Phase 1 — VISIBILITY (Days 1–5)

Key question:

👉 *Do we actually know what is connected to our environment?*

The first phase is not about deploying new tools, running vulnerability scans, or buying additional security products. It is about an honest reality check. Most organizations believe they understand their IT environment — until they are forced to prove it under pressure. The first five days are designed to remove assumptions and replace them with clarity.

At the executive level, visibility means understanding **where real attack surfaces exist**, not where they could theoretically exist.

A primary focus is **edge and legacy systems**. These typically include old laptops still in occasional use, legacy NAS devices, network printers, displays, IoT devices, test environments, or operational systems that have "always been there." These assets often share the same characteristics:

- They are outdated or unpatched
- Poorly documented
- Rarely monitored
- And ownership is unclear or assumed

From a ransomware perspective, they are ideal entry points precisely because they sit outside day-to-day attention.

The second focus area is **customer and third-party access**. The question here is not whether external access is necessary — it usually is — but whether it is **intentional, controlled, and understood**. This phase examines:

- Existing VPN connections
- RMM tools installed by service providers
- Shared or legacy accounts
- API connections and automation credentials
- Old "temporary" access that was never revoked

In many environments, access paths exist that nobody actively uses anymore — but attackers certainly can.

A third critical element is **identity reality**. Policies often look clean on paper, but real access rights tell a different story. The goal here is to understand what actually exists today:

- Who currently holds administrative privileges?
- Where is multi-factor authentication missing?
- Which service accounts operate with excessive permissions?
- Which cloud or SaaS identities represent high-impact risk if compromised?

Modern ransomware rarely relies on brute force alone. Identity abuse is the fastest way to escalate and spread.

Finally, **ownership clarity** is addressed. For every relevant system, leadership must be able to answer:

- Who is responsible for it?
- Is that responsibility internal or external?
- Who makes decisions when something goes wrong?
- Who understands how the system actually works?

“This is handled by our provider” is not ownership — it is a blind spot.

Output of Phase 1:
The result is not a long report, but a focused outcome:

✓ A short list of the **3–5 most realistic ransomware entry points**
✓ Based on actual usage, access paths, and responsibility
✓ Likely scenarios — not theoretical risks

Phase 2 — CONTAINMENT (Days 6–10)

Key question:
👉 *If something breaks, how far can it spread?*

Once likely entry points are understood, the focus shifts from prevention to impact control. The question is no longer *“How do we stop every attack?”* but *“How do we prevent one compromised system from taking everything down?”*

The first area examined is **segmentation reality**. On architecture diagrams, networks often appear neatly separated. In reality, they are frequently not. Common issues include:

- Production systems with unrestricted access to office networks
- Backups residing in the same segment as live systems
- Administrative access spanning multiple zones
- Cloud environments without meaningful workload or role separation

This phase is not about redesigning the architecture — it is about identifying what *should* be isolated but currently is not.

Next is **access minimization**. Persistent access is convenient, but it dramatically increases risk. This step evaluates:

- Which access rights are genuinely required on a daily basis
- Which administrative permissions could be time-bound
- Which accounts exist only "just in case"

Every unnecessary standing privilege increases the potential blast radius.

A third priority is **MFA prioritization**. This is not about universal coverage overnight, but about strategic impact. MFA must be enforced where compromise would be catastrophic:

- Administrative accounts
- Email systems
- Remote access pathways
- Cloud control planes and core management portals

Effectiveness matters more than completeness.

Finally, **external trust paths** are reviewed. This includes access held by MSPs, customers, automation tools, and integrations that may not be visible in daily operations. The key question is simple:

- If this external access is compromised, how far can an attacker go?

Output of Phase 2:

✓ 3–5 concrete containment actions
✓ Feasible within days, not quarters
✓ Effective even if one system is already compromised

PAGE 2 — RECOVER & DECIDE

Phase 3 — BACKUP & RECOVERY REALITY (Days 11–12)

Key question:

👉 *Could we recover without negotiating?*

This phase is often the most uncomfortable — because it challenges assumptions. Most organizations have backups. Far fewer have proven recovery.

What matters here is not what sounds reassuring, but what actually works under pressure. Key questions include:

- Are backups truly isolated or immutable?
- Can production systems reach them?
- Do offline or write-protected copies exist?

Equally important is the **last successful restore** — not in theory, but in reality:

- When was the last restore test performed?
- Which systems were involved?
- How long did recovery take?

Many organizations cannot answer these questions with confidence — and that uncertainty is itself a risk.

Another focus is **time to recovery**. The critical issue is not eventual restoration, but operational impact:

- How long until critical systems are usable again?
- Which business processes are unavailable during that time?
- What does that mean financially and operationally?

Finally, a frequently overlooked question is addressed: **Who can actually restore systems?**

- What if internal IT or the MSP is unavailable or affected?
- Is recovery knowledge centralized or distributed?
- Are procedures documented and accessible during an incident?

Output of Phase 3:

✓ A **realistic RTO/RPO estimate**
✓ Identification of the **single largest recovery risk**

Phase 4 — DECISION PLAYBOOK (Days 13–14)

Key question:
👉 *Who decides what — under pressure?*

Technical failures can be fixed. Decision paralysis causes lasting damage. The final phase therefore focuses entirely on leadership, authority, and decision-making under stress.

First, the **ransom decision** is defined *before* it becomes urgent:

- Who has authority to decide pay vs. don't pay?
- On what criteria is that decision based?
- What role do legal counsel, insurers, and executives play?

Next, the **first 24 hours** are clarified at a decision level, not operational detail:

- Who must be informed?
- Who communicates internally and externally?
- Who engages insurers, legal advisors, regulators, or customers?
- Who coordinates the response overall?

At the center is **leadership clarity**. There must be:

- **One** decision owner
- **One** source of truth
- No parallel command structures

Output of Phase 4:

✓ A **one-page incident decision framework**
✓ Clear enough to function under pressure and uncertainty

EXECUTIVE SUMMARY (for internal sharing)

After 14 days, leadership should be able to clearly answer:

- What are our **top three ransomware risks**?
- What are our **top five actions for the next 90 days**?
- Can we realistically recover **without paying**?
- Who decides **what** in the first 24 hours?

Final Note

This roadmap is not a checklist.
It is a **decision framework**.