IstroSec

GRYPHON CASE STUDY: PREVENTION ACTIVE BREACH RESPONSE

Date: 14/06/2026

TABLE OF CONTENTS

SITUATION AT A GLANCE

*Illustrative scenario. **"Example Hospital"** is a fictional composite created to demonstrate documented GRYPHON capabilities; it is not a real customer. The incident is representative of real post-breach ransomware engagements. Every product capability referenced is part of GRYPHON's published feature set.*

Attribute	Detail
Organization	**Example Hospital** — standalone community hospital with around 250 beds (fictional)
Endpoints	Approximately 1,400 Windows endpoints, including legacy Windows 7 / Server 2008 R2 on imaging and lab systems
Event	Ransomware detonated around 02:00 on a Saturday; EHR, scheduling and shared drives encrypted or offline
Attacker state	Domain admin obtained; incumbent antivirus and EDR partially disabled or tampered with; backup integrity uncertain
Constraint	Identity infrastructure could not be trusted; the team had to contain, eradicate, recover and investigate at the same time
GRYPHON role	Deployed into the live, compromised environment as the incident response platform to regain control

THE REALITY OF POST-BREACH RESPONSE

When ransomware has already detonated in a hospital, the clock is both clinical and regulatory. Clinicians revert to paper, imaging and lab turnaround slows, and the organization is now inside a 60-day window in which any breach of unsecured protected health information (PHI) affecting 500 or more individuals must be reported to the HHS Office for Civil Rights and posted publicly.[1] Healthcare incidents already take an industry-worst 279 days on average to fully identify and contain;[2] every hour of uncontrolled access widens both the clinical and the reportable-data impact.

The deeper problem is architectural. Most security tools assume they are being installed into a clean, trusted environment. In an active breach the opposite is true: the attacker may hold domain admin, the existing EDR may be blinded or switched off, the identity directory itself is suspect, and a staged two-weeks rollout with a learning baseline is a luxury nobody has.[3]

WHY GRYPHON WAS CHOSEN FOR AN ACTIVE BREACH

GRYPHON is architected to deploy **inside** a compromised environment — it is the product's defining use case, not an afterthought. Several properties made it deployable on the Saturday morning:

- **Self-contained, no baseline period.** The pure C++ service and kernel driver begin protecting first boot. Both engines — behavioral policy and integrated anti-malware — work immediately, with no learning window to wait out.
- **Zero-trust onboarding.** Fresh agents land in a quarantine-like "new" group with minimal policy and no control surface; promoting an endpoint into active monitoring is an explicit operator decision. In a compromised estate, this prevents a rogue or attacker-planted agent from being automatically trusted.
- **Co-exists with compromised tools.** GRYPHON runs safely alongside whatever the attacker may already have subverted, and its kernel-level anti-tamper protections — covering shutdown, service, uninstall and network hardening — keep the attacker from removing it. Uninstall attempts to surface as alerts whether or not they succeed.
- **No cloud dependency, universal deployment.** Detection, response, quarantine and recovery all run on the endpoint, so even air-gapped OT systems are covered, and the single MSI deploys through GPO, SCCM, Intune or any standard tool — across modern and legacy Windows alike.
- **Recovery starts at deployment.** Backup (VSS) protection and automatic backup (VSS) creation begin on install, so anything the attacker damages after GRYPHON's arrival is recoverable on the endpoint.

THE FIRST 24 HOURS

1.1. Hours 0–2 — Mass deployment and first triage

The incident response team pushed the **IstroSetup.msi** agent by GPO and SCCM across the reachable estate and installed it manually on isolated OT systems, including the legacy Windows 7 imaging and lab terminals the incumbent EDR had never covered. Every new agent entered the **'new'** group under zero-trust onboarding. As each came online, GRYPHON's **Automatic Initial Scan** ran, and the integrated anti-malware engine quarantined known-malicious files — each detection (**MaliciousFileDetected)** retained with its SHA256 and path for the investigation.

1.2. Hours 2–6 — Live investigation from one console

Using the Incident Response Console, the team worked every endpoint from a single view: process tree, network context, user context, file changes and alerts from both engines, each tagged with the engine that fired, the action taken, and the account involved. The **Remote Process Manager** exposed live processes with their identifiers (PIDs), command lines, executable hashes and a reputation score derived from environment prevalence; **Process Reputation** in Asset Management ranked rare, unsigned binaries to the top. A remote **(PowerShell) terminal** and **remote file manager** let analysts examine and pull artefacts in place, and the malicious processes were killed and blocked across the estate.

