INDEPENDENT RESEARCH

2026 EDITION

# The State of Enterprise Recoverability 2026

*From Backup and Recovery to Demonstrated Business Resilience*

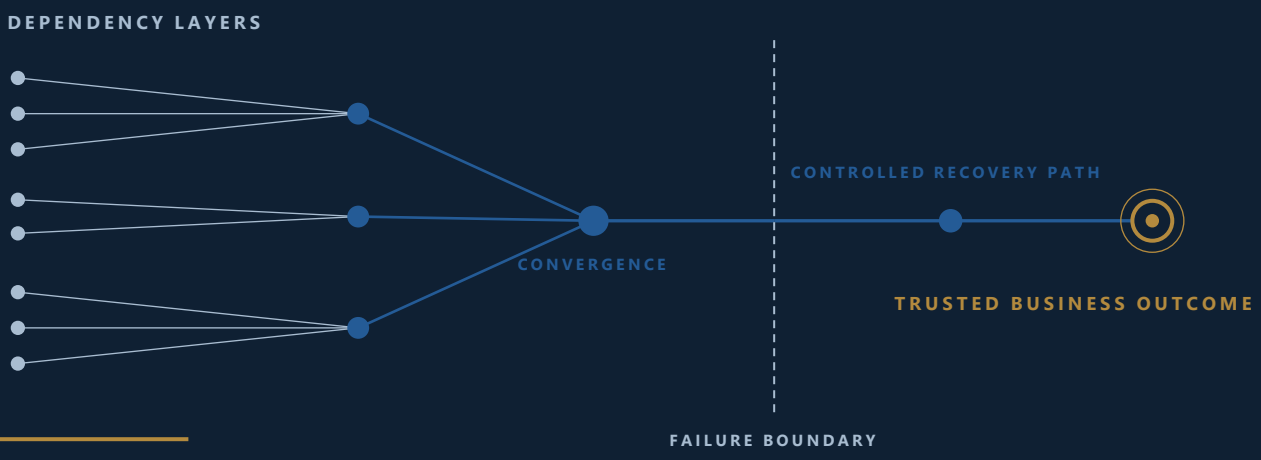


KARIM GABER SALEH

Independent Research

September 2026

EXECUTIVE THESIS

# Backup is getting better. Recovery is becoming harder to prove.

Enterprise backup and recovery capabilities have advanced materially, particularly in immutability, cyber recovery and cloud-native protection.
The remaining challenge is increasingly one of demonstrated recoverability: proving that trusted critical business services can be restored, at sufficient capacity, within acceptable disruption tolerances across complex identity, technology and third-party dependencies.

THE SHIFT

| | | |
|---|---|---|
| 01 | **BACKUP** | Protect the data |
| | ↓ | |
| 02 | **CYBER RECOVERY** | Protect the ability to reconstruct |
| | ↓ | |
| 03 | **OPERATIONAL RESILIENCE** | Protect the business outcome |

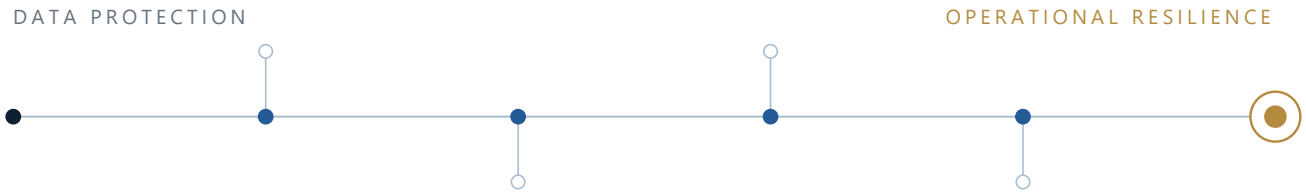


## RECOVERABILITY

*The bridge between data protection and operational resilience.*

RESEARCH METHODOLOGY

# How this research was built

## RESEARCH PROCESS

01 Define the research questions and hypotheses

02 Review current market and technology evolution

03 Prioritize regulatory and primary evidence

04 Examine real disruption and recovery incidents

05 Compare enterprise adoption and testing evidence

06 Test findings across cloud, SaaS, identity and cyber recovery

07 Deliberately challenge the thesis with contradictory evidence

08 Retain only conclusions that survived the challenge

## EVIDENCE HIERARCHY

| | |
|---|---|
| **Tier A** | Regulators, legislation, government guidance, SEC/corporate disclosures |
| **Tier B** | Independent research, incident-response research, industry foundations |
| **Tier C** | Vendor technical evidence |
| **Tier D** | Vendor surveys and commissioned research |

> *Major conclusions were triangulated across multiple evidence types. Vendor evidence was used primarily for technology capabilities and directional market signals, not as sole support for market-wide conclusions.*

RESEARCH MAP

# The research question

> How recoverable are modern enterprises in 2026, and where does the gap lie between possessing backup and resilience technologies and demonstrating recovery of critical business services under realistic disruption conditions?

THE RESEARCH PATH

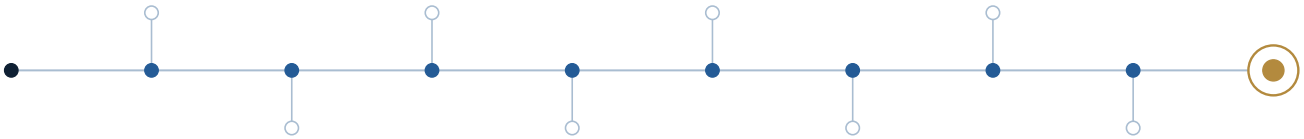

# 01

# Backup is not standing still.

Modern data protection has moved well beyond scheduled copies and offsite storage.

| TRADITIONAL BACKUP | | MODERN DATA PROTECTION |
|---|---|---|
| Scheduled copies | → | Immutability |
| Offsite storage | | Isolation |
| Retention | | Threat detection |
| Restore | | Cloud and SaaS protection |
| | | Recovery orchestration |
| | | Cyber recovery |

**The baseline has moved from “Can we restore the data?” toward “Can the recovery path survive disruption?”**

# Backup has become a resilience platform.

The modern category now spans protection, cyber recovery, cloud, SaaS and increasingly security-oriented capabilities.

| | | |
|---|---|---|
| 01 | **PROTECT** | Backup Snapshots Replication CDP |
| 02 | **SURVIVE** | Immutability Offline copies Air gaps Isolated vaults |
| 03 | **DETECT** | Anomaly detection Ransomware indicators Malware / IOC scanning |
| 04 | **RECOVER** | Clean recovery points Recovery orchestration Isolated recovery environments |
| 05 | **COVER** | Hybrid infrastructure Public cloud SaaS Kubernetes and modern applications |

Current Gartner Backup and Data Protection Platforms requirements include hybrid and multicloud protection, SaaS/data-service protection, ransomware detection, immutable storage integration and recovery capabilities.

Source: Gartner, Backup and Data Protection Platforms, 2026

# Recovery protection is becoming mainstream.

**76%**

of respondents reported using air-gapped backups

UK RESILIENCE SURVEY

**63%**

reported using immutable backups

UK RESILIENCE SURVEY

**66%**

of ransomware incidents involving encrypted data used backups for recovery

GLOBAL RANSOMWARE SURVEY

**The evidence does not support a narrative that enterprise backup is broadly failing. Organisations are adopting stronger survivability controls, and backups remain a major recovery mechanism.**

Survey methodologies and populations differ. These figures should be read as directional evidence, not as global adoption rates.

