilities

02

FACILITIES MANAGEMENT

> **Priority Functions:** Critical environmental controls, physical access controls

03

IDENTITY AND NETWORK INFRASTRUCTURE

> **Priority Functions:** Core identity and authentication functions in the Isolated Recovery Environment, basic connectivity for critical applications, user access and permissions, internal connectivity, standby emergency connectivity for critical vendors

04

PATIENT TELEMETRY

These systems directly support patient care delivery and clinical decision-making.

> **Priority Functions:** Patient monitoring, clinical decision support and clinical alerts, care coordination

05

ELECTRONIC HEALTH RECORD (EHR) SYSTEMS

At the heart of modern healthcare delivery, EHR systems represent the most critical digital asset for any hospital:

> **Priority Functions:** Patient historical information, medication lists, allergies, etc. Documentation for coding and billing, workflow automation.

06

PHARMACY SYSTEMS

Priority Functions: Ability to re-order, inventory and distribute medications, basic dispensing by clinicians, drug interaction checks, compound medications if necessary

07

RADIOLOGY INFORMATION SYSTEMS

Priority Functions: Order and take images, workflow automation, storage/retrieval and diagnostics

08

HEALTHCARE INFORMATION SYSTEMS

Priority Functions: Patient tracking and bed management, basic appointment management, patient registration, insurance/out-of-network information

09

HEALTHCARE CODING AND BILLING

Priority Functions: Documentation of billable services, coding, essential processes necessary to submit claims in a timely manner, timely denial appeal. While coding a billing may not be an immediate priority, how long can your system operate on cash reserves without some amount of income? Remember that full “normalcy” may require an extended period of time

10

SUPPLY CHAIN SYSTEMS

Priority Functions: Critical supplies tracking, basic procurement functions, distribution management with supply allocation to critical areas

11

EMPLOYEE MANAGEMENT SYSTEMS

Priority Functions: Basic shift coverage tracking, basic payroll operations, critical role coverage management, credential verification

APPLICATION DEPENDENCY MAPPING

As with all forms of contingency planning, understanding the interdependencies between critical applications is crucial and more easily said than done. For each critical application, it's necessary to know:

Supporting infrastructure requirements

Database dependencies

Authentication systems needed

Dependencies on integration points with other systems

Degradation modes—can some functionality operate in a degraded mode if dependencies aren't met?

Who owns this application and can provide testing post-recovery

This mapping exercise identifies the true scope of recovery requirements to operate on an interim basis, and those core processes and supporting applications for which no downtime procedure exists. With this prioritized list of applications and dependencies, you have defined the Minimum Viable Hospital around which you will build your IRE and recovery processes.

UC SAN DIEGO'S PROJECT CRASHCART

Acknowledging the risk of rootkits and how vulnerable connected medical devices can be, some organization's recovery plans involve extending the loss of trust to every existing device and piece of network infrastructure, until proven otherwise. While Rubrik focuses on applications in the datacenter and cloud, parallel thinking is being done on how to address bedside care and provide reliable, clean networking to the medical devices necessary for monitoring patients and dispensing medications. A notable example of this is the UC San Diego Center for Healthcare Cybersecurity Project CRASHCART.

THE ISOLATED RECOVERY ENVIRONMENT

The uncertainty challenge and loss of trust makes it critical to have a trusted location for your restored systems. If your usual production and disaster recovery environments are unavailable during forensic analysis, where can you quickly host your recovery?

An isolated recovery environment (IRE) provides a secure foundation for restoring critical applications without risk of reinfection.

SIX KEY ELEMENTS OF AN IRE:

Isolated infrastructure without dependencies on anything in production or DR environments. The IRE must assume all else is unusable, including connectivity

Carefully segmented, logically air-gapped isolated infrastructure, with tightly restricted access procedures, no shared credentials, careful attention to using Privileged Access Workstations to avoid leaking cached credentials

Sufficient capacity (or a plan to rapidly add capacity) necessary to run the set of applications identified as the MVH

The ability to rapidly restore necessary infrastructure services like Active Directory and DNS into the IRE, back to a previous point in time

Standby alternate connectivity and plans for how clean endpoints/clients given to critical roles and personnel will connect to restored applications without reliance on existing network equipment

Regular automated testing and validation

A well designed and clearly defined IRE significantly reduces uncertainty and accelerates the process of restoring applications in order to minimize disruption while the forensic work continues in other environments.

