Industry Solutions

How autonomous prevention applies to specific sectors - healthcare, manufacturing and OT, financial services, government, education, maritime, and managed service providers.

- How does Cyber Crucible protect manufacturing, OT, and ICS environments?
- How does Cyber Crucible protect financial services firms?
- How does Cyber Crucible support government and high-security environments?
- Can Cyber Crucible protect ships and remote or disconnected operations?
- Why do MSPs and resellers partner with Cyber Crucible?
- We were breached recently. Can Cyber Crucible be deployed into an environment that may already be compromised?

How does Cyber Crucible protect manufacturing, OT, and ICS environments?

Short answer: It runs entirely on the endpoint with no cloud dependency, so it protects operational technology and industrial control systems that are offline, air-gapped, or intermittently connected — and it suspends only the malicious process, so production keeps moving.

Why conventional tools struggle here

OT and ICS environments can tolerate neither a cloud round trip nor a full-network lockdown. Many are deliberately disconnected. Security models that assume constant connectivity and a remote SOC simply don't apply.

What changes

- **100% offline capable** — full protection with no signal required, air-gap ready.
- **No lockdowns** — malicious processes are suspended in memory while production continues.
- **One agent across the environment** — the same lightweight agent protects PLCs on the plant floor, vehicles across a logistics network, and vessels at sea, with no analyst in the loop and no connectivity assumptions.

How does Cyber Crucible protect financial services firms?

Short answer: It stops infostealers and session-token theft at the kernel level, before exfiltration begins, and it analyzes everything locally so keys and tokens never have to be uploaded to a third-party cloud.

The threat financial firms actually face

Hyper-automated "smash and grab" attacks target session tokens and API keys. These move from infiltration to self-deletion in seconds — faster than a cloud alert can leave the building. A stolen session token often bypasses passwords and MFA entirely, giving an attacker a legitimate-looking way back in at any time.

The documented case

A major financial services firm ran three industry-leading EDRs and an outsourced Top-50 MSSP. Cyber Crucible was deployed to find the blind spot. Within 60 days it had autonomously intercepted nearly 10,000 malicious processes — 98% of them on the firm's highly privileged Microsoft SQL server farm — with no alerts from the legacy stack.

The root cause traced to a likely compromised remote monitoring credential. Months later, the industry reported a global wave of ransomware backdooring SQL servers worldwide. The financial sector suffered severe breaches. This client was never a victim.

How does Cyber Crucible support government and high-security environments?

Short answer: Through complete data sovereignty — fully on-premises installation in physically secured racks, air-gapped operation, and multi-tenant architecture — with no public cloud, no third-party SOC, and no data sharing.

Sovereignty by architecture

For organizations requiring absolute control, nothing needs to leave the environment. There are no third-party platforms in the path, no external access, and no vendor-hosted analysis of your data.

This matters because many vendors upload sensitive identity data — private keys, session tokens — to their hosted SOC teams. That centralized storage is a single point of failure and a high-value target for state-sponsored attackers, and it can be reached through legal compulsion without your knowledge.

Quiet by default

Threats are neutralized silently at the kernel level. Only you know an attack was attempted, and if no data was touched there are typically no external alerts or required disclosures.

Multi-tenant

The architecture supports multi-tenant deployments, which suits managing multiple departments or agencies from a single secure on-premises instance.

Can Cyber Crucible protect ships and remote or disconnected operations?

Short answer: Yes, and it has. In a maritime deployment, Cyber Crucible operated autonomously across multi-day voyages with no IT staff aboard, adapting to low-bandwidth satellite links and long disconnected stretches while protecting everything from passenger terminals to mission-critical maritime controls.

What the deployment demonstrated

- **Autonomous response** with no IT team present for days at a time.
- **Satellite-aware** operation across bursty, low-bandwidth, weather-interrupted links.
- **Forensic retention** — rich logs preserved through disconnection for post-voyage analysis.
- **Resilience under attack** — it held when adversaries tried to shut it down or degrade its performance.

The unusual part

Logs revealed activity consistent with nation-state involvement, and the targeting was selective: only vessels operating in international waters were affected, while domestic sightseeing vessels were untouched.

When conventional intrusion attempts failed, the adversaries escalated — hijacking Windows login credentials, then reaching BIOS-level settings to lock systems before boot. Even with systems halted before the operating system could load, the deployment continued to hold the line where traditional security stacks had already failed.

Why do MSPs and resellers partner with Cyber Crucible?

Short answer: Because autonomous prevention lets a partner own the margin, the customer relationship, and the SLA — instead of reselling a vendor-hosted SOC and absorbing the blame when that SOC misses a machine-speed attack.

The trap in reselling detection

The industry pushed partners toward reselling vendor-hosted MDR and SOC services. When human-driven monitoring inevitably misses an automated attack, the vendor misses the breach — and the customer doesn't fire the vendor. They fire the partner.

What changes operationally

- Threats are neutralized instantly, before analysts receive an alert.
- Alert volume drops dramatically from the 100+ per endpoint per day that drives alert fatigue.
- No YARA rule tuning and no manual threat hunting.
- Partners capture the revenue stream rather than settling for vendor commissions.

And on the risk side

A partner whose security offering prevents attacks rather than documenting them has a fundamentally different conversation with clients about outcomes and accountability.

We were breached recently. Can Cyber Crucible be deployed into an environment that may already be compromised?

Short answer: Yes. Cyber Crucible evaluates what programs are doing right now rather than searching for known-bad files, so it begins stopping malicious behavior on deployment — including from an attacker who already has a foothold.

Why prior compromise doesn't blind it

Because prevention is based on behavioral intent at the kernel level, an attacker who is already resident still has to *act* — accessing identity data, injecting into processes, beginning encryption, or moving data. Those actions are exactly what triggers interception.

This is how the financial services deployment surfaced an active, ongoing intrusion the existing stack had never alerted on. Cyber Crucible wasn't brought in to investigate a known breach; it was deployed to look for a blind spot, and it found nearly 10,000 malicious processes already in progress.

What to expect

Deploying into a compromised environment often produces immediate, high-volume interception activity. That is the tool working as intended, and it is frequently the first hard evidence an organization has that something was already underway.