SOURCES

Databarracks, Data Health Check / State of Resilience 2026, 500 UK IT and resilience professionals

Sophos, State of Ransomware 2026, 2,158 respondents across 17 countries

# Better protection does not automatically mean proven recovery.

**PROTECTION MATURITY ↑**

- Immutable copies
- Air gaps
- Cloud vaults
- Ransomware recovery
- Broader workload coverage

**RECOVERY EVIDENCE ↑**

- Full-service testing
- Identity recovery
- Dependency recovery
- Third-party recovery
- Recovery at required scale and time

**65.1%**

said they were confident they could recover from ransomware within 48 hours

VERSUS

**35.4%**

had actually met RPO and RTO targets in full recovery testing

One 2026 survey found a substantial difference between recovery confidence and demonstrated performance. It is not globally representative, but it illustrates the question this research investigates.

Source: Arcserve, State of Data Resilience 2026, 200+ IT respondents

**The next resilience challenge may therefore be less about acquiring recovery technology, and more about proving that it works under realistic conditions.**

# 02

# A backup copy is the beginning of recovery. Not the end.

Modern recovery depends on a chain of capabilities. Data must exist, survive the disruption, be trusted, be reconstructed into a viable environment, be restored in dependency order, and ultimately return the business service to an acceptable operating level.

**BACKUP**

"Can we restore the data?"

↓

**RECOVERABILITY**

"Can we restore the trusted business service?"

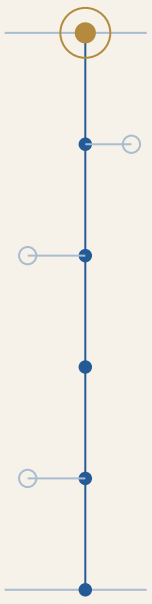

# The Recoverability Stack

Six conditions must hold before recovery becomes a business outcome.

01 **COPY**

Do the required recovery assets exist?

↓

02 **SURVIVE**

Will those assets remain available after the disruption?

↓

03 **TRUST**

Can a known-good recovery state be identified?

↓

04 **RECONSTRUCT**

Can trusted identity, control planes, compute capacity, data and application components be recreated in a viable target environment?

↓

05 **ORCHESTRATE**

Can dependencies be restored correctly and at operational scale?

↓

06 **OPERATE**

Can the critical business service function at an acceptable operating level within its impact tolerance?

Analytical framework developed for this research. The underlying recovery and continuity principles draw on established disaster recovery, cyber recovery and operational resilience practice.

# Immutability solves one problem. Not the whole recovery problem.

**IMMUTABLE BACKUP** — SURVIVE ✓

Recovery copies survived the attack.

**TRUST** — UNRESOLVED ?

Which recovery point is actually clean?

**RECONSTRUCT** — UNRESOLVED ?

Can identity, control planes and trusted compute be restored?

**ORCHESTRATE** — UNRESOLVED ?

Can applications and dependencies be rebuilt in the correct sequence?

**OPERATE** — UNRESOLVED ?

Can the business service actually function within tolerance?

**Modern immutable and isolated backup architectures materially reduce recovery risk. But preserving the recovery assets does not establish that they are clean, provide a viable execution environment, resolve dependencies or restore the business outcome.**

NIST recovery guidance distinguishes protected recovery copies from identifying and restoring a last-known-good state.

Source: NIST NCCoE, Data Integrity: Recovering from Ransomware and Other Destructive Events

# What exactly are we trying to recover?

**DATA**

Files, databases, application state

↓

**SERVER / VM**

Compute environment

↓

**APPLICATION**

Usable technology service

↓

**DEPENDENCY CHAIN**

Identity, DNS, PKI, network, databases, APIs, SaaS, third parties

↓

**CRITICAL BUSINESS SERVICE**

The outcome customers, employees or markets actually depend upon

## A technically successful restore can still be a failed business recovery.

| | |
|---|---|
| Database restored | **30 minutes** |
| Application available | **90 minutes** |
| Business service usable | **6 hours** |

ILLUSTRATIVE EXAMPLE, NOT RESEARCH DATA

The relevant recovery boundary depends on the organisation and service. The principle is that resilience ultimately has to be evaluated against the business outcome being protected.

Operational-resilience frameworks from the FCA/PRA assess resilience at the level of important business services and their supporting dependencies.

Sources: Financial Conduct Authority, Operational Resilience guidance and 2026 observations
Bank of England / PRA, Operational Resilience framework

# 03

# Systems recover. Businesses recover differently.

Major incidents show that recovery is rarely a single technical event. Containment, trust, dependencies, workarounds and operational capacity can extend disruption well beyond the restoration of individual systems.


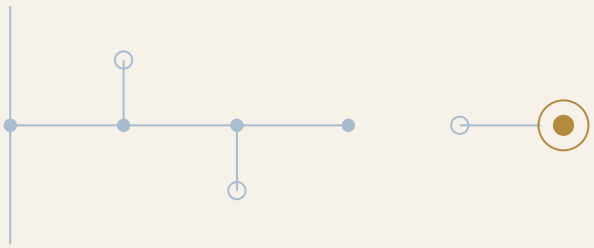

**The recovery clock experienced by the business may continue running after the technology is technically available.**

# One provider failure can become an ecosystem recovery problem.

## CHANGE HEALTHCARE / 2024

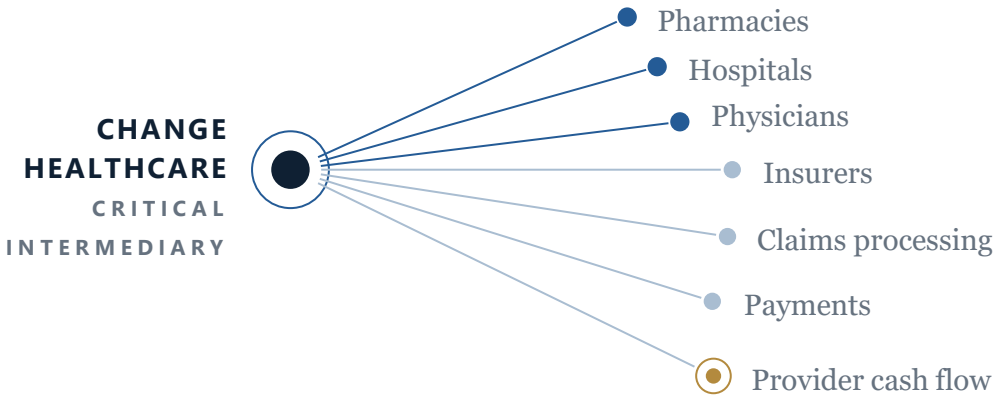


### DISCLOSED RECOVERY TIMELINE

**21 FEBRUARY 2024**

UnitedHealth isolates Change Healthcare systems after detecting the cyberattack.

**MARCH 2024**

Pharmacy services progressively return. Claims testing and customer reconnection begin.

**22 APRIL 2024**

UnitedHealth reports approximately 80% of Change functionality restored across major platforms and products.

**$9B+** accelerated funding and interest-free loans provided to healthcare providers by June 2024.

**The incident demonstrates that recovery of a critical intermediary can propagate across organisations whose own technology environments were not compromised.**

### RESTORED ≠ RECONNECTED

Some customers performed their own security assessment before reconnecting to restored Change systems.

Sources: UnitedHealth Group / SEC filings, February–June 2024; relevant healthcare-provider SEC disclosures

# Cyber recovery can require deliberately stopping the business.

### MGM RESORTS / 2023

Systems were deliberately shut down as part of containment.

**Approx. $100M**

estimated negative impact to Adjusted Property EBITDAR.

**Less than $10M**

estimated one-time incident expenses.

**A technically functioning system may still need to remain unavailable when continued operation cannot be trusted.**

Source: MGM Resorts International SEC disclosures, October 2023 and FY2023.

### CLOROX / 2023

Manual ordering and processing procedures were activated after systems were taken offline. Operations continued at reduced capacity.

**23–28%**

expected year-on-year Q1 sales decline reported during the disruption.

**Approx. $49M**

incremental cyberattack-related costs through December 2023.

**Continuity can preserve operation without preserving normal business capacity.**

Source: The Clorox Company SEC disclosures, 2023–2024.

