OPERATOR MANUAL — CONFIDENTIAL

HYVE RAPTOR

Post-Quantum AI Agent Defense Platform
Complete Installation, Operation & Reference Guide

PUBLISHER	DOCUMENT	VERSION	CLASSIFICATION
Vibe Software Solutions	HYR-OPS-001	1.0 — 2026	Client Confidential

Table of Contents

01 Platform Overview

HYVE Raptor is a post-quantum AI agent defense platform. It monitors AI agents running on your clients' machines, detects behavioral anomalies in real time, and gives you — the operator — the visibility and control to respond to threats immediately from a full desktop command center.

Everything runs on infrastructure you own. No data leaves your network. No licensing authority exists outside your server. No internet connection is required for operations once deployed.

How the Platform Is Organized

HYVE Raptor is composed of three components that work together. Understanding what each one does and where it runs will help you deploy and troubleshoot the platform with confidence.

COMPONENT	WHAT IT DOES	WHERE IT RUNS
HYVE Backend	Issues and validates license tokens. Every shield must authenticate here before it can connect to the command server. You control who gets access.	Your server — port 9443
Raptor Command	The threat intelligence hub. Receives encrypted threat data from all connected shields, scores events, stores them, and exposes your operator dashboard.	Your server — ports 8443 & 8080
Raptor Shield	Watches every AI agent on the client machine. Detects behavioral changes, capability drift, and unauthorized activity. Transmits encrypted threat events back to your server.	Client machines (Windows service)
Raptor Command Center	Your operator desktop application. Full threat feed, shield management, defense command dispatch, compliance posture, and Alpha AI Advisor — all in one window.	Your machine (Electron desktop app)

The Data Flow

A threat event originates on a client machine, flows to your server encrypted with post-quantum cryptography, is stored and scored, and surfaces in your Command Center. When you issue a defense command, it flows back through the same channel to the client shield. At no point does any event data touch HYVE's infrastructure.

About Post-Quantum Encryption

The channel between each raptor-shield and your raptor-command server uses ML-KEM-768 — a NIST-standardized post-quantum key encapsulation algorithm — applied at the application layer. This protects your threat intelligence against adversaries who may record traffic today and attempt to decrypt it in the future using quantum computers. Standard transport security is not sufficient for this threat model. Raptor addresses it.

02 Pre-Installation Checklist

Before running the installer, verify the following requirements are met on your server machine. Skipping this step is the most common cause of failed installations.

Server Requirements

REQUIREMENT	MINIMUM	RECOMMENDED
Operating System	Windows 10 / Windows Server 2016	Windows 11 / Windows Server 2022
RAM	8 GB	16 GB (required for Alpha AI model)
Disk Space	5 GB free	20 GB free (Alpha model is ~2 GB; logs grow over time)
CPU	4-core x64	8-core x64 (Alpha inference is CPU-bound)
Architecture	x86-64 (AMD64) only	
Internet (install only)	Required during initial install to download Ollama and the Alpha AI model. Not required for ongoing operation.	

Ports That Must Be Free

Run the following in an elevated Command Prompt before installing to confirm no other application is already using these ports:

```
netstat -an | findstr "8080 8443 9443"
```

The output should be empty. If any port shows as `LISTENING`, identify and stop the conflicting application before proceeding.

PORT	PURPOSE	EXPOSE TO INTERNET?
8080	Dashboard REST API (Command Center connects here)	No — localhost / LAN only
8443	Shield-to-server encrypted channel	Yes — if clients are outside your LAN
9443	License validation	Yes — if clients are outside your LAN
11434	Ollama / Alpha AI (local inference)	Never — localhost only

Firewall Configuration

If your clients connect from outside your local network, open ports **8443** and **9443** inbound on your server's firewall. Port 8080 should *never* be exposed externally — it is the dashboard API and should only be accessible from your operator workstation.

> **Windows Firewall**
> The installer will prompt for administrator access (UAC). This is required to register the firewall rules. If you deny the UAC prompt the server will still start but firewall rules will not be created — you will need to add them manually.