BUDGETING FOR AN IRE MAY BE EASIER THAN YOU THINK

Many health systems were already operating at the very limits of their budgets before factoring in the new necessity of having an IRE. However, some leaders are re-thinking their approach to continuity and how they deploy the budget they already have. Rather than thinking of an IRE as an entirely new and additive set of requirements, some organizations are considering de-prioritizing a fast failover in a physical disaster and instead focusing on the cyber context and a strict isolation of those compute and storage resources. Here's why it matters.

The main reason disaster recovery environments aren't of much use in a cyber recovery is that they are tightly coupled to production in order to achieve a fast failover in the event of a physical failure at the production site. Attackers gain access to both sites concurrently when they compromise the right credentials. However, these physical failures (connectivity, power, fire, flood, earthquake) are happening less frequently than cyber attacks. Hurricanes are usually predicted well in advance.

A debate is taking place whether the answer to affording an IRE is solved by simply re-architecting DR environments into IREs, and making failovers for physical disasters into a secondary use case. The emphasis would be on isolation and preserving trust in the environment. The organization would then need to accept that a physical DR event would take a longer period of time to accomplish, with automation needed to speed up that process. Prioritizing cyber recoveries over physical recoveries arguably more directly addresses the higher probability threat to patient safety.

PLAN OF ACTION:

HOW TO PREPARE FOR THE NEXT CRISIS

After defining the critical applications necessary to maintain interim operations, organizations should consider what assets, training, resourcing, and routine drills/validation need to be established. This will help everyone involved with the response to know their roles when the next crisis arrives—and have confidence that processes and resources will be there to support them in that time of need.

01

REQUIRED RESOURCING

> Identify minimum resources needed for critical functions to implement the plan—people, training, and physical resources.

02

EMERGENCY RESPONSE PROTOCOLS

> Develop clear procedures for various disruptions and understand the key differences between a disaster recovery event and a cyber recovery event. Also identify who makes decisions at the time of the crisis.

03

COMMUNICATION PLANS

> Establish protocols for stakeholder communication during crises that aren't dependent on technology systems that could be affected by a cyber attack.

04 DATA PROTECTION AND RECOVERY

Implement robust data protection measures that will withstand a determined and well-executed attack. Many legacy solutions with a large attack surface fail this critical test. In fact, 74% of organizations said that malicious actors were at least partially able to harm backup and recovery options.

05 AUTOMATED APPLICATION RECOVERY RUNBOOKS

Pre-developed application restoration runbooks that can be executed for a predictable recovery for each of the MVH applications. Focus on building automation in advance that results in a testable outcome, not on heroics in the wake of an incident when personnel are exhausted.

06 TESTING AND DRILLS

Regularly assess and improve the plan through end-to-end testing and simulations, which are thankfully easier to conduct than DR tests, as the target is an isolated environment that can't interact with production systems.

07 MAINTENANCE AND ADAPTATION

Ensure the plan is revisited and is updated and scaled as your environment grows or changes.

[9] Rubrik Zero Labs, A Distributed Crisis, 2025

BECOME CYBER RESILIENT WITH RUBRIK

Implementing the Minimum Viable Hospital framework requires more than just traditional backup. Rubrik Security Cloud provides the comprehensive data security foundation needed to support MVH initiatives through:

ZERO TRUST DATA SECURITY

- Isolated, immutable backups that are truly resistant to encryption and tampering, that will survive the attack and be available for recovery
- Mandatory multi-factor authentication for any type of access
- Quorum authorization, to limit the impact of stolen administrator credentials and to negate rogue actor threats
- Automated, multi-level anomaly detection to identify attacks in progress
- Automated threat monitoring to identify potential introduction of malware

RAPID RECOVERY CAPABILITIES

- Instant access to critical application data without requiring a rebuilding of the backup infrastructure
- The ability to very rapidly threat hunt through the entire data estate to determine which backups are malware free and can be safely restored
- Live mount capabilities for accelerated system restoration

RUBRIK CUSTOMER VOICE:

“We have 2.5PB of data and millions of patient records in our environment we have to secure every day. Cyber resilience to St. Luke’s is absolutely crucial to ensure that we have the right security foundation. Moving to a system like Rubrik that was much more dynamic and integrated with our environment was essential for us.