## AVAILABILITY IS NOT BINARY.

Recovery can move through:


Reported sales decline reflects overall business impact during the period and should not be read as a direct measure of cyberattack cost.

# Your recovery time can be determined by someone else.

CDK GLOBAL ECOSYSTEM / 2024

**CDK PLATFORM**

**Dealer management**

**CRM / inventory / financing / service / accounting**

**Automotive dealerships**

**Customer transactions**

**19 JUNE 2024**

Outage begins.

Dealerships remain physically open using alternative/manual processes.

**8 JULY 2024**

Asbury reports full primary CDK functionality restored, with additional integrations returning subsequently.

**$47.2M** Sonic Automotive's estimated negative pre-tax impact from the CDK outage during 2024.

**INTERNAL SYSTEMS HEALTHY**
**+ CRITICAL THIRD PARTY UNAVAILABLE**
**= BUSINESS SERVICE IMPAIRED**

An enterprise cannot independently restore a service operated by a critical external provider. Continuity therefore depends on workarounds, alternative providers, architectural diversity or waiting for the provider to recover.

Some dealerships using different dealer-management platforms were unaffected, illustrating the relationship between concentration and resilience. Diversification can reduce concentration exposure but introduces cost and complexity. It should be evaluated against business-service impact tolerance.

Sources: Asbury Automotive Group, Sonic Automotive, Group 1 Automotive and Lithia Motors SEC disclosures, 2024

**Across these incidents, the recurring failure boundary was not simply the backup copy. It was the path from restored technology to trusted business operation.**

# 04

# The recovery plan says four hours. But what exactly recovered?

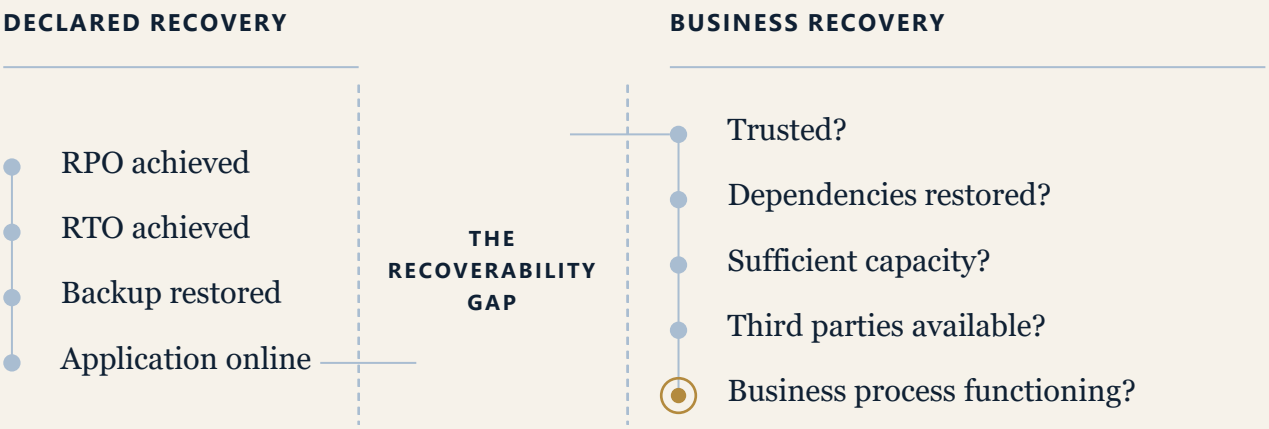


The gap between stated or technically demonstrated recovery capability and the evidenced ability to restore a trusted critical business service to an acceptable operating level within its impact tolerance.

RESEARCH CONSTRUCT DEVELOPED FOR THIS PUBLICATION

**A successful restore is evidence of recovery capability. It is not necessarily evidence of business recoverability.**

# Recovery confidence can exceed recovery evidence.

**65.1%**

said they were confident they could recover from ransomware within 48 hours

VERSUS

**35.4%**

had actually met RPO and RTO targets in full recovery testing

**29.7**
percentage-point difference

CONFIDENCE-TESTING DIFFERENCE IN THIS SURVEY

**24.1%** had never performed a full recovery test

**35.9%** reported successful recovery but outside their required recovery timeframe

**This survey does not establish a global recovery-readiness rate. It does illustrate a recurring assurance problem: confidence and architecture can exist without equivalent evidence that recovery objectives will be achieved under full testing.**

Source: Arcserve, State of Data Resilience 2026, 200+ IT respondents. Vendor-sponsored survey. Used as directional evidence and triangulated with regulatory observations.

> The FCA has separately observed firms asserting that they could recover from any scenario without providing evidence from sufficiently severe testing.
> Source: Financial Conduct Authority, Operational Resilience Insights and Observations, 2026

# Recovery has more than one clock.

01 **INFRASTRUCTURE RECOVERY TIME**

When the required technology environment is available.

↓

02 **APPLICATION RECOVERY TIME**

When the application is technically operational.

↓

03 **MINIMUM VIABLE SERVICE TIME**

When the business can operate at an acceptable degraded level.

↓

04 **FULL OPERATIONAL RECOVERY TIME**

When the service returns to normal operating capacity.

**TRUSTED RECOVERY TIME / SPANS THE TIMELINE**

When the environment is sufficiently validated to reconnect and operate safely.

These are analytical distinctions used in this research, not proposed replacements for established RTO/RPO terminology.

| | |
|---|---|
| Infrastructure restored | **01:00** |
| Application online | **02:00** |
| Minimum viable business service | **05:00** |
| Normal operating capacity | **18:00** |

ILLUSTRATIVE ONLY, NOT RESEARCH DATA

**Traditional RPO and RTO remain essential. But a technical RTO alone may not reveal when the business outcome has actually been restored.**

FCA operational-resilience guidance distinguishes recovery time objectives from wider business-service impact tolerances.

# Recovery is not binary.

The question is not only when a service returns, but at what capacity.

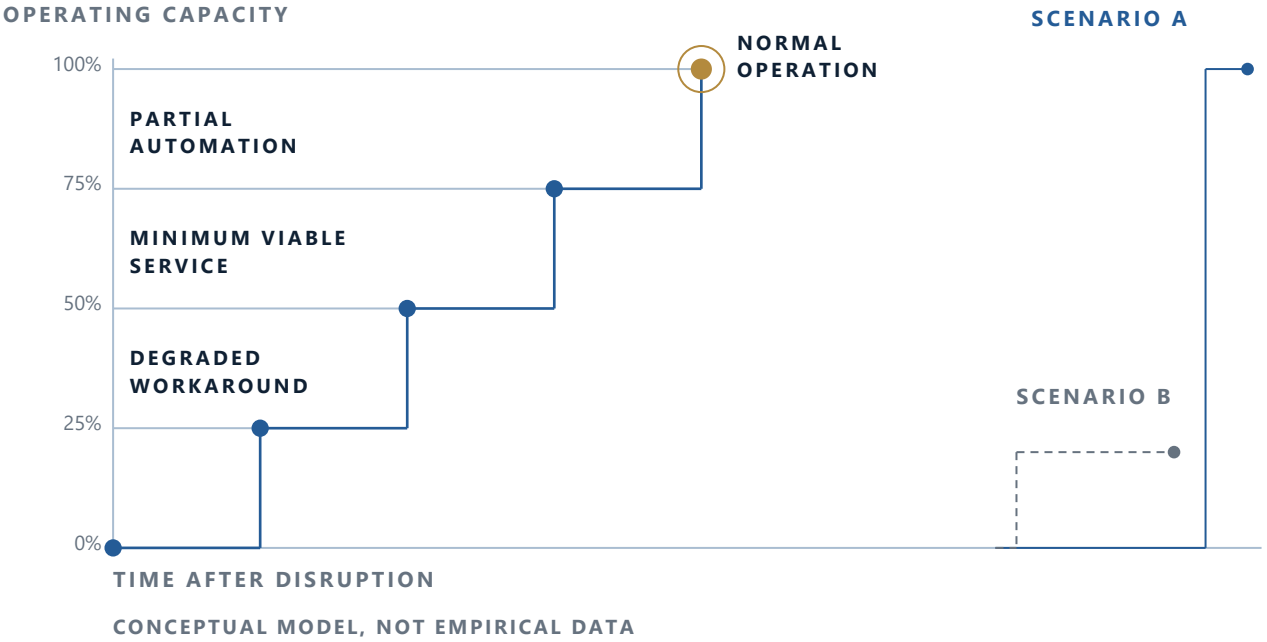


Service unavailable for 4 hours, then returns at 100%.