What Gets Installed and Where

ITEM	LOCATION
Server binaries	`bin\` folder in your Raptor package
Runtime data, logs, and certificates	`%APPDATA%\HyveRaptor\`
License keys	`%APPDATA%\HyveRaptor\license-keys.txt`
Shield identity vault	`%ProgramData%\HyveRaptor\vault\`
Shield service log	`%ProgramData%\HyveRaptor\raptor-shield.log`
Ollama engine	`%LOCALAPPDATA%\Programs\Ollama\`
Alpha AI model	Managed by Ollama (typically `%USERPROFILE%\.ollama\models\`)

03 Server Installation

The server installation covers the two backend processes (raptor-command and hyve-backend) as well as the Ollama AI engine that powers Alpha Advisor. This is a one-time procedure on your server machine.

1

Double-click "1 - Install Server.exe"

Right-click is not required — the installer handles elevation automatically. A Windows UAC prompt will appear asking for administrator access. Click **Yes** to proceed. If you click No, the installation will not complete.

2

The installer configures your environment

Behind the scenes the installer: generates a unique cryptographic signing secret for your deployment, creates the required data directories, starts the raptor-command and hyve-backend processes, then begins the Ollama and Alpha AI model setup.

3

Ollama & Alpha AI setup

If Ollama is not already installed, a dialog will appear explaining that it will be downloaded (~50 MB installer) along with the Alpha AI model (~2 GB). Click **OK** to proceed. If you click Cancel, Alpha Advisor will not be available. You can always re-run the installer to set it up later.

4

Alpha model download confirmation

A second dialog will appear confirming that the AI model is downloading in the background. Click **OK**. The server starts immediately and is fully operational. Alpha Advisor will show "DOWNLOADING MODEL" in its status bar until the download completes (5–20 minutes depending on your connection). You do not need to wait.

5

Confirm the server is running

A final confirmation dialog appears: *"HYVE Raptor is running."* This confirms both server processes started successfully. Your operator dashboard is now live at `http://localhost:8080` and your shield endpoint is accepting connections on port 8443.

Verifying the installation independently

Double-click
4 - Check Status.exe
at any time to see which processes are running, their PIDs, and which license keys are currently active. This is the quickest way to confirm everything is healthy without opening the Command Center.

What Happens if the Installer Fails

Partial start — one process failed

If the confirmation dialog says "raptor-command started OK but hyve-backend failed" (or vice versa), check the error log. The dialog will tell you exactly which log file to look at. The most common cause is a port conflict. Resolve the conflict (see Section 14) and run **3 - Restart Server.exe**.

Both processes failed

Check logs in `%APPDATA%\HyveRaptor\logs\`. The files `hyve-backend-err.log` and `raptor-command-err.log` will contain the specific error. Both files are plain text and can be opened in Notepad.

Ollama download failed

This does not prevent the server from starting. HYVE Raptor is fully operational without Alpha Advisor. The error dialog will acknowledge the failure. Alpha Advisor will show "OFFLINE" in the Command Center. To retry, download Ollama manually from **ollama.com**, install it, then re-run the Raptor installer.

Subsequent server restarts

After the initial installation, use
3 - Restart Server.exe
to restart the server — not the install script. The install script is safe to re-run if needed, but the restart script is faster and the correct tool for day-to-day operations.

04 Opening the Command Center

The Raptor Command Center is your primary operator workstation — the desktop application from which you monitor all connected shields, review threats, issue defense commands, check compliance, and consult Alpha.

Launching the Application

Double-click **Raptor Command Center.exe**. The application will open directly to the Dashboard tab. The Command Center connects to your raptor-command server on port 8080.

> **Server must be running first**
>
> The Command Center connects to the raptor-command server on port 8080. If the server is not running, the application will attempt to start it automatically. If it cannot find the server binary, it will display an error dialog. Run
> **1 - Install Server.exe**
> if you have not already done so.

What You See When It Opens

The interface is a dark-themed desktop application with seven tabs along the left navigation bar. Each tab represents a distinct operational function. The status bar at the bottom shows server connection state, active shield count, and Alpha Advisor status at all times.

First Launch Notes

- If this is your first install, the Dashboard will show zero shields and zero events — this is correct. Shields appear as clients install raptor-shield and connect.
- The Compliance tab will show failing controls on a fresh deployment. This is expected. Controls are satisfied as shields connect and as you issue defense commands.
- Alpha Advisor will show "DOWNLOADING MODEL" if the Ollama model is still downloading. This can take 5–20 minutes on a first install. It will not block any other functionality.

Running the Command Center Alongside an Existing Server

If you started the server with **1 - Install Server.exe** and then open the Command Center, it will detect the already-running server and connect to it without launching a second instance. You will never have a port conflict from opening the Command Center while the server is running.