DAVID FINKELSTEIN

CISO, St. Luke’s

St Luke’s UNIVERSITY HEALTH NETWORK

- Granular recovery options with automated recovery workflows
- Hybrid recovery across on-premises, cloud, and SaaS workloads, including cross-region restores

THREAT INTELLIGENCE

- Continuous threat monitoring to automatically identify and alert on millions of known indicators of compromise (IoCs), and determine scope and time of infection
- Anomaly detection to proactively alert teams to potential threats
- Pre-emptive Recovery Engine pre-hashes all files protected, to speed threat hunts — scan up to 75,000 backups within an estimated 60 seconds for new or emerging IoCs using pre-compiled tables of file hashes
- Ability to natively threat hunt using YARA rules, file hashes, or file patterns

SENSITIVE DATA DISCOVERY

- Automatically discover PHI and PII across on-premises, cloud, and SaaS environments
- Get alerts on overexposed, misconfigured, obsolete, or misplaced data to mitigate exposure risks and HIPAA violations
- Increase governance over sensitive data across your organization with readiness reporting for compliance purposes

IDENTITY RECOVERY AND RESILIENCE

- Accelerated recovery of an entire Active Directory Forest to new servers, and Entra ID recovery, inclusive of Enterprise Apps and App Registrations
- Repeatable, wizard-driven recovery
- Identify anomalous privilege escalations and mitigate risks from human and non-human identities

RUBRIK CUSTOMER VOICE:

Rubrik helps us secure our crown jewels — EMR and the privacy of all the data within that EMR. Not only are we subject to regulatory requirements through HIPAA and other regulations, but we have an obligation to our patients to make sure that we maintain the integrity as well as security of their data and ensure their privacy.

DARIN PRILL

CTO, OU Health

OU Health

OVERCOME THE MOUNTAIN OF UNCERTAINTY

When building out crisis plans, it's important to note the psychological impacts of planning. When starting out, you might find both anxiety and resistance from team members as you start to uncover some of the gaps in your preparation and understanding. However, asking these questions and getting started is critical to gaining confidence that you will know what to do when disaster strikes.

WHAT APPLICATIONS ARE MOST IMPORTANT?

CAN THOSE RUN IN ISOLATION?

WHAT ORDER DO I BRING THINGS UP IN?

WHAT HAPPENED?

HOW LONG WERE THEY IN OUR ENVIRONMENT?

WHAT MALWARE WAS INTRODUCED, AND WHEN?

WHAT'S AFFECTED?

WHAT DOES EACH APPLICATION CONSIST OF?