Service returns after 1 hour but operates at 20% capacity for 5 days.

## Which service actually recovered better?

The answer depends on the business service, demand, backlog, customer impact and defined impact tolerance.

The Clorox cyber incident illustrates why this matters: manual procedures allowed operations to continue, but at reduced capacity while automated systems were restored.

Source: The Clorox Company SEC disclosures, 2023–2024

**TIME TO RESTORE**
**+ CAPACITY RESTORED**
**+ TRUST**
**+ BUSINESS IMPACT**

**together provide a richer picture of recovery than binary availability alone.**

# 05

# A recovery plan is a claim. A recovery test is evidence.

Recovery architecture describes what an organisation has built. Recovery testing describes what an organisation has shown. Resilience assurance depends on the difference between the two.

**RECOVERY PLAN**

"We have documented how we would recover."

↓

**RECOVERY DESIGN**

"We have built the capability to recover."

↓

**RECOVERY TEST**

"We have shown that part of it works."

↓

**BUSINESS-SERVICE EVIDENCE**

"We have shown the service itself can be restored."

**The strength of a resilience claim should depend on the strength of the evidence behind it.**

# The Recovery Evidence Ladder

Not all evidence of recoverability is equal.

LEVEL 1

### ASSUMED

"We believe we could recover."

↑

LEVEL 2

### DECLARED

"We have a documented recovery plan."

↑

LEVEL 3

### DESIGNED

"We have built recovery capability."

↑

LEVEL 4

### COMPONENT TESTED

"We restored a system successfully."

↑

LEVEL 5

### SERVICE DEMONSTRATED

"We restored a critical business service end to end."

↑

LEVEL 6

### ADVERSARIALLY DEMONSTRATED

"We recovered under realistic severe-but-plausible conditions."

↑

LEVEL 7

### CONTINUOUSLY VALIDATED

"We can evidence recoverability repeatedly over time."

**Confidence in a recovery claim should increase only as the evidence supporting it increases.**

An analytical framework developed for this research. It rests on established recovery-testing and operational-resilience practices, but it is not an industry standard.

# The test should cross the same failure boundaries as the incident.

### COMPONENT RESTORE TEST

- Restore a server
- Restore a database
- Restore a file set
- Validate backup integrity

### SERVICE RECOVERY EXERCISE

- Recover identity and access
- Recover dependencies in sequence
- Validate data integrity and trust
- Recover at production scale
- Reconnect third parties
- Restore the business process
- Meet impact tolerance

## Both tests are useful. They answer different questions.

### COMPONENT TEST ASKS

“Did the restore work?”

### SERVICE EXERCISE ASKS

“Did the business service return within tolerance?”

### EVIDENCE CONTEXT

> The FCA has observed firms asserting broad recovery capability without evidence drawn from sufficiently severe testing.
>
> Source: Financial Conduct Authority, Operational Resilience Insights and Observations, 2026

> NIST guidance on cyber event recovery emphasises exercising and testing recovery capability rather than relying on documented plans alone.
>
> Source: NIST SP 800-184, Guide for Cybersecurity Event Recovery

# Confidence should follow evidence.

Justified confidence rises with demonstrated evidence, not with intent.

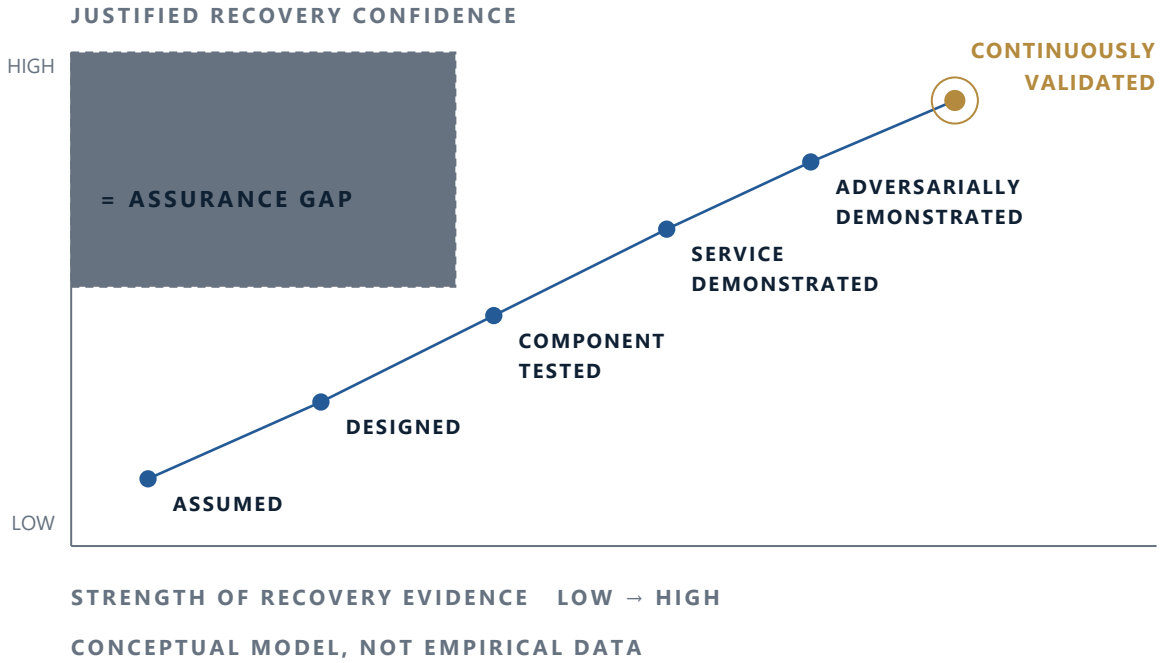


## Where confidence sits above evidence, the organisation is relying on assumption rather than demonstration.

This assurance gap is distinct from the Recoverability Gap described in Section 04. It concerns what has been proven, not what can be achieved.

Leading practice moves recovery assurance from documented intent toward demonstrated business-service recovery under realistic conditions.

Sources: Financial Conduct Authority, Operational Resilience Insights and Observations, 2026; NIST SP 800-184, Guide for Cybersecurity Event Recovery. Neither source uses this framework.

THE QUESTION FOR LEADERS IS NOT

"Do we have a recovery plan?"

IT IS

**"What have we actually demonstrated?"**

# 06

# What do you recover first when recovery itself depends on it?

Modern enterprise services depend on foundational systems that authenticate administrators, control infrastructure, resolve services, issue trust and coordinate recovery.


**The lower the dependency, the greater its potential recovery blast radius.**

# Attackers increasingly target the ability to recover.

| | |
|---|---|
| **IDENTITY** | Credential theft<br>Domain privilege<br>Authentication infrastructure |
| **CONTROL PLANES** | vCenter / ESXi<br>Cloud IAM<br>Management platforms |
| **RECOVERY INFRASTRUCTURE** | Backup administration<br>Stored credentials<br>Recovery repositories |

**~10%** of ransomware intrusions investigated by Mandiant in 2025 involved attackers targeting Veeam Backup & Replication to harvest stored credentials.

**This does not indicate a Veeam-specific recovery weakness. It demonstrates the value attackers place on privileged credentials accessible through recovery infrastructure.**

Mandiant's 2026 threat research describes ransomware actors increasingly targeting backup infrastructure, identity services and virtualization layers to deny recovery.

Sources: Google Cloud / Mandiant, M-Trends 2026; Mandiant, Ransomware TTPs: Shifting Threat Landscape, 2025/2026

**Where recovery infrastructure holds privileged access, it becomes part of the attack surface it is meant to protect.**

# Recovery dependencies can become circular.

↑ CIRCULAR RECOVERY DEPENDENCY

THE SAME PATTERN CAN EXIST ACROSS

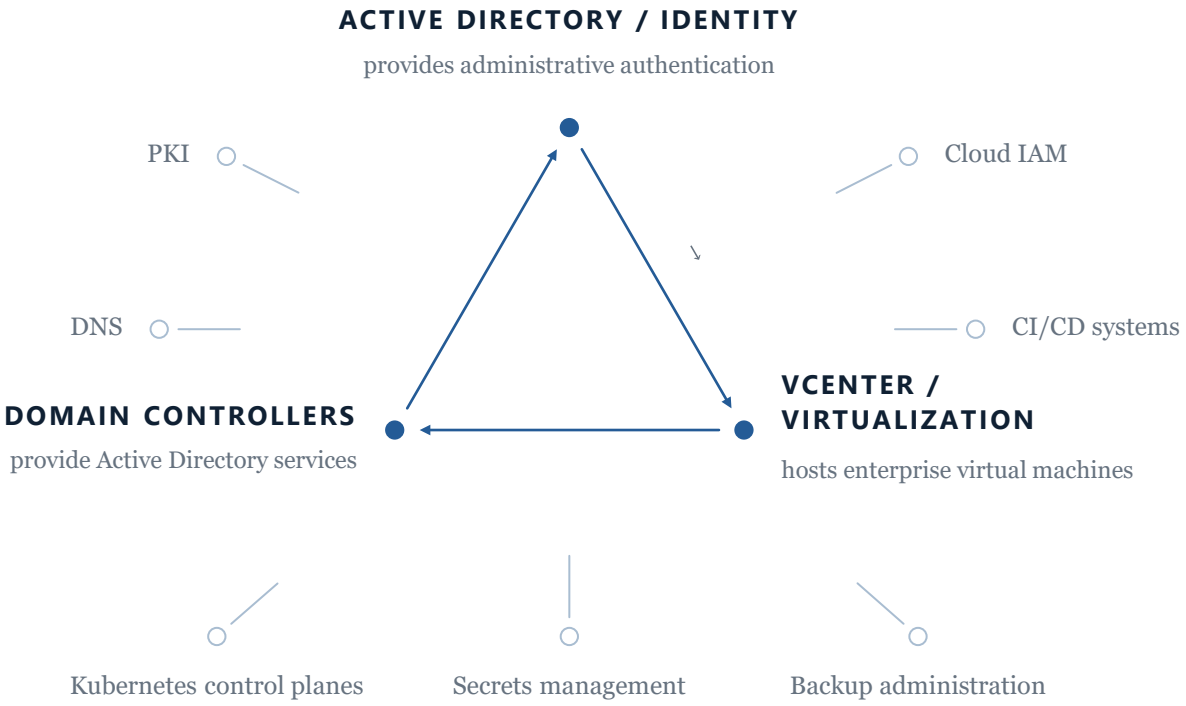


**If virtualized domain controllers depend on a compromised virtualization platform, while administration of that platform depends on Active Directory, recovery sequencing becomes significantly more complex.**

## Shared control planes can create shared recovery failure domains.

EVIDENCE

> Microsoft ransomware incident-response guidance recommends preserving known-good domain controllers and explicitly considers the dependency between virtual domain controllers and virtualization infrastructure.

> Mandiant has separately documented the risk created when compromised Active Directory credentials provide administrative access to vSphere.

Sources: Microsoft, Ransomware Incident Response Playbook; Google Cloud / Mandiant, vSphere Active Directory Integration Risks

# A clean backup needs somewhere trusted to run.

**TRUSTED DATA**

**+ TRUSTED IDENTITY**

**+ TRUSTED CONTROL PLANE**

**+ TRUSTED COMPUTE**

**+ CORRECT DEPENDENCY SEQUENCE**

**= VIABLE RECOVERY ENVIRONMENT**

## TRUSTED RECOVERY COMPUTE

Compute capacity and infrastructure that can safely host recovered workloads when the original production environment cannot yet be trusted.

- Isolated DR infrastructure
- Separate cloud account or subscription
- Alternate region
- Clean hypervisor hosts
- Cyber recovery environment

Not every enterprise requires all of these. The architectural requirement is appropriate separation of recovery failure domains based on service criticality and risk.

The independent review of Ireland's HSE cyberattack identified the need for validated recovery of infrastructure, including Active Directory, protected backups, mass-rebuild capability and prioritized application recovery.

Source: HSE, Conti Cyber Attack on the HSE, Independent Post Incident Review

**BACKUP SURVIVAL ANSWERS**

"Did the recovery assets survive?"

**TRUSTED RECONSTRUCTION ANSWERS**

"Can we safely build the service again?"

**SURVIVE** **TRUST** **RECONSTRUCT**

# 07

# Cloud did not remove recovery risk. It changed its boundaries.

Modern platforms provide powerful native resilience and recovery capabilities. But recovery increasingly depends on understanding where the failure domain ends, and whether recovery assets, identities, capacity and dependencies exist outside it.

### PRODUCTION FAILURE DOMAIN

- Account / subscription / project
- Applications
- Control plane
- Identity
- Data

↓ CROSS THE FAILURE BOUNDARY

### TRUSTED RECOVERY DOMAIN

- Protected recovery assets
- Viable recovery environment
- Independent administrative access
- Required capacity

## Geographic separation is not always administrative separation.

Diversification can reduce concentration exposure but introduces cost, complexity and operational overhead. The appropriate response depends on service criticality and impact tolerance.

# Cloud-native recovery is becoming cyber-aware.

Three platforms, independently moving in the same architectural direction.

| AWS | AZURE | GOOGLE CLOUD |
|---|---|---|
| Logically air-gapped vaults | Immutable backup vaults | Backup Vault |
| Compliance-mode Vault Lock | Locked immutability | Immutable / indelible protection |
| Cross-account recovery | Cross-subscription restore | Project-level identity isolation |
| Multiparty approval options | Isolated recovery boundary | Isolated recovery environments |

## THE COMMON ARCHITECTURAL PATTERN

**PROTECT** — **ISOLATE** — **RESTORE INTO A TRUSTED DOMAIN**

## The major hyperscalers are independently implementing recovery patterns that separate protected recovery assets from the production failure domain.

This indicates that isolation and alternate administrative recovery boundaries are becoming mainstream cloud-recovery architecture.

These services are not equivalent and are not ranked here. Their architectures differ. The purpose is to identify a shared direction of travel.

### SOURCES

AWS Backup documentation, Logically Air-Gapped Vault
Microsoft Azure Backup documentation, Data Protection and Immutable Vault guidance
Google Cloud Backup and DR documentation, Backup Vault

# The SaaS question is no longer: “Does the provider have backups?”

OLD QUESTION

- Does Microsoft or the SaaS provider back up my data?

BETTER QUESTION

- **Do the available recovery capabilities cover the scenarios my business actually requires?**

THE DIMENSIONS THAT DECIDE IT

- POINT-IN-TIME RECOVERY
- RESTORE GRANULARITY
- ADMINISTRATIVE INDEPENDENCE
- COMPLIANCE / LEGAL REQUIREMENTS
- RETENTION
- RECOVERY SPEED
- CROSS-SERVICE RECOVERY
- PROVIDER FAILURE ASSUMPTIONS

**10-MINUTE RPO** for recent restore points in supported Microsoft 365 workloads.

Microsoft 365 Backup now provides dedicated first-party backup and high-speed restore capabilities for supported Exchange, SharePoint and OneDrive scenarios.

Microsoft also explicitly distinguishes disaster-recovery copies, which preserve current service state, from backup, which enables return to a previous healthy state.

Source: Microsoft Learn, Microsoft 365 Backup Overview and FAQ

**Native SaaS protection may be entirely sufficient for some organisations and scenarios.**
Others may justify independent protection for different retention, administrative separation, cross-SaaS consistency, governance or portability requirements.

# Some critical dependencies are outside your recovery control.

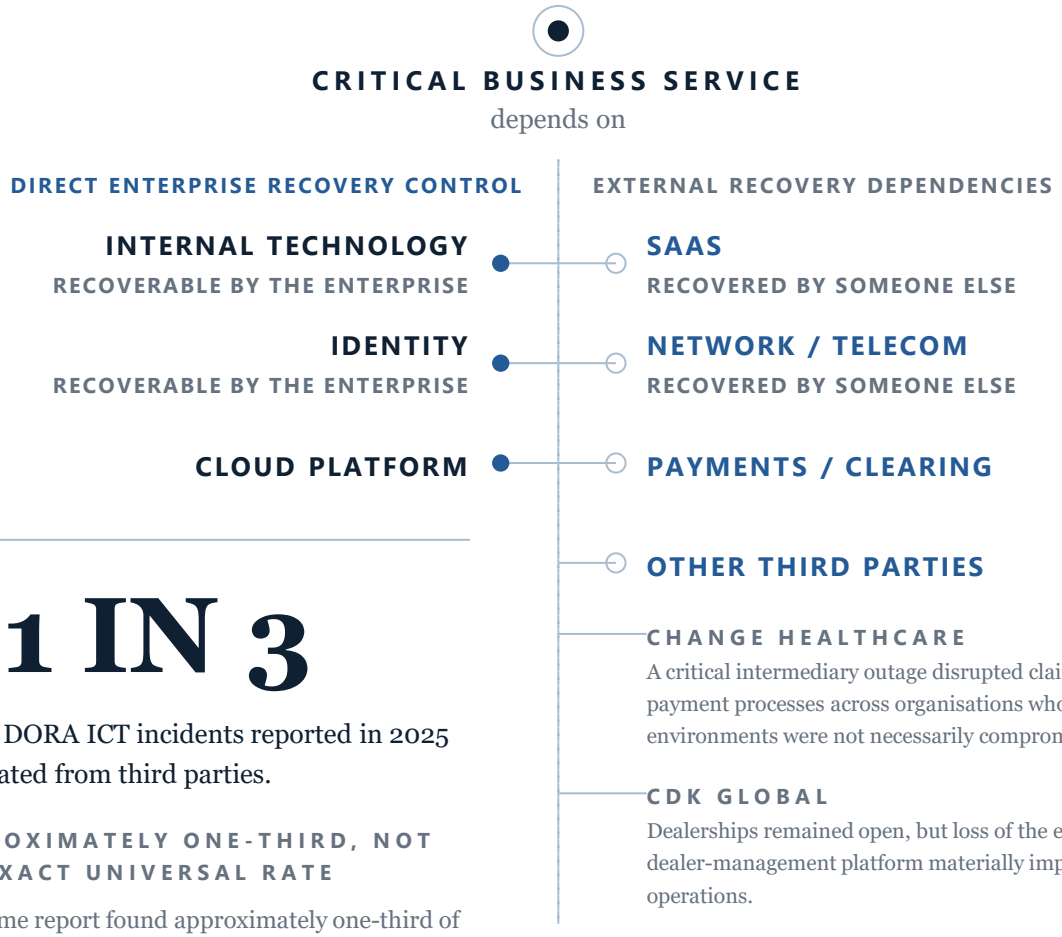


## ~1 IN 3

major DORA ICT incidents reported in 2025 originated from third parties.

**APPROXIMATELY ONE-THIRD, NOT AN EXACT UNIVERSAL RATE**

The same report found approximately one-third of major incidents had cross-border impact.

Source: European Supervisory Authorities, Joint Report on Major ICT-Related Incidents in 2025

**Modern recoverability is increasingly ecosystem-dependent. An organisation may possess excellent internal recovery capability while remaining unable to independently restore a critical external dependency.**

**THE ENTERPRISE CAN CONTROL**

Architecture · Workarounds · Alternative pathways · Contracts · Testing · Concentration exposure

**BUT NOT**

The recovery execution of every external provider.

**This makes dependency mapping and severe third-party scenario testing part of recoverability, not merely procurement risk.**

# 08

# Design recovery from the business service backward.

The recovery architecture should be determined by what the business must restore, how quickly it must operate, which dependencies are required, and which failure boundaries the recovery path must survive.


## Business requirements design downward. Recovery executes upward.

The design logic moves down from business requirements. The recovery path moves up from trusted recovery assets toward restored business operation.

# A target architecture for demonstrated recoverability


↓ CONTROLLED RECOVERY PATH ACROSS THE FAILURE BOUNDARY

**TRUSTED RECOVERY DOMAIN**

- Trusted identity
- Trusted compute
- Isolated administration
- Trusted control plane
- Protected recovery assets

A separate failure boundary, not a copy.

**OUTSIDE DIRECT ENTERPRISE RECOVERY CONTROL**

- Cloud
- SaaS
- Telecom
- Payments
- Critical providers

**Recovery is designed from the business outcome downward, and executed from protected recovery assets upward into a trusted domain.**

The Recoverability Stack maps onto these layers, from copies that survive at the protected layer to operation proven at the business outcome.

Analytical target architecture developed for this research. Not a regulatory requirement, and no guarantee of recovery. Implementation depends on criticality, impact tolerance, threat model, regulation and cost.

# The recovery environment should not inherit the same failure.

### 01 IDENTITY COMPROMISE

**RISK** Production identity plane is compromised.

**WEAK PATTERN** Recovery administration depends entirely on the same identity plane.

**PRINCIPLE** **Maintain an appropriately isolated path to administer and reconstruct recovery.**

### 02 CLOUD ACCOUNT / SUBSCRIPTION COMPROMISE

**RISK** Production administrative boundary cannot be trusted.

**WEAK PATTERN** All recovery assets remain accessible only through the compromised boundary.

**PRINCIPLE** **Where justified, enable protected recovery assets and restoration into a separate trusted boundary.**

### 03 PRODUCTION COMPUTE COMPROMISED OR UNAVAILABLE

**RISK** Clean recovery points exist, but no trusted execution environment does.

**WEAK PATTERN** Assume the original production environment will always be available for restoration.

**PRINCIPLE** **Define where recovered workloads can safely execute, and whether recovery capacity exists.**

## Recovery independence should match the failure scenario.

Independence does not necessarily require separate physical infrastructure. Depending on architecture and risk, separation may be:

- Administrative
- Logical
- Account-based
- Subscription-based
- Project-based
- Regional
- Physical

Greater isolation can improve resilience but also increases cost and operational complexity. The appropriate level should be driven by business-service criticality and impact tolerance.

# Recoverability is a system of capabilities.

Not a product to purchase, but a continuous operating model.


↑ CONTINUOUS, NOT A ONE-TIME PROGRAMME

**THE RECOVERABILITY STACK**

**COPY → SURVIVE → TRUST → RECONSTRUCT → ORCHESTRATE → OPERATE**

**THE RECOVERY EVIDENCE LADDER**

**ASSUMED → DECLARED → DESIGNED → COMPONENT TESTED → SERVICE DEMONSTRATED → ADVERSARIALLY DEMONSTRATED → CONTINUOUSLY VALIDATED**

THE STACK DESCRIBES THE CONDITIONS. THE LADDER DESCRIBES THE EVIDENCE.

The objective is not perfect redundancy.

## The objective is evidenced recovery of the required business outcome within its tolerance.

Architecture, protection, orchestration and testing therefore operate as one recoverability system.

# AI adds state, scale and scarce capacity to the recovery equation.

An AI service depends on more than application code and data. Recoverability can extend across models, training state, vector stores, orchestration and specialized infrastructure.