Multiple operator workstations

The Command Center connects to port 8080 on localhost by default. If your operator workstation is a different machine from your server, the Command Center can connect to the server over the network. Open
Settings
in the Command Center and update the server address to your server's IP or hostname.

05 License Key Management

License keys control which shield deployments are authorized to connect to your server. You are the licensing authority — HYVE has no involvement in key generation, distribution, or revocation. Keys are generated locally and stored locally.

Generating a New Key

1

Double-click "2 - Generate License Key.exe"

A dialog appears showing your newly generated key in the format `HYVE-XXXX-XXXX-XXXX`. The key is generated using a cryptographically secure random number generator on your machine.

2

Copy the key and store it

Copy the key from the dialog. Keep a record of which key was issued to which client — Raptor does not store this association. The key file at `%APPDATA%\HyveRaptor\license-keys.txt` stores all active keys as a comma-separated list.

3

Restart the server to load the new key

Run **3 - Restart Server.exe**. The server reads the key file on startup. A key that exists in the file will not be recognized until the server has been restarted with it present.

4

Deliver the key to your client

Send the client: (1) the license key, (2) your server's IP address or hostname, (3) the shield installer package. These three pieces are all they need to install and connect.

Revoking a Key

To revoke a key and disconnect a shield deployment:

1. Open `%APPDATA%\HyveRaptor\license-keys.txt` in Notepad.
2. Delete the key you want to revoke from the comma-separated list.
3. Save the file.
4. Run **3 - Restart Server.exe**.

After the restart, any shield using the revoked key will fail authentication and stop reporting. It will not receive defense commands. The revocation takes effect the next time the shield attempts to authenticate.

> **Key file format**
>
> The license-keys.txt file must contain keys as comma-separated values with no line breaks, e.g.: `HYVE-3A9F-C214-88B7,HYVE-7B2E-F31A-09CD`. Adding extra spaces, line breaks, or characters other than the keys and commas will cause the server to fail to load keys correctly.

How Many Keys Can I Generate?

There is no hard limit on the number of keys you can generate. Your license plan determines how many simultaneous shield connections are authorized, not the number of keys. Generate one unique key per client organization — this allows you to revoke individual clients without affecting others.

Viewing Active Keys

Double-click **4 - Check Status.exe** to see all currently loaded license keys alongside the running process status.

06 Client Shield Installation

Raptor Shield runs on your clients' machines as a Windows service. Once installed it starts automatically on every boot, monitors AI agents continuously, and sends encrypted threat events to your server without any user interaction required.

What to Give Your Client

Before a client can install the shield, they need three things from you:

Their License Key

The HYVE-XXXX-XXXX-XXXX key you generated for them in Section 5.

Your Server Address

Your server's IP address or hostname, including the port. Example: `http://203.0.113.42:8443`

The Shield Installer

The file **6 - Install Shield (Client).exe** from the Raptor package, or the full Raptor package folder including the `bin\` directory.

Client Installation Steps

1 **Double-click "6 - Install Shield (Client).exe"**

A setup window appears with three fields. A UAC prompt will appear — the client must click Yes to allow the service to be installed. Administrator access is required only for this one-time installation step.

2 **Fill in the setup form**

FIELD	WHAT TO ENTER
Server Address	Your server's address, e.g. `http://203.0.113.42:8443`
License Key	Their assigned key, e.g. `HYVE-3A9F-C214-88B7`
Shield Name	A label for this machine. Defaults to the computer name. This is what appears in your Command Center.

3

Click Install

The shield binary is copied to the system, the Windows service **RaptorShield** is registered, and the service starts immediately. A success dialog confirms the installation.

4

Verify the shield appears in your Command Center

Within 30 seconds of installation, the shield will authenticate against your license server, establish an encrypted connection to raptor-command, and appear in the **Shields** tab of your Command Center with a green heartbeat indicator.

Testing on a single machine

If you are testing on the same machine where the server is running, use `http://localhost:8443` as the server address. The server and shield can run on the same machine.