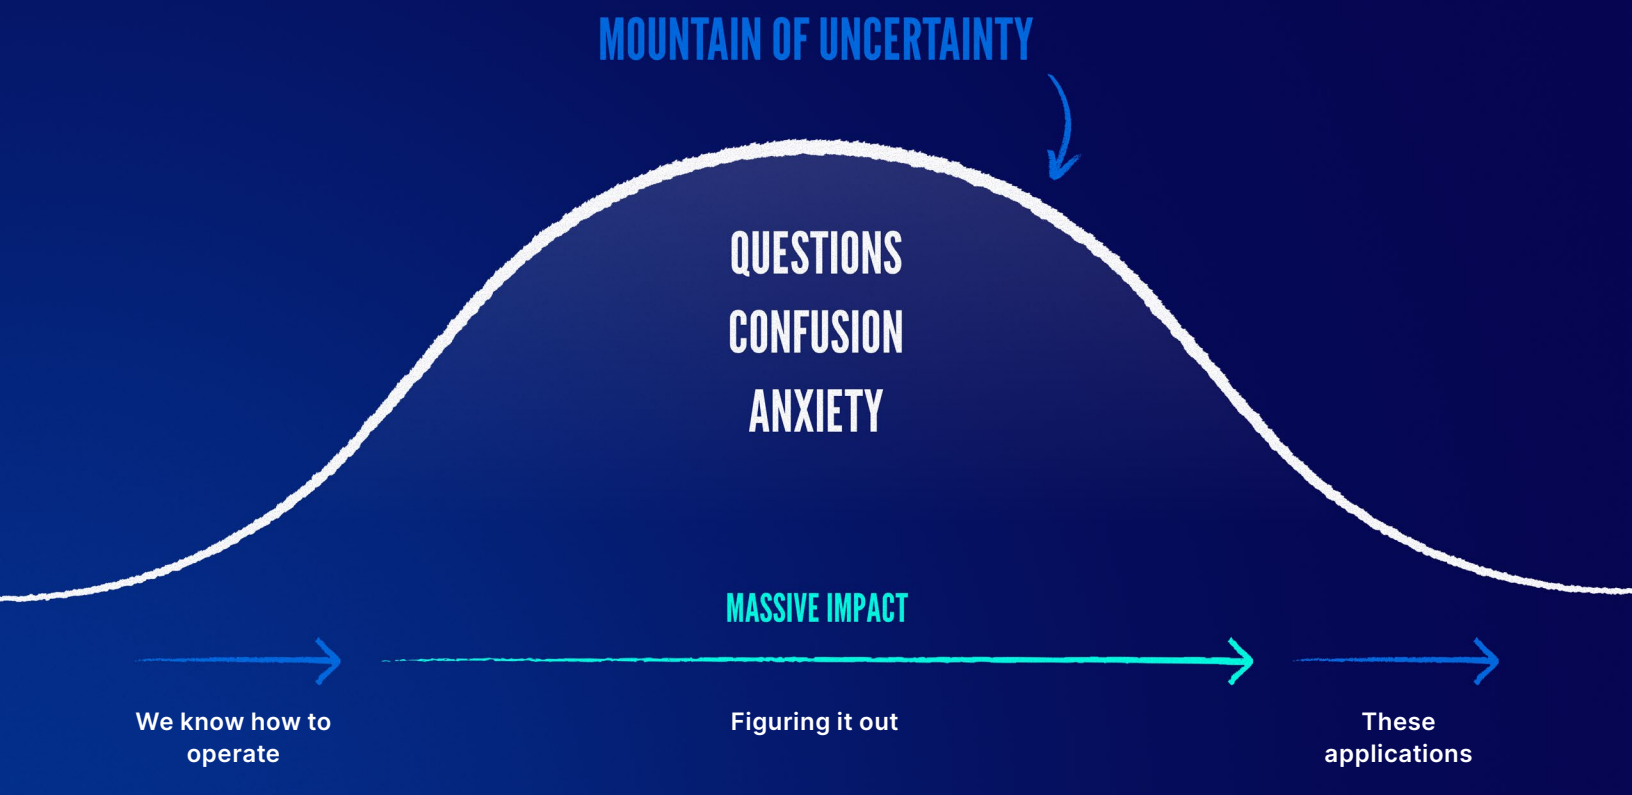

BUILDING YOUR MVH:

PRACTICAL STEPS TO GET STARTED

GOAL	WHAT TO DO	WHY IT MATTERS
IDENTIFY YOUR TOP CRITICAL SERVICES	Convene a brief working session with key stakeholders (e.g., executives, operations, IT) to quickly identify the three services that the organization literally cannot function without.	Starting small focuses your efforts on the most patient-critical processes and prevents "analysis paralysis." You can expand from there once you have clarity on your foundational services.
IDENTIFY ESSENTIAL DEPENDENCIES & REQUIRED RESOURCES	For each of your identified services, list the people, technologies, and third-party providers required to keep them operational. Consider dependencies like network infrastructure, cloud resources, or specialized skill sets.	It reveals single points of failure and areas that need extra attention or redundancy. It also helps you prioritize data protection and recovery strategies.
ESTABLISH A SIMPLIFIED COMMUNICATION FRAMEWORK	Create a quick-reference contact list and escalation path, including primary, secondary, and tertiary channels (email, phone, messaging platform) for both internal and external stakeholders.	Even if you haven't formalized a complete MVH plan yet, a solid communication structure ensures clear, authoritative messaging during a crisis—a pivotal first step in mitigating damage and confusion.

Rubrik is designed to help you maintain patient care, even during cyber attacks. To explore what Rubrik has to offer, check out our product demo today.