**AI-ENABLED BUSINESS SERVICE**

depends on

- APPLICATION / API
- MODEL
- MODEL WEIGHTS / CHECKPOINTS
- DATA / VECTOR STORES
- ORCHESTRATION + CONFIGURATION
- IDENTITY / SECRETS / CONTROL PLANE
- GPU / ACCELERATOR CAPACITY
- NETWORK + STORAGE

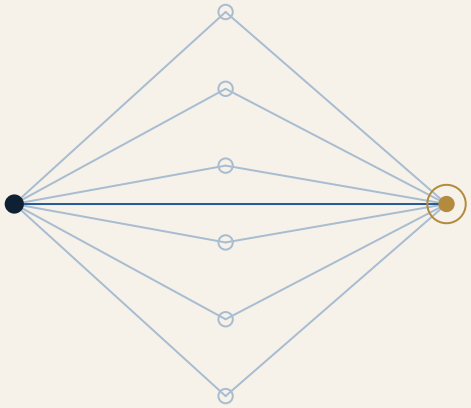

## THREE EMERGING RECOVERY QUESTIONS

| | |
|---|---|
| **STATE** | What must be preserved to resume rather than restart? |
| **TRUST** | Can the model, data, checkpoint and configuration be validated as known-good? |
| **CAPACITY** | Is sufficient accelerator, storage and network capacity available to restore the service at the required operating level? |

## The Recoverability Stack still applies.

**COPY → SURVIVE → TRUST → RECONSTRUCT → ORCHESTRATE → OPERATE**

What changes is the composition, scale and scarcity of some dependencies.

AI architectures vary significantly. Not every AI workload requires recovery of every component shown.

# A recovery point is useful only if there is somewhere viable to run it.

**RECOVERABLE STATE**

**+ TRUSTED ENVIRONMENT**

**+ AVAILABLE COMPUTE CAPACITY**

**= VIABLE AI SERVICE RECOVERY**

### TRADITIONAL STATELESS SERVICE

Application can often be recreated from code, configuration and conventional infrastructure capacity.

### COMPUTE-INTENSIVE AI SERVICE

Recovery may additionally depend on

- Large model state
- Training checkpoints
- Distributed workload state
- GPU / accelerator availability
- High-performance storage
- Network topology
- Data and model provenance

### OBSERVED CURRENT EVIDENCE

Kubernetes established a dedicated Checkpoint/Restore Working Group in 2026. Its stated use cases include:

- Long-running distributed model training
- Interruption-aware scheduling
- Forensic checkpointing after incidents
- Workload migration
- AI inference services

Kubernetes storage development is also advancing group-consistent volume snapshots and changed-block tracking, while state management, data mobility and distributed recovery remain active engineering challenges.

Sources: Kubernetes, Introducing the Checkpoint/Restore Working Group, January 2026; Kubernetes SIG Storage, 2026

### EMERGING RECOVERY IMPLICATIONS

The issue is not that AI invalidates existing recovery practice. It is that some AI services may make trusted state reconstruction and recovery capacity materially more expensive, scarce or complex.

## AI does not replace the recoverability problem. It extends it.

# 10

# Shift the question from “Are we protected?” to “What can we prove we can recover?”

The next step in recovery maturity is stronger alignment between business-service requirements, recovery architecture, dependency understanding and demonstrated recovery evidence.


## SIX PRIORITIES

01 Start with critical business services

02 Map the real recovery dependency chain

03 Engineer recovery across failure boundaries

04 Treat identity and control planes as recovery-critical

05 Test business recovery, not only component restoration

06 Measure and improve demonstrated recoverability

# Design the recovery architecture from the business backward.

## 01 / START WITH THE BUSINESS SERVICE

Critical business service
Impact tolerance
Minimum viable operating capacity
Acceptable data loss
Required recovery sequence

**QUESTION**

What business outcome must be restored, and at what operating level?

## 02 / MAP THE REAL RECOVERY DEPENDENCY CHAIN

Applications
Data
Identity
DNS / PKI
Network
Virtualization / cloud
Secrets
SaaS
Third parties
Recovery capacity

**QUESTION**

Which dependency can prevent the service from operating even when its application and data are restored?

## 03 / ENGINEER RECOVERY ACROSS THE FAILURE BOUNDARY

Identity compromise
Administrative compromise
Production infrastructure loss
Regional disruption
Recovery-system attack
Critical-provider outage

**PRINCIPLE**

Recovery independence should match the failure scenario and service criticality.

Greater isolation and diversification introduce cost and complexity. They should be justified by business impact.

# Turn recovery capability into recovery evidence.

## 04 / TREAT IDENTITY AND CONTROL PLANES AS RECOVERY-CRITICAL

Identity
Cloud IAM
Virtualization management
DNS
PKI
Secrets
Backup administration
Recovery compute

**QUESTION**

Can we securely administer recovery if the normal control plane cannot be trusted?

## 05 / TEST THE BUSINESS SERVICE, NOT ONLY THE COMPONENT

Compromised identity
Untrusted production
Dependency failure
Third-party outage
Degraded operating capacity
Alternative recovery environment

**PROGRESSION**

From “Can we restore the VM?” toward “Can the business operate?” Measure time, trust, capacity, dependency restoration and business impact.

## 06 / MEASURE AND IMPROVE DEMONSTRATED RECOVERABILITY

Declared objectives
Demonstrated performance
Infrastructure recovery time
Application recovery time
Minimum viable service time
Full operational recovery time
Trusted recovery time

**NOTE**

These analytical distinctions supplement rather than replace established RPO/RTO measures.

# Six questions every technology leadership team should be able to answer.

01 Which critical business services are we actually protecting?

02 What identities, control planes, infrastructure, data and third parties must recover before each service can operate?

03 Which failure scenarios could compromise both production and our recovery path?

04 Where will workloads run if the normal environment cannot be trusted or lacks sufficient recovery capacity?

05 What is the strongest evidence that each critical service can recover within its impact tolerance?

06 Where does recovery confidence currently exceed demonstrated evidence?

## IF THE INCIDENT HAPPENED TOMORROW, WHAT HAVE WE ACTUALLY PROVED?

CONCLUSION

# Backup is getting better. The recovery question is getting bigger.

01 **PROTECTION HAS ADVANCED**

Immutability, isolation, cloud-native protection and cyber-recovery capabilities have materially strengthened the enterprise recovery toolkit.

02 **SURVIVAL IS NOT THE SAME AS RECOVERY**

Protected recovery assets still need to be trusted, reconstructed into a viable environment and restored through the required dependency chain.

03 **THE BUSINESS SERVICE IS THE ULTIMATE RECOVERY BOUNDARY**

Technical restoration can precede meaningful business recovery by hours or days. Recovery therefore needs to consider trust, dependencies and operating capacity, not only system availability.

04 **RECOVERABILITY IS INCREASINGLY ECOSYSTEM-DEPENDENT**

Identity, control planes, SaaS, cloud platforms and critical third parties can determine recovery outcomes beyond the primary application or backup environment.

05 **EVIDENCE MATTERS**

Documented objectives and recovery architecture provide important assurance. Realistic testing provides stronger evidence of what the organisation can actually recover.

AN EXPANDING RECOVERY SCOPE


AN EXPANDING SCOPE OF RECOVERY, NOT A SEQUENCE OF REPLACEMENT TECHNOLOGIES

The evidence does not show that enterprise recovery is broadly failing.

**It shows that the definition of successful recovery is becoming more demanding.**

CONCLUSION

THE STATE OF ENTERPRISE RECOVERABILITY 2026

# Enterprise backup and recovery capabilities have advanced materially.

The remaining challenge is increasingly one of demonstrated recoverability:

proving that trusted critical business services can be restored, at sufficient capacity, within acceptable disruption tolerances across complex identity, technology and third-party dependencies.