What the Shield Service Does

Once installed, raptor-shield runs as a Windows service named **RaptorShield**. It:

- Starts automatically on every Windows boot — no user login required
- Runs as the LocalSystem account with no interactive session
- Generates and stores a unique cryptographic identity key the first time it runs, protected by Windows DPAPI machine-scoped encryption
- Authenticates against your license server on startup using the assigned license key
- Maintains a persistent encrypted connection to your raptor-command server
- Scans AI agent behavior continuously and transmits anomalies as structured threat events
- Listens for and executes defense commands issued from your Command Center

Updating the Shield on a Client Machine

When a new version of raptor-shield is available, the client runs **7 - Restart Shield Service.exe** after the updated binary has been placed in the installation directory. This stops the service, replaces the binary, and restarts it. No re-configuration is required.

07 Command Center Reference

The Raptor Command Center has seven tabs. Each tab is a distinct operational module. Here is what each one contains and how to use it.

DASHBOARD

Live operational overview

The first thing you see on launch. Provides at-a-glance situational awareness across your entire network.

- Total shields connected / active count
- Threat event counts by severity (critical, high, medium, low)
- Overall compliance score across all frameworks
- Recent threat events list with timestamp and severity
- Active defense commands in progress

SHIELD AGENTS

All connected shields

A live list of every raptor-shield deployment that is currently authenticated and reporting to your server.

- Shield name, machine hostname, and IP address
- Last heartbeat timestamp — turns red if a shield goes silent
- Current status: Active, Degraded, Isolated, Locked Down
- Threat event count per shield
- Click any shield to view its full event history

THREAT FEED

Live attack event stream

A real-time stream of all threat events received from all shields, ordered by most recent. Each event includes:

- Severity score and color-coded level
- Attack vector classification
- Affected AI agent identifier
- Description of the behavioral anomaly detected
- Timestamp and originating shield
- Recommended defense action

DEFENSE COMMAND

Issue commands to any shield

Select one or more shields and issue a defense command. Commands are encrypted and delivered to the target shield within 100ms. See Section 8 for full command descriptions.

- ISOLATE — cut network access for the affected agent
- CONTAIN — restrict agent to a behavioral sandbox
- NEUTRALIZE — terminate the agent process
- LOCKDOWN — full endpoint quarantine

COMPLIANCE

Six-framework compliance posture

Real-time scoring across six regulatory frameworks simultaneously. Failing controls surface with evidence and remediation context.

- CMMC 2.0, HIPAA, FedRAMP, CJIS, NIST 800-171, SOC 2
- Pass / Fail status per required control
- Evidence: what data satisfied or failed the control
- Remediation guidance for each failing control

ALPHA ADVISOR

Offline AI defense analyst

An embedded AI advisor with live access to your entire system state. See Section 9 for full documentation. Alpha is not a generic chatbot — it knows your network specifically.

- Ask about specific threats, incidents, or shields
- Request compliance remediation paths
- Incident response guidance with your actual event data
- Runs 100% offline — no data leaves your machine

SETTINGS

Platform configuration

Server connection settings, display preferences, and cryptographic information.

- Server address and port configuration
- Alpha Advisor status and model refresh
- Encryption information (your ML-KEM key fingerprint)
- Log file paths

08 Defense Commands

Defense commands are operator-initiated actions dispatched from the Command Center to any connected shield. Each command is encrypted with the shield's unique session key and delivered in under 100ms. The shield executes the command autonomously — no user action is required on the client machine.

Commands escalate in severity. ISOLATE is reversible. LOCKDOWN is the most aggressive response and should be reserved for confirmed compromise scenarios.

ISOLATE

Effect: Cuts the target AI agent's network access. The agent process continues running but all outbound connections from it are blocked.

Use when: You see a suspicious network destination in the threat feed and want to contain the activity without disrupting the client's other operations.

Reversible: Yes — issue a RELEASE command to restore network access.

CONTAIN

Effect: Restricts the agent to a behavioral sandbox — limits its API call scope, reduces its capability set, and flags all subsequent activity for elevated scrutiny.

Use when: Capability drift is detected and you want to restrict the agent while you investigate, without stopping it entirely.

Reversible: Yes — issue a RELEASE command to restore normal operation.

NEUTRALIZE

Effect: Terminates the target AI agent process immediately. The agent stops running. The shield continues monitoring the endpoint and reporting events.

Use when: You have confirmed malicious behavior and want to stop the agent without locking down the entire endpoint.

Reversible: The command cannot be undone, but the agent can be restarted by the client if appropriate.

LOCKDOWN

Effect: Full endpoint quarantine. Network access is severed, all AI agent processes are terminated, and the shield enters a restricted reporting-only mode.

Use when: You have confirmed active compromise and need to immediately stop all activity on the endpoint while you coordinate a response.