1.3. Hours 4–8 — Containment without going dark

Rather than pulling cables, the team used GRYPHON's **Network Isolation**, which activates a separate **Isolation Firewall** defaulting to BlockAll. They edited it to permit only what the investigation needed — DNS, patch infrastructure, the SIEM destination and IR tooling — building network islands while GRYPHON retained management access to every isolated host. Patient-zero and the lateral-movement hosts were contained while clinical systems that had to keep communicating were kept reachable. Forensic collection scaled through **custom response scripts** (structured PowerShell/cmd output across many endpoints at once) and established third-party tools (KAPE and Sysinternals), pushed through GRYPHON's own agent channel. Because the directory was untrusted, responders worked through **temporary IR accounts** created on the endpoints instead of the hospital's normal administrative identities.

1.4. Hours 8–24 — Eradication, recovery and root cause

With scope established, eradication proceeded: persistence mechanisms inventoried in Asset Management were removed, malicious files quarantined, and GRYPHON's anti-tamper protection kept the attacker from re-disabling protection during clean-up. For recovery, files damaged **after** GRYPHON's arrival were restored automatically from protected backups (shadow copies) on the endpoint; systems encrypted before deployment were rebuilt clean and restored from off-host backups, with GRYPHON preventing any further encryption throughout. **File Integrity Monitoring** turned 'what changed during the incident window?' into a quarriable answer, pivotable to the process that made each change, and the **Vulnerabilities** and **Pending Updates** views pinpointed the unpatched, internet-facing service — flagged with a high GRYPHON Severity score and an exploited-in-the-wild marker — that had been the entry point, so it could be closed before recovery completed.

OUTCOME

Phase	Result with GRYPHON
Containment	Active spread halted within hours via Isolation-Firewall network islands; management access retained throughout
Eradication	Malicious processes killed and blocked, persistence removed, malware quarantined; anti-tamper protection blocked re-compromise during clean-up
Recovery	Post-deployment damage automatically recovered on the endpoint; pre-existing encryption rebuilt from clean backups with no further loss
Root cause	Entry vector identified via Vulnerabilities and Pending Updates closed; persistence and foothold documented
Evidence	Forensically organized timeline, process trees, hashes and file-integrity records for the OCR notification and possible litigation

HOW EACH CAPABILITY CONTRIBUTED

GRYPHON capability	Role in this response
Self-contained C++ agent with no baseline and kernel detection on first bot	Protection began the moment it was deployed into the live breach
Zero-trust 'new' group onboarding	Prevented attacker-planted or rogue agents from being trusted
Universal deployment (MSI / GPO / SCCM / Intune) with legacy and OT support	Covered the whole estate fast, including systems the EDR could not reach
IR console with Remote Process / PowerShell / File / Desktop tools	Single-pane live investigation and response across endpoints
Customizable Isolation Firewall (network islands)	Contained hosts while preserving needed connectivity and management
Custom response scripting and third-party tool channel (KAPE, Sysinternals)	Scaled forensic collection across thousands of endpoints
Temporary IR accounts	Operated without trusting compromised identity infrastructure
Kernel-level anti-temper	Kept the attacker from removing GRYPHON during eradication
Backup (VSS) protection and autonomous recovery	Recovered everything damaged after deployment; blocked further encryption
Asset management: vulnerabilities, pending updates, persistence, process reputation	Found the entry vector, persistence and rogue binaries
File integrity monitoring with SIEM forwarding and API access	Answered 'what changed' and fed the team's SIEM in real time

THE COMPLIANCE AND RECOVERY PAYOFF

Two outcomes mattered most to **Example Hospital**. First, control was regained in hours rather than the days a from-scratch toolchain would have cost — and faster containment is the single biggest lever on breach cost in the IBM data, with a tested incident response capability worth roughly US$2.66 million. Second, the response produced a defensible, time-stamped evidentiary record: a clear scope of which systems and records were actually affected, which is exactly what counsel and the HHS Office for Civil Rights require to make — and substantiate — the 60-day breach determination.

> ***Walking into a breach and walking out with control, recovery and evidence is the scenario GRYPHON was built for.***

[1]HIPAA Breach Notification Rule (45 CFR §§164.400–414), enforced by the HHS Office for Civil Rights: breaches of unsecured PHI affecting 500+ individuals must be reported to HHS OCR without unreasonable delay and no later than 60 calendar days from discovery, are published on the HHS public breach portal, and trigger an OCR compliance review (hhs.gov; hipaajournal.com).

[2]IBM, Cost of a Data Breach Report 2025: healthcare highest-cost sector for the 14th consecutive year at US$7.42M average; US average US$10.22M; healthcare breaches averaged 279 days to identify and contain; a tested incident-response plan is associated with roughly US$2.66M lower breach cost (ibm.com; hipaajournal.com).

[3]All GRYPHON capabilities described in this document are drawn from the published product feature set at gryphon.istrosec.com (Incident Response, Advanced Defense, Zero Trust, Deployment, Assets & Network, Integrations, and FAQ — including the “deploy inside an active breach” use case).