The objective is not simply to possess recoverable data.

**It is to demonstrate a recoverable business.**

THE QUESTION IS NO LONGER ONLY

"Do we have backups?"

OR EVEN

"Can we restore?"

IT IS

# "What have we proved we can recover?"

Karim Gaber Saleh
Independent Research
September 2026

# Sources & research notes

This research triangulated regulatory guidance, primary corporate disclosures, incident-response research, technical documentation, industry research and enterprise surveys.

## EVIDENCE HIERARCHY

A **TIER A / PRIMARY & REGULATORY**

Regulators · Government guidance · Official standards guidance · SEC and corporate disclosures · Primary platform documentation

B **TIER B / INDEPENDENT & FRONTLINE RESEARCH**

Incident-response research · Industry foundations · Independent professional research

C **TIER C / VENDOR TECHNICAL EVIDENCE**

Used primarily to establish current product or architectural capabilities.

D **TIER D / VENDOR & COMMISSIONED SURVEYS**

Used directionally and with methodological qualification.

## RESEARCH NOTES

**Major conclusions were triangulated across multiple evidence types wherever possible.**

Survey percentages were not treated as globally representative unless the underlying methodology supported such an interpretation.

The Recoverability Stack, Recovery Evidence Ladder and Recoverability Gap are analytical constructs developed for this research. They synthesize established concepts from backup, disaster recovery, cyber recovery, business continuity and operational resilience.

No organisation cited in this publication endorses these frameworks. They are the author's analytical constructs and are presented as such.

# Primary resilience and regulatory sources

TIER A / PRIMARY & REGULATORY EVIDENCE

### FINANCIAL CONDUCT AUTHORITY

Operational Resilience guidance
Operational Resilience: Insights and Observations, 2026

**USED FOR**

Important business services · Impact tolerances · Scenario testing · Third-party dependencies · Evidence of recoverability

### BANK OF ENGLAND / PRA

Operational Resilience framework
Financial Market Infrastructure Annual Report 2025–26
Financial Stability Report, July 2026

**USED FOR**

Important business services · Severe disruption · Third-party concentration · Operational resilience maturity

### EUROPEAN SUPERVISORY AUTHORITIES

Joint Report on Major ICT-Related Incidents in 2025

**USED FOR**

Major DORA ICT incidents reported in 2025 · Third-party incident origination · Cross-border impact

### NIST

SP 800-34 Rev.1, Contingency Planning Guide for Federal Information Systems
SP 800-184, Guide for Cybersecurity Event Recovery
NCCoE Data Integrity guidance, Recovering from Ransomware and Other Destructive Events

**USED FOR**

RTO / RPO · Recovery strategy · Testing · Known-good recovery · Cyber recovery principles

### CISA

Ransomware and VMware ESXi-related cybersecurity guidance

**USED FOR**

Virtualization-layer attack evidence and defensive context

Regulatory requirements vary by jurisdiction and sector. References are used here to identify resilience principles and observed practices, not to imply universal applicability.

# Incident and threat evidence

PRIMARY CORPORATE DISCLOSURES · INDEPENDENT AND FRONTLINE RESEARCH

### CHANGE HEALTHCARE / 2024

UnitedHealth Group corporate and SEC disclosures
Healthcare-provider SEC disclosures

**USED FOR**

Ecosystem dependency · Claims and payment disruption · Customer reconnection · Provider liquidity effects · Recovery progression

### CDK GLOBAL ECOSYSTEM / 2024

Asbury Automotive Group disclosures
Sonic Automotive disclosures
Group 1 Automotive disclosures
Lithia Motors disclosures

**USED FOR**

Critical SaaS dependency · Manual workarounds · Recovery duration · Integration dependencies · Concentration exposure · Sonic Automotive's estimated $47.2M negative pre-tax impact

### MGM RESORTS / 2023

MGM Resorts International SEC disclosures

**USED FOR**

Containment-driven shutdown · Operational disruption · Approx. $100M negative Adjusted Property EBITDAR impact · Incident expenses

### HSE / IRELAND

Conti Cyber Attack on the HSE
Independent Post Incident Review

**USED FOR**

Active Directory recovery · Protected backups · Mass rebuilding · Prioritized application recovery

### THE CLOROX COMPANY / 2023

Corporate and SEC disclosures

**USED FOR**

Manual workarounds · Reduced operating capacity · 23–28% expected Q1 year-on-year sales decline during disruption · Approx. $49M incremental cyberattack-related costs through December 2023

### MANDIANT / GOOGLE CLOUD

M-Trends 2026
Ransomware TTPs: Shifting Threat Landscape
vSphere Active Directory Integration Risks

**USED FOR**

Recovery-denial tactics · Identity compromise · Virtualization risk · Backup-infrastructure targeting · Control-plane dependencies

Figures are taken from the organisations' own disclosures and post-incident reviews.

# Market, adoption and survey evidence

### BACKUP / RECOVERY MARKET

Gartner, Backup and Data Protection Platforms, 2026

**USED FOR**

Current category scope · Cyberattack detection · Immutable storage · Hybrid, multicloud and SaaS protection · Recovery capabilities

## TIER D / SURVEY AND DIRECTIONAL EVIDENCE

### DATABARRACKS

Data Health Check / State of Resilience 2026
500 UK IT and resilience professionals

**USED FOR**

Directional. Air-gapped and immutable backup adoption

### SOPHOS

State of Ransomware 2026
2,158 respondents across 17 countries

**USED FOR**

Backup-based ransomware recovery · Recovery duration and cost context

### QUEST SOFTWARE

Identity Disaster Recovery Testing Survey, 2026
650 global IT and security professionals

**USED FOR**

Directional. Identity recovery testing practices

### KPMG

Global Third-Party Risk Management Survey, 2026
851 organizations

**USED FOR**

Third-party risk governance maturity

### ARCSERVE

State of Data Resilience 2026
200+ IT respondents

**USED FOR**

Directional. Recovery confidence versus full-test performance · Full recovery testing

## SURVEY CAUTION

**Samples, definitions, geographies and methodologies differ.**

Figures from different surveys should not be combined into a single global adoption rate.

# Platform evidence and analytical definitions

## TIER C / VENDOR AND PLATFORM TECHNICAL DOCUMENTATION

### AWS BACKUP

Logically Air-Gapped Vault · Vault Lock · Cross-account recovery · Multiparty approval

### AZURE BACKUP

Immutable vaults · Cross-subscription restore · Data-protection guidance

### GOOGLE CLOUD BACKUP AND DR

Backup Vault · Identity isolation · Isolated recovery environments

### MICROSOFT 365 BACKUP

Microsoft 365 Backup documentation

**USED FOR**

First-party backup · Supported 10-minute recent restore-point RPO · High-speed restore · Distinction between disaster-recovery copies and historical backup recovery

### KUBERNETES

Introducing the Checkpoint/Restore Working Group, 2026 · SIG Storage Spotlight, 2026 · Self-Healing documentation

**USED FOR**

Checkpoint / restore · Stateful workload recovery · Group-consistent snapshots · Data mobility · Emerging AI recovery considerations

## ANALYTICAL DEFINITIONS

### RECOVERABILITY STACK

**COPY → SURVIVE → TRUST → RECONSTRUCT → ORCHESTRATE → OPERATE**

Analytical framework developed for this research.

### RECOVERY EVIDENCE LADDER

**ASSUMED → DECLARED → DESIGNED → COMPONENT TESTED → SERVICE DEMONSTRATED → ADVERSARIALLY DEMONSTRATED → CONTINUOUSLY VALIDATED**

Analytical framework developed for this research.

### RECOVERABILITY GAP

The gap between stated or technically demonstrated recovery capability and the evidenced ability to restore a trusted critical business service to an acceptable operating level within its impact tolerance.

Research construct developed for this publication.

## RESEARCH CUT-OFF

August 2026. The publication reflects sources and evidence reviewed up to the research cut-off.