

A FIELD MANUAL

Getting Your Data Back on Your Own Terms

Backup architecture, restore testing and rehearsed recovery, the discipline that decides whether ransomware is a catastrophe or a bad Tuesday.

Bhavin G Patel

C|CISO · CISSP-ISSAP · CISM

cybersecleader.com

Ten runbooks · Eight templates · Built to be adapted, not admired

About this manual

Ransomware isn't a malware problem. It is a problem of not being able to recover, and that reframe is the whole reason this manual exists. The attacker's entire business model is the gap between "pay us" and "rebuild from nothing." Close that gap and the ransom becomes a rude email you delete.

Almost every organisation already knows it should have backups. Very few can say, without flinching, how long recovery would take and whether they are certain of the number. This manual is about turning that from a hope into a measured, rehearsed, signed-off answer.

Who it is for

The person who would be held responsible if the answer to "how long until we are running again?" turned out to be wrong. A CISO, a head of infrastructure, a resilience lead, or the engineer who owns backups and suspects nobody has tested them.

How to use it

If you have limited time, do three things in this order: run **RR-01** to find out what is actually in scope, run **RR-02** to discover whether a restore works and how long it takes, and run **RR-03** to confirm your last-resort copy is genuinely beyond an attacker's reach. Most organisations get an unpleasant surprise from at least one of those three, and finding it on a quiet Tuesday is the entire point.

A DELIBERATE OMISSION

This manual is about recovery, not prevention. Endpoint detection, email filtering and privileged access management all matter and are covered well elsewhere. The argument here is that every dollar spent guaranteeing nothing ever gets in is a bet you eventually lose, while the spend that makes getting in not matter pays off regardless. Treat this as the half of the problem that is usually under-resourced.

© 2026 Bhavin G Patel. All rights reserved.

First published 2026 at cybersecleader.com. *Getting Your Data Back on Your Own Terms: A Field Manual for Ransomware Recovery*

You may reproduce and adapt this material inside your own organisation, including deriving your own policies, standards, runbooks and templates from it, provided attribution to the author is retained in any derived document that is circulated beyond your immediate team. You may not sell it, bundle it into a commercial or consulting deliverable offered to third parties, republish it in whole or in substantial part, or present it as your own work. For any use beyond that, ask me first.

All trademarks, product names and regulatory framework names referenced belong to their respective owners and are used here for identification only. No affiliation or endorsement is implied.

Disclaimer

This manual is general information, not legal, regulatory or incident-response advice, and reading it does not create an advisory relationship between us. Decisions taken during a live ransomware incident, particularly around payment, notification and law enforcement, carry legal and regulatory consequences that vary by jurisdiction and by the sanctions status of the actor involved. Engage your own counsel early. I accept no liability for decisions taken on the basis of this document. Verify before you rely.

Contents

PART I · THE POSITION	6
1 Ransomware Is a Recovery Problem	7
2 What the Attacker Is Actually Selling	8
3 Backups That Survive an Attacker With Domain Admin	9
4 Immutability, Air Gaps and What “Offline” Really Means	11
5 Segmentation and Blast Radius	12
PART II · PROVING YOU CAN RECOVER	13
6 A Backup You Have Never Restored Is a Hypothesis	14
7 Restore Testing: Scope, Cadence, Stopwatch	15
8 The Number You Must Be Able to Say Out Loud	16
9 The Isolated Recovery Environment	17
PART III · THE RECOVERY ITSELF	18
10 Recovery Is a Plan, Not a Button	19
11 Restoring Without Restoring the Attacker	20
12 The Payment Decision	21
13 Notification, Law Enforcement and the Clock	22
PART IV · RUNBOOKS	23
RR-01 Backup Scope and Coverage Review	24
RR-02 Restore Test	25
RR-03 Immutability and Air-Gap Verification	26
RR-04 Backup Infrastructure Hardening Review	27
RR-05 Ransomware Incident: The First Sixty Minutes	28
RR-06 Building the Clean Room	29
RR-07 Recovery Sequencing and Execution	30
RR-08 Reinfection Prevention During Restore	31
RR-09 Ransom Decision Governance	32

RR-10 Full Recovery Rehearsal	33
PART V · TEMPLATES AND REFERENCE	34
A Recovery Tier Definitions	35
B Backup Coverage Register	36
C Restore Test Record	37
D Recovery Sequencing Plan	38
E Tabletop Scenario Pack	39
F Board Reporting: Recovery Readiness	40
G Where This Gets Hard	41
H Glossary	42

PART ONE

The Position

1 · Ransomware Is a Recovery Problem

The thing that decides whether ransomware is a catastrophe or a bad Tuesday is almost always the oldest, most boring discipline available: can you get your data back on your own terms.

When people ask how to defend against ransomware they usually want to talk about detection and endpoint tooling. All of it is useful. None of it is the control that decides the outcome, because given enough attempts something eventually gets in, and the organisations that come through intact aren't the ones that were never breached. They are the ones that could shrug, restore, and move on.

This isn't defeatism. It is a statement about where leverage actually sits. Prevention spend is a bet that has to win every time. Recovery spend pays off regardless of whether prevention held, which makes it the only part of the budget with an unconditional return.

The consequence of the reframe

Once you accept that recovery is the deciding control, several things that looked like infrastructure hygiene become security controls of the first order, and they change owner accordingly. Backup scope is a security question. Restore duration is a security metric. The credentials on your backup system are among the most sensitive in the organisation. And the recovery plan is an incident-response artefact, not an IT operations document.

That reassignment is usually the hardest organisational part of this. Backups typically live with infrastructure, are funded as operations, and are reported on — if at all — as job success rates. None of that survives contact with a ransomware incident, because job success isn't restore success and nobody in the room will care how many jobs completed.

THE QUESTION THIS MANUAL EXISTS TO ANSWER


If everything were encrypted tonight, how long until we are running again, and are we certain of that number? If you dislike the answer, or you are not certain of it. That's your next project, and it is worth more than the best tool on the show floor.

2 · What the Attacker Is Actually Selling

Ransomware works because the alternative to paying is worse. That gap — not the encryption — is the entire product.

Understanding the business model tells you exactly where to intervene. The attacker isn't selling decryption. They are selling the difference between two recovery paths: theirs, which is fast and uncertain, and yours, which they are betting is slow and untested. Everything in this manual is an attempt to make your path the obviously better one, at which point their product has no value.

THE LEVERAGE THEY ARE SELLING



Each of their three moves has a counter. Point three is the one most recovery plans ignore entirely.

Double extortion changes one thing, not everything

Exfiltration before encryption means restoring cleanly no longer guarantees the incident ends, they still hold the data and can threaten publication. This is real and it has shifted the calculus. It is also the reason some organisations conclude that recovery capability no longer matters. That conclusion is wrong.

Recovery capability removes the operational catastrophe: the outage, the revenue loss, the inability to serve customers. What remains is a data-breach problem, which is serious but is a category your organisation already has a process for. Paying to prevent publication buys a criminal's promise, which is worth what you would expect. Recovery converts an existential event into a bad and expensive one, and that's the whole objective.

3 · Backups That Survive an Attacker With Domain Admin

Modern crews are not smash-and-grab. They get in, look around, and one of the first things they hunt is your backup infrastructure — because they know it is the only thing standing between you and ignoring them.

This is the single most important design constraint and the one most backup estates fail. The question isn't "do we have backups?" It is "**which copy survives an attacker who already has domain administrator and a fortnight to look around?**" If the honest answer is none, you have an operational backup system and no ransomware control, and those are different things that get confused constantly.

3-2-1, and what each number is actually for

Element	What it protects against
3 copies	Ordinary loss — corruption, deletion, hardware failure. The original discipline, and not the part that saves you here.
2 media types	A defect or vulnerability affecting one technology. Increasingly interpreted as two distinct platforms rather than tape versus disk.
1 offsite	Site loss. Necessary, and by itself insufficient: an offsite copy reachable over the network is reachable by the attacker.
+1 immutable or offline	The ransomware addition. A copy that cannot be altered or deleted, by anyone, including someone holding your most privileged credentials.
+0 verified errors	The other modern addition: a restore that has been performed and validated, not a job that reported success.

Fate sharing is the failure

The design principle underneath all of it: **a backup that shares fate with production isn't a safety net. It's part of the thing being held hostage.** Fate sharing takes several forms and organisations usually have more than one.

Credential fate sharing is the most common and the most dangerous. If your backup system authenticates against the same directory as everything else, then domain admin is backup admin, and the attacker acquired both at the same moment. Backup infrastructure needs its own identity plane — separate accounts, separate MFA, and no trust relationship that lets production identity assume backup privilege.

Network fate sharing means the backup repository is reachable from the general estate. If a workstation can route to it, so can whatever is running on that workstation.

Platform fate sharing means backups live on the same hypervisor, storage array or cloud tenancy as production. A tenancy-level compromise takes both.

Vendor fate sharing means your backup product and your production platform come from the same supplier, so one vulnerability class can affect both. This one is a judgement call rather than a rule, and it trades against the operational benefit of an integrated stack.

4 · Immutability, Air Gaps and What “Offline” Really Means

“Immutable” is a word vendors use loosely and buyers accept uncritically. The specific question is: who could delete this, and by what mechanism?

STRENGTH OF PROTECTION — ASK WHO CAN STILL DELETE IT

SOFT DELETE / RECYCLE BIN

An administrator can purge it. Protects against accident, not an adversary.

RETENTION LOCK, ADMIN-OVERRIDABLE

The backup admin can shorten retention. Attacker with that role can too.

COMPLIANCE-MODE OBJECT LOCK

Nobody can shorten it — not the admin, not the vendor, not the account root.

PHYSICALLY DISCONNECTED MEDIA

Tape on a shelf. Unfashionable, slow to restore, and genuinely unreachable.

If you cannot name the mechanism that would stop a domain admin deleting a copy, you do not have one.

Most estates believe they are in the bottom two bands and are actually in the top two. RR-03 is how you find out which.

The questions that separate real immutability from marketing

Can the retention period be shortened by anyone, and by whom? Governance-mode locks can be overridden by a privileged user. Compliance-mode locks can’t be overridden by anyone, including the account owner and the vendor. That distinction is the entire control, and it is the one most often glossed over during procurement.

What happens if the account or subscription is deleted? A cloud immutable copy inside a tenancy the attacker controls may be immune to object deletion and still lost with the tenancy. Immutability at the object level doesn’t survive destruction at the account level.

Who holds the encryption keys, and are they in the blast radius? An immutable copy you can’t decrypt isn’t a recovery capability. If keys live only in the system being encrypted, you have built an elaborate way to lose data slowly.

What is the restore path when the primary environment no longer exists? Some immutable stores can only be read back through management infrastructure that’s itself part of production. Chapter 9 covers the environment you restore into.

5 · Segmentation and Blast Radius

Segmentation does not prevent ransomware. It determines how much of the organisation is in scope when it happens, and therefore how long recovery takes.

The recovery argument for segmentation is more concrete than the prevention one. Every additional system in the encryption blast radius is a system that must be restored, validated and returned to service before the business is whole. A firm that loses one business unit has a difficult week. A firm that loses everything simultaneously has a queue — and the queue, not the restore speed, is what sets the outage duration.

The boundaries that matter most for recovery

Boundary	Why it changes the recovery picture
Backup infrastructure	Non-negotiable. Its own segment, its own identity plane, inbound connections only from defined sources.
Identity infrastructure	If the directory is encrypted, you cannot authenticate to begin recovery. Domain controllers need their own protection and their own tested restore path.
Hypervisor and storage management	Encryption at this layer takes every workload above it at once, and is the difference between losing servers and losing the estate.
Business units and environments	Limits how much has to be recovered simultaneously. This is the boundary that most directly shortens outage.
OT / production systems	Where recovery is slowest and consequences physical. Usually the segment that most justifies its cost.

PART TWO

Proving You Can Recover

6 · A Backup You Have Never Restored Is a Hypothesis

The trap that catches serious, well-funded organisations: they have backups, they feel covered, and then on the worst day they discover the backups do not restore.

Every experienced practitioner has a version of this story. The job that had been silently failing for months. The critical system that was never in scope because it was built after the backup policy was written. The restore process so slow that “we will be back” means three weeks rather than three hours. The database that restores technically but comes back inconsistent because nobody had ever tested it with the application attached.

JOB SUCCESS IS NOT RESTORE SUCCESS

Backup software reports on whether it finished writing. That tells you nothing about whether the result is restorable, complete, consistent, or fast enough to matter. A dashboard of green jobs is one of the most reassuring and least informative artefacts in enterprise IT, and it is what most boards are shown.

The four failure modes, and what finds each

Failure	What it looks like	Found by
Coverage gap	A system nobody ever added. Almost always found during the incident.	RR-01, reconciled against the asset inventory
Silent job failure	Backups stopped weeks ago; alerting was to a mailbox nobody reads.	RR-01, plus alerting on absence
Unrestorable data	Restores, but the application will not start or data is inconsistent.	RR-02 with application-level validation
Too slow to matter	Restores correctly, over eleven days, against a tolerance of two.	RR-02 with a stopwatch

7 · Restore Testing: Scope, Cadence, Stopwatch

Tested on a schedule, timed, and signed off, so that “can we recover?” is a known answer rather than a prayer.

Three properties a restore test must have

It restores to somewhere other than where it came from. Restoring in place proves the data is readable. It doesn’t prove you could rebuild after losing the platform, which is the scenario that matters.

It is validated at the application layer, by someone who would notice. A database that mounts isn’t a working system. Validation means the application starts, a transaction completes, and a person familiar with the system confirms the data looks right. That person is usually not in the infrastructure team.

It is timed end to end and the number is recorded. From decision-to-restore through to service-usable, including the parts nobody counts: locating media, provisioning capacity, network configuration, and waiting for someone with the right access to wake up.

A workable cadence

Tier	Restore test	Full rehearsal
Tier 1 — critical	Monthly, application-validated	Annually, full sequence under RR-10
Tier 2, important	Quarterly	Included in the annual rehearsal scope on rotation
Tier 3 — standard	Semi-annually, sampled	Not individually rehearsed
Identity infrastructure	Quarterly, into an isolated environment	Annually, as the first step of the rehearsal

The identity row is deliberate. Directory services are simultaneously the most likely to be targeted, the most complicated to restore correctly, and the prerequisite for restoring anything else. Organisations routinely discover during an incident that their domain controller recovery procedure was written for a different forest topology three migrations ago.

8 · The Number You Must Be Able to Say Out Loud

If everything were encrypted tonight, how long until we are running again? A leadership team that cannot answer that in a sentence does not have a recovery capability — it has backup software.

The number needs three qualities. It must be **measured**, from rehearsal rather than estimated from throughput arithmetic. It must be **expressed per business service** rather than per system, because leadership cares about the service and the dependency chain is where the time goes. And it must have a stated **confidence**: tested last month is different from tested last year is different from never.

Why measured and calculated numbers diverge so badly

Calculated recovery time takes data volume and divides by restore throughput. Measured recovery time includes everything the calculation omits, which is most of the elapsed duration: deciding to invoke, assembling people, building somewhere clean to restore into, restoring identity first, sequencing dependencies, validating each service, and discovering that two systems each need the other to be up first. Organisations that compare the two typically find the measured number is three to ten times the calculated one, and the first rehearsal is where that gets discovered.

THE HONEST REPORTING FORMAT

“Tier 1 services: 38 hours to first service restored, 4 days to full operation. Measured in the March rehearsal, scope covered 6 of 9 Tier 1 services, identity restore included. Confidence: moderate, the payments platform was excluded and has never been rehearsed.” That is a useful sentence. “Our RTO is 24 hours” is an aspiration written in a policy.

9 · The Isolated Recovery Environment

You cannot restore into an environment you no longer trust. The place you recover into has to exist, be reachable, and be clean, and almost nobody has one until they have needed one.

This is the most commonly missing piece of an otherwise mature capability. The backups are immutable, the tests pass, and then the incident arrives and there's nowhere to put the data: production is compromised and can't be trusted, the standby site was joined to the same directory, and building somewhere clean under pressure adds days.

What the clean room needs

Requirement	Notes
Separate identity	Its own authentication, no trust to the compromised directory. This is what makes it clean.
Network isolation	No routing to production during rebuild and validation. Controlled egress only when a service is declared clean.
Independent access path	Administrators must be able to reach it when the corporate estate is down. Out-of-band, with credentials stored outside production.
Capacity, arranged in advance	Either standing capacity or a documented, tested commercial path to it. "We will spin it up in cloud" is an assumption until someone has done it with the actual account limits in place.
Read path to backups	Must be able to read the immutable store without depending on production management infrastructure.
Forensic capability	You will be examining restored systems for the attacker before releasing them. Chapter 11.

PART THREE

The Recovery Itself

10 · Recovery Is a Plan, Not a Button

Good backups do not save you if nobody knows how to turn them into a running business under pressure. The middle of a crisis is the worst possible time to invent a recovery process.


Three questions have to be answered before the incident, because answering them during costs hours you don't have.

What order do systems come back in? Not by importance — by dependency. Identity first, then name resolution and certificate services, then core data platforms, then the applications on top, then the integrations between them. Most organisations have never written this down, and the order isn't obvious even to the people who built the systems.

Who is allowed to make the call to start restoring? Restoration destroys forensic state and commits you to a course of action. If that authority is undefined, the first hours are consumed by a debate about whether to preserve evidence or restore service. Name the role now.

How do you avoid restoring the same infection? Chapter 11, and the answer determines your recovery point far more than backup frequency does.

RECOVERY SEQUENCE — DEPENDENCY ORDER, NOT IMPORTANCE ORDER



EVERY STEP HAS A VALIDATION GATE

Nothing moves to the next stage until the previous one is confirmed clean and functional.

The most common sequencing error is restoring applications before identity, discovering nothing can authenticate, and rolling back — losing a day to a dependency that was written down nowhere.

Integrations last, deliberately: reconnecting to partners before you are clean re-exposes them to you.

Sequence is set by dependency. Appendix D is the template for recording yours before you need it.

11 · Restoring Without Restoring the Attacker

Dwell time is the reason recovery point is harder than it looks. If they were inside for six weeks, a backup from last night contains them.

This is the part of ransomware recovery that surprises people who have planned only for data loss. The intuitive move — restore the most recent clean-looking backup — frequently restores the persistence mechanism along with the data, and the organisation is re-encrypted days later, sometimes during the recovery itself.

Recovery point is a forensic decision, not a backup decision

Choosing what to restore from requires knowing roughly when the attacker gained access, which requires investigation running *in parallel* with recovery rather than before it. That parallelism is uncomfortable and it is the only workable answer: waiting for a complete forensic picture takes weeks, and restoring blind risks repeating the whole event.

In practice this means recovery proceeds with a working hypothesis about initial access, that hypothesis is revised as investigation continues, and the recovery plan must tolerate being told on day four that the chosen restore point was too recent.

What reduces the risk

Rebuild rather than restore, wherever it is feasible. Operating systems and applications rebuilt from known-good images with data restored on top carries far less risk than restoring whole systems. It is slower per system and dramatically safer, and it is the reason infrastructure-as-code pays off in ways nobody costed.

Scan restored systems before they join anything. In the clean room, with current indicators, before network connectivity is granted.

Treat credentials as universally compromised. Every password, key, certificate and token in the affected scope. This is a large, unglamorous programme of work that runs alongside recovery, and skipping it is how organisations get hit twice by the same actor.

Keep the immutable copies from before you knew. Do not let retention expire on backups spanning the suspected dwell period. You may need to reach further back than you first thought, and that decision often arrives after the older copies would otherwise have rotated out.

12 · The Payment Decision

This manual exists to make the question uninteresting. It is worth being clear about how the decision should be governed, because organisations that improvise it under pressure make choices they cannot defend afterwards.

Payment is a business decision with legal consequences, not a technical one. It must not be made by the incident team, and it must not be made in the first hours. Decide the governance now: who decides, who must be consulted, what must be established before the question is even considered.

What has to be established before the question is live

Question	Why it comes first
Is payment lawful here?	Sanctions exposure depends on the actor's attribution and your jurisdiction. This is a counsel question with criminal consequences and it is not optional.
What is the realistic recovery path without paying?	The decision is a comparison. Without a measured recovery estimate you are not comparing anything — which is precisely the leverage the attacker is selling.
What are we actually buying?	A decryption tool of unknown quality, or a promise not to publish. Decryptors are frequently slow and partial; the promise is unenforceable.
Who must be told, and when?	Regulator, insurer, board, law enforcement. Several have clocks that have already started.
Does the insurer have a position?	Policies increasingly constrain this. Discover it now, not during.

DECIDE THE GOVERNANCE, NOT THE ANSWER

A pre-committed “we will never pay” is a position many organisations hold and some abandon at the worst moment, which is worse than not having held it. What genuinely helps is deciding *who decides*, what must be established first, and who must be in the room, and rehearsing that in the tabletop under RR-09.

13 · Notification, Law Enforcement and the Clock

Several clocks start before you understand the incident, and at least one of them is shorter than the time it takes to know what happened.

Regulatory notification obligations for a material incident are frequently measured in hours, not days, and they typically begin at awareness rather than at understanding. Waiting for a complete picture will breach them. The workable approach is an initial notification that says what is known and what isn't, followed by updates, which is what most regimes expect and what most organisations feel unable to do.

Who to engage and roughly when

Party	Timing and note
Legal counsel	Immediately, before technical containment decisions. Privilege considerations affect how investigation is structured.
Cyber insurer	Within hours. Policies often require notification and may mandate specific responders; using your own first can prejudice cover.
Regulators	Per your obligations, from awareness. Initial partial notification beats a late complete one.
Law enforcement	Early. They may hold intelligence on the actor, including decryption capability, and their involvement can matter to sanctions questions.
Customers and partners	Once impact is understood well enough to be accurate. Partners connected to you may need to act before you are ready to speak publicly.
Board	Immediately for a material event, and on a defined cadence after. Set the cadence in advance or it consumes the incident team.

PART FOUR

Runbooks

RUNBOOK RR-01

Backup Scope and Coverage Review

PURPOSE	Establish what is actually backed up, reconciled against what exists — the gap nobody finds until the incident.
TRIGGER	Quarterly, and after any material infrastructure change or acquisition.
OWNER	Infrastructure runs it. Security reviews the exceptions.
EVIDENCE	Coverage register, unreconciled asset list, exception decisions.

1 Start from the asset inventory, not the backup console.

Reconciling the backup system against itself only proves it is internally consistent. The population must come from an independent source — CMDB, cloud inventory, hypervisor list.

2 Reconcile and list the difference explicitly.

Every asset with no corresponding backup job is a finding with a named owner, not a footnote. Expect the first run to be uncomfortable; systems built after the policy was written are the usual gap.

3 Check job outcomes over a period, not the latest run.

A job that succeeded last night and failed for the eleven nights before is a failure. Report on consecutive success streaks, not point-in-time status.

4 Alert on absence, not only on failure.

A job that stopped being scheduled reports nothing at all, and nothing looks like success on most dashboards. This is the same silent failure mode that affects control testing generally.

5 Verify retention against the recovery point you actually need.

Given dwell time, a 14-day retention may not reach back beyond the intrusion. Retention is a security parameter, not a storage-cost parameter.

6 Confirm what is out of scope, and that someone decided it.

Undocumented exclusions are the difference between an accepted risk and a nasty discovery.

7 Record coverage as a percentage of Tier 1 services, not of servers.

Ninety-five per cent of servers can still mean one Tier 1 service is unrecoverable.

RUNBOOK RR-02

Restore Test

PURPOSE	Convert a backup from a hypothesis into a measured, validated, signed-off capability.
TRIGGER	Per the cadence in Chapter 7. Also after any backup platform change.
OWNER	Infrastructure executes. Application owner validates. Security signs off the record.
EVIDENCE	Restore test record (Appendix C) with elapsed time, validation result and defects raised.

- 1 **Restore somewhere other than the source.**
In-place restore proves readability only. The scenario that matters is rebuilding after losing the platform.
- 2 **Start the clock at the decision, not at the data transfer.**
Include locating media, provisioning target capacity, obtaining access and assembling people. This is where most of the real elapsed time lives and where estimates are most wrong.
- 3 **Restore without consulting the person who built the system, where you can.**
If recovery depends on undocumented knowledge held by one individual, you have a single point of failure that no backup product will fix. Note it as a defect.
- 4 **Validate at the application layer with someone who would notice.**
Service starts, a representative transaction completes, and a business-side reviewer confirms the data looks correct. Infrastructure confirming that a disk mounted is not validation.
- 5 **Record the elapsed time honestly, including the delays that felt unfair.**
Waiting four hours for someone with the right access is part of your recovery time. Removing it from the record makes the number a fiction.
- 6 **Raise every friction point as a defect with an owner.**
The purpose of the test is to generate this list. A test that produced no defects usually means the scope was too easy.
- 7 **Sign off and publish the number.**
Into the recovery readiness report (Appendix F). A test whose result nobody sees does not change decisions.

DO NOT TEST ONLY THE EASY SYSTEMS

The natural drift is towards restoring the same well-understood, well-documented system every quarter because it passes. Rotate scope deliberately towards the systems people are least confident about — that is where the information is.

RUNBOOK RR-03

Immutability and Air-Gap Verification

PURPOSE	Prove that at least one copy cannot be destroyed by an attacker holding your most privileged credentials.
TRIGGER	Semi-annually, and after any change to backup platform, retention policy or cloud account structure.
OWNER	Security owns the assertion. Infrastructure provides access.
EVIDENCE	Attempted-deletion test result, lock mode confirmation, credential separation evidence.

- 1 Identify which specific copy is the last resort.**
Not the backup estate in general, the one copy you would rely on if everything else were gone. If nobody can name it immediately, that's the finding.
- 2 Establish the lock mode precisely.**
Governance mode (privileged users can override) or compliance mode (nobody can, including the account owner and vendor). Read the configuration; do not accept the product name as evidence.
- 3 Attempt to delete it, deliberately, with the highest privilege available.**
In a controlled test against a non-critical object under the same policy. The attempt must fail. This is the only step that actually proves the control, and it is the one most programmes skip.
- 4 Attempt to shorten the retention period.**
Same principle. A lock you can shorten to zero is not a lock.
- 5 Verify credential separation.**
Confirm that compromise of the production directory does not yield backup administrative access. Trace the trust relationships rather than asking whether they exist.
- 6 Verify tenancy-level survivability.**
Establish what happens to the copy if the account or subscription containing it is deleted. Object-level immutability inside a destroyable container is not last-resort protection.
- 7 Confirm the decryption keys are outside the blast radius**
and that someone other than one individual can retrieve them.
- 8 Confirm the read path works without production.**
If reading the immutable store requires management infrastructure that would be encrypted, the copy is unreachable exactly when needed.

RUNBOOK RR-04

Backup Infrastructure Hardening Review

PURPOSE Treat the backup estate as the high-value target it is, because the attacker already does.

TRIGGER Annually, and on platform change.

OWNER Security, jointly with infrastructure.

EVIDENCE Hardening checklist result, identity separation evidence, network exposure assessment.

1 Map who can administer the backup system, including indirectly.

Nested groups, delegated rights, hypervisor administrators who can reach the appliance, and anyone able to alter the underlying storage. The real list is always longer than the intended one.

2 Confirm a separate identity plane.

Backup administration must not be reachable from production directory compromise. Separate accounts, separate MFA, no trust path.

3 Assess network reachability from the general estate.

Backup infrastructure should accept connections from defined sources only. A workstation that can route to the repository is a workstation whose malware can too.

4 Check management-plane exposure.

Web consoles, APIs and agents. Backup product vulnerabilities are actively exploited and patching is frequently deferred because the system is considered internal.

5 Verify logging goes somewhere the backup administrator cannot alter.

Deletion of backup data must be detectable after the fact by someone other than the person who could perform it.

6 Alert on the destructive operations specifically.

Retention reduction, policy deletion, large-scale expiry, disabling of jobs. These are the attacker's preparation steps and they are visible if anyone is watching for them.

7 Confirm the patching position

for the backup platform, and treat it at the same tier as internet-facing infrastructure.

RUNBOOK RR-05

Ransomware Incident: The First Sixty Minutes

PURPOSE	Contain, preserve and mobilise without foreclosing recovery options.
TRIGGER	Confirmed or strongly suspected ransomware activity.
OWNER	Incident commander. Recovery lead appointed within the hour as a separate role.
EVIDENCE	Timeline, decision log, containment actions, notification record.

1 Declare, and appoint two distinct leads.

An incident commander and a recovery lead. One person cannot investigate and rebuild simultaneously, and the attempt is why hour one is usually wasted.

2 Protect the backups before anything else.

Before containment, before analysis. Verify the immutable copy is intact, confirm retention has not been altered, and if there is any doubt, isolate backup infrastructure from the network immediately. Everything else is recoverable from here; this is not.

3 Preserve before you contain.

Memory and disk images from representative affected systems. Containment destroys the evidence that establishes dwell time, and dwell time determines your recovery point.

4 Contain by identity and network together.

Disable compromised accounts, invalidate active sessions, and isolate affected segments. Remember that disabling an account does not terminate sessions already issued.

5 Do not power off encrypted systems indiscriminately.

It destroys memory-resident evidence and occasionally destroys keys that would otherwise be recoverable. Isolate rather than shut down unless there is active spread.

6 Engage counsel and the insurer within the first hour.

Both have consequences for how the rest proceeds, and the insurer may mandate specific responders.

7 Establish out-of-band communications.

Assume email and chat are compromised or unavailable. This must be pre-arranged; setting it up during is a wasted hour.

8 Start the notification clock consciously.

Record the moment of awareness. Regulatory obligations run from it and you will be asked to evidence it.

9 Do not restore anything yet.

Recovery point is a forensic decision (Chapter 11). Premature restoration reinstates the attacker and destroys the state needed to determine when they arrived.

RUNBOOK RR-06

Building the Clean Room

PURPOSE Establish a trusted environment to restore into, isolated from the compromised estate.

TRIGGER Incident declared. Also built and torn down annually as part of RR-10.

OWNER Recovery lead.

EVIDENCE Build record, isolation verification, access path test.

1 Provision isolated capacity.

From the pre-arranged path. If this is the first time anyone has tried, expect account limits, procurement friction and identity dependencies to cost a day, which is why it belongs in the annual rehearsal.

2 Stand up independent identity.

New directory or tenancy with no trust to the compromised environment. Administrative credentials created fresh, held by named individuals, MFA enrolled.

3 Verify isolation before any data arrives.

Confirm no routing to production. Restoring into an environment that can still reach the compromised estate defeats the purpose entirely.

4 Establish the read path to immutable backups

without traversing production management infrastructure.

5 Deploy detection and forensic tooling first,

before restoring anything, so restored systems are examined as they arrive rather than retrospectively.

6 Confirm administrator access works from outside the corporate network.

Assume the normal access path is unavailable.

7 Define and document the egress gate.

The specific criteria under which a validated service is permitted to talk to the outside world, and who authorises it.

RUNBOOK RR-07

Recovery Sequencing and Execution

PURPOSE	Return services in dependency order, with a validation gate at each stage.
TRIGGER	Clean room ready and recovery point selected.
OWNER	Recovery lead. Application owners validate their own services.
EVIDENCE	Sequencing plan with timestamps, validation sign-offs, deviation log.

- 1 Confirm the recovery point with the investigation lead.**
In writing. This is the decision most likely to be revisited on day four, so record the basis for it.
- 2 Restore identity infrastructure first and validate it fully.**
Directory, DNS, certificate services, time. Nothing else can authenticate until this is right, and it is the restore most likely to reveal a procedure written for an older topology.
- 3 Rebuild rather than restore wherever feasible.**
Operating systems and applications from known-good images, with data restored on top. Slower per system, dramatically lower reinfection risk.
- 4 Scan every restored system before granting it connectivity.**
With current indicators, in the clean room, per RR-08.
- 5 Rotate credentials as services come up,**
not afterwards. Everything in the affected scope is compromised: passwords, keys, certificates, service accounts, API tokens.
- 6 Validate each service with its business owner before declaring it recovered.**
Technically running is not recovered. The person who uses it decides.
- 7 Reconnect integrations last, and deliberately.**
Partners and external feeds after your own estate is validated clean — reconnecting early re-exposes them to you and turns your incident into theirs.
- 8 Log every deviation from the plan.**
The deviations are the improvement backlog and they are also what the post-incident review and any examination will focus on.

RUNBOOK RR-08

Reinfection Prevention During Restore

PURPOSE	Ensure the recovery does not reinstate the attacker along with the data.
TRIGGER	Runs continuously alongside RR-07.
OWNER	Investigation lead, jointly with recovery lead.
EVIDENCE	Scan results per system, credential rotation record, indicator list and revision history.

- 1 Maintain a living indicator list**
and treat it as versioned. Systems cleared against an early list may need re-examining as investigation matures.
- 2 Hold every restored system in quarantine until cleared.**
No connectivity beyond the restore path until it has been examined.
- 3 Hunt for persistence specifically, not only for the ransomware payload.**
Scheduled tasks, service creation, registry run keys, web shells, new local accounts, altered group membership, and legitimate remote-access tooling the attacker installed.
- 4 Treat every credential in scope as compromised.**
Including service accounts, machine accounts, API keys, certificates and anything cached. Partial rotation is how organisations get hit a second time by the same actor.
- 5 Re-examine the recovery point if new evidence extends the dwell period.**
Have a defined mechanism for this — it is disruptive. It's sometimes necessary, and improvising it costs more.
- 6 Preserve backups spanning the suspected intrusion window.**
Suspend normal retention expiry immediately; you may need to reach further back than currently believed.
- 7 Monitor restored systems at elevated sensitivity**
for a defined period after recovery, and decide in advance how long that period is.

RUNBOOK RR-09

Ransom Decision Governance

PURPOSE	Ensure that if the question arises it is decided by the right people, on the right basis, lawfully.
TRIGGER	A demand is received. The governance is defined in advance and rehearsed.
OWNER	Named executive decides. Counsel advises. Incident team does not decide.
EVIDENCE	Decision record, legal advice reference, sanctions check, board notification.

- 1 **Route the demand to the pre-defined decision-maker.**
Never to the incident team, who are the least able to weigh it and the most under pressure.
- 2 **Obtain the sanctions and legality position from counsel before anything else.**
Payment to a sanctioned entity carries criminal consequences independent of the incident, and attribution is frequently uncertain.
- 3 **Notify the insurer and establish their position.**
Cover may be conditional on their involvement in this decision.
- 4 **Produce a measured recovery estimate from the recovery lead.**
The decision is a comparison, and without this number there is nothing to compare against, which is exactly the position the attacker is trying to create.
- 5 **State explicitly what payment would and would not buy.**
A decryptor of unknown quality and speed, or an unenforceable promise not to publish already-exfiltrated data. Write it down so the decision is made against reality.
- 6 **Engage law enforcement**
before deciding. They may hold decryption capability or actor intelligence that changes the calculation.
- 7 **Record the decision, the basis and the participants.**
This will be examined afterwards by regulators, insurers and possibly litigants.

RUNBOOK RR-10

Full Recovery Rehearsal

PURPOSE Convert the recovery plan from a document into a rehearsed capability with a measured duration.

TRIGGER Annually. Scope rotates across Tier 1 services.

OWNER Recovery lead. Observed by second line or internal audit.

EVIDENCE Rehearsal report with measured timings, defect list, and the recovery estimate it produces.

- 1 **Scope to a full business service and its dependency chain,**
not to a single system. The dependencies are where the surprises are.
- 2 **Build the clean room as part of the exercise.**
Per RR-06. Do not use a pre-built environment — the build is one of the longest and least understood steps.
- 3 **Assume the corporate estate is unavailable.**
No access to the normal documentation store, the normal chat tool, or the normal credential vault. Organisations routinely discover their recovery runbook is only reachable from the environment that is down.
- 4 **Run it with the people who would actually be available,**
not only the experts. If recovery depends on three named individuals, discover that now.
- 5 **Time every stage and record the elapsed total honestly.**
Including delays that felt unrepresentative — they will recur.
- 6 **Inject at least one complication.**
A second restore point required, a key person unavailable, an immutable copy that turns out to be the wrong scope. Perfect-conditions rehearsals produce optimistic numbers.
- 7 **Produce the recovery estimate and its confidence level,**
for Appendix F and the board.
- 8 **Turn every defect into an owned action with a date,**
and check them at the next rehearsal. A rehearsal whose findings are not closed is theatre with a stopwatch.

PART FIVE

Templates and Reference

A · Recovery Tier Definitions

Tiers are set by business impact of unavailability, agreed with the business, and drive every downstream decision. Adjust the tolerances to your sector.

Tier	Definition	Implications
Tier 1	Unavailability is existential or triggers regulatory consequence within hours.	Immutable copy mandatory. Monthly restore test with application validation. Annual full rehearsal. Recovery sequence documented and dependency-mapped.
Tier 2	Material operational impact within days.	Immutable copy. Quarterly restore test. Included in rehearsal on rotation.
Tier 3	Tolerable for a week or more.	Standard backup. Semi-annual sampled restore test.
Identity	Prerequisite for restoring anything else, regardless of its own tier.	Treated as Tier 1 always. Quarterly isolated restore test. First item in every rehearsal.

B · Backup Coverage Register

Field	Notes
Business service	The register is organised by service, not by server. Leadership recovers services.
Recovery tier	Per Appendix A.
Component systems	Everything the service depends on, including ones owned by other teams.
Backed up?	Yes / No / Partial — with the gap named where partial.
Immutable copy exists?	Yes / No, and where. “No” on a Tier 1 service is a reportable finding.
Retention	Assessed against dwell-time assumptions, not storage cost.
Last successful restore test	Date, elapsed time, validated by whom.
Measured recovery time	From rehearsal. Blank means unproven — and blank on Tier 1 is the headline finding.
Dependencies	What must be restored before this can be. Drives the sequencing plan.
Owner	Named individual, business side.

C · Restore Test Record

Field	Notes
Test date / system / tier	
Restore source	Which copy — primary, offsite, or immutable. Testing only the fastest copy proves the least.
Restore target	Must not be the source environment.
Clock start / clock stop	Decision to restore → service usable. Not data-transfer duration.
Elapsed total	The number that goes to the board.
Breakdown	Locate, provision, transfer, configure, validate. Shows where the time actually went.
Validation performed by	Named business-side reviewer, not infrastructure.
Validation result	Service started, transaction completed, data correct, each confirmed separately.
Defects raised	With owners and dates. The primary output of the test.
Undocumented knowledge relied on	Anything only one person knew. A single point of failure no product fixes.
Sign-off	

D · Recovery Sequencing Plan

One row per system, ordered by dependency. Written before the incident; used during it.

Order	System	Depends on	Validation gate — who confirms, and what
0	Clean room	—	Recovery lead: isolation verified, backup read path confirmed
1	Directory services	Clean room	Identity team: authentication works, replication healthy
2	DNS / certificate services / time	Directory	Infrastructure: resolution and certificate issuance functioning
3	Core data platforms	Identity	DBA: consistency checks pass, scanned clean
4	Tier 1 applications	Data platforms	Business owner: representative transaction completes
5	Tier 2 applications	Varies	Business owner per service
6	External integrations	All above validated	Security: estate declared clean before reconnection

E · Tabletop Scenario Pack

Four scenarios that surface different failures. Run one per quarter, ninety minutes, with the people who would actually be in the room.

Scenario	What it is designed to expose
1 · Friday night	Encryption begins at 19:00 Friday. Tests out-of-hours mobilisation, who has authority when executives are unreachable, and whether anyone can reach the recovery documentation from outside the network.
2 · The backups are gone	Attacker deleted the primary backup repository two weeks before encrypting. Tests whether the immutable copy is real, whether anyone can name it, and whether the read path works.
3 · Six-week dwell	Investigation reveals access began six weeks ago, after recovery has already started from a two-week-old restore point. Tests the mechanism for revising recovery point mid-recovery and whether older backups still exist.
4 · Exfiltration and a deadline	Data was taken and publication is threatened in 72 hours. Tests payment governance, notification decisions, and whether recovery capability actually changes the conversation.

F · Board Reporting: Recovery Readiness

One page, quarterly. Four numbers and a sentence.

Measure	Current	Commentary prompt
Tier 1 services with a verified immutable copy		Which are missing one, and when is that closed?
Tier 1 services with a measured recovery time		Unmeasured means unproven. Which are unproven?
Measured time to restore the first Tier 1 service		From the last rehearsal, with the date and scope stated.
Restore tests completed vs scheduled		Which slipped, and why?

The sentence: “If our estate were encrypted tonight, we would restore [service] within [measured hours] and full Tier 1 operation within [measured days], based on the [month] rehearsal covering [scope]. Confidence is [level] because [the honest caveat].” If that sentence can’t be completed truthfully, that’s the report.

G · Where This Gets Hard

The judgement calls this manual does not resolve for you — because the right answer depends on your sector, your risk appetite and your balance sheet. These are the conversations worth having properly rather than defaulting through.

How far back does retention need to reach? Dwell times of several months are documented. Retaining immutable copies that long is expensive and the cost is real. Somewhere between fourteen days and a year is your answer, and it is a risk decision rather than a storage decision. Most organisations set it by cost and then discover the constraint during an incident.

Rebuild or restore? Rebuilding from known-good images is safer and materially slower per system. Under a live outage with revenue stopping, the pressure to restore whole systems is enormous and sometimes correct. Where that line sits for you should be decided before the pressure exists, and it differs by tier.

What is a defensible recovery time? There is no universal answer, and benchmarks are unhelpful because scope definitions vary wildly. What matters is that the number is measured, that leadership has accepted it, and that the gap between it and what the business assumed is visible.

How much segmentation is worth the operational cost? Segmentation shortens outages and slows delivery. Firms that segment aggressively pay for it every day in friction against an event that may never come. This is genuinely a judgement call and anyone who tells you otherwise is selling something.

Do you pre-commit on payment? A public “we will never pay” is clarifying and occasionally abandoned at the worst possible moment, which is worse than never having said it. The alternative, deciding the governance without pre-committing the answer — is less satisfying and more durable.

Where does recovery sit organisationally? Backups usually live with infrastructure and are funded as operations. Treating recovery as a security capability means moving budget, ownership and reporting — which is an organisational fight, not a technical one, and it is the fight that most determines whether any of this actually happens.

H · Glossary

Air gap	A copy with no network path from the production estate. Increasingly logical rather than physical; the test is whether an attacker with maximum privilege could reach it.
Clean room	Isolated environment with independent identity, used to restore into when production cannot be trusted.
Compliance mode	Object lock that nobody can shorten or override, including the account owner and vendor. Distinct from governance mode.
Double extortion	Exfiltration before encryption, so recovery alone does not end the incident.
Dwell time	How long the attacker was present before acting. Determines how far back a safe restore point sits.
Fate sharing	When a backup depends on the same credentials, network, platform or vendor as production, and is therefore lost in the same event.
Governance mode	Object lock a privileged user can override. Protects against accident, not an adversary with privilege.
Measured recovery time	Elapsed duration observed in rehearsal, decision to service-usable. Distinct from calculated RTO, usually by a large multiple.
Recovery point	The moment in time you restore to. In ransomware a forensic decision, not a backup-schedule one.

USING THIS MANUAL

Take it apart. Rename the runbooks, renumber them, delete what does not apply, and argue with the parts you disagree with — particularly the judgement calls in Appendix G, where your sector and your appetite should produce different answers than mine would.