Reversible: Requires deliberate manual release from the Command Center after the situation is resolved.

Command confirmation

NEUTRALIZE and LOCKDOWN commands require a confirmation step in the Command Center before they are dispatched. ISOLATE and CONTAIN execute immediately. This is intentional — NEUTRALIZE and

LOCKDOWN have operational impact on the client machine and should be deliberate decisions.

Issuing a Command

1

Navigate to the Defense Command tab

The tab shows all connected shields. Each shield displays its name, current status, and a row of command buttons.

2

Select the target shield

Click the shield you want to act on. You can select multiple shields to issue a command to all of them simultaneously.

3

Click the command button

Click ISOLATE, CONTAIN, NEUTRALIZE, or LOCKDOWN. For NEUTRALIZE and LOCKDOWN, a confirmation dialog appears. Confirm the command. The shield status updates in the Shields tab within seconds.

Command History

All defense commands are logged in the threat feed with a timestamp, the issuing operator, the target shield, and the command type. This creates an immutable audit trail of all defense actions taken through the platform.

09 Alpha Advisor

Alpha is HYVE Raptor's embedded AI analyst. It is categorically different from AI assistants you may have used elsewhere, and understanding this difference is important for using it effectively.

A general AI assistant — general-purpose chatbots — answers questions based on training data and whatever context you manually provide in the conversation. It knows about security in general. It knows nothing about your network.

Alpha knows your network. Before answering any question, Alpha reads your live system state: every shield currently connected, every threat event in your database, every compliance control that is passing or failing right now, the behavioral history of every AI agent your shields have reported on. When you ask Alpha a question, it is not reasoning from a description of your environment. It is reasoning from your environment itself.

What Alpha Can Do

Incident analysis

Ask Alpha to walk you through what happened during a specific event or time period. It will pull the relevant events, identify the behavioral pattern, describe the probable attack vector, and tell you what the shield did in response — using your actual data, not a hypothetical template.

Compliance remediation

Ask Alpha which compliance controls are currently failing and what you need to do to fix them. It reads your live compliance engine output and gives you specific, prioritized remediation steps based on your actual posture — not a generic checklist.

Threat prioritization

When you have dozens of events across multiple shields, ask Alpha which ones matter most. It cross-references severity scores, attack vector classifications, and historical behavior across your entire network to surface the events that warrant your immediate attention.

Operational guidance

Ask Alpha which defense command to issue for a given situation. It knows the event, the shield's history, and the available response options — and it will recommend the least disruptive appropriate action.

Cross-client pattern recognition

If you manage multiple shield deployments, Alpha can identify patterns that span clients — a new attack vector appearing across multiple shields, similar behavioral drift in agents that serve the same industry — that would be invisible to an operator reviewing dashboards one at a time.

Alpha Status Indicators

STATUS	WHAT IT MEANS	WHAT TO DO
ONLINE	Alpha is ready. Type your question.	Nothing required.
DOWNLOADING MODEL	The AI model is still downloading from the initial install. The download started when you ran the server installer and continues in the background.	Wait 5–20 minutes. Click REFRESH in the Alpha tab once the download completes.
THINKING	Alpha is generating a response. A blinking indicator and elapsed-seconds counter confirm it is working.	Wait for the response. Click STOP if you want to cancel.
OFFLINE	Ollama is not running on the server machine.	Re-run **1 - Install Server.exe** on the server machine. Ollama will be detected and started automatically.

Important: Alpha Is Fully Offline

Every query you send to Alpha, and every response it generates, stays on your machine. Alpha uses a locally running language model (powered by Ollama). No question, no threat data, no compliance gap, no incident detail is ever sent to any external server — including HYVE's. This is not a configuration option. It is how the system is built. Alpha works in air-gapped environments with no internet access, and it works identically in classified facilities where outbound traffic is prohibited or monitored.

> **Getting the best answers from Alpha**
>
> Alpha already knows your system state, so you do not need to describe it. The most effective questions are direct operational queries: "What are my three highest-priority threats right now?", "Which shields have gone quiet in the last hour?", "Walk me through the 14:32 event on Shield-Atlanta", "What do I need to fix to pass CMMC Level 2?". Treat Alpha as a senior analyst who has read every alert — ask it operational questions, not general ones.

Alpha During Model Download

The Alpha AI model is approximately 2 GB and downloads in the background when you first run the server installer. During this time, Alpha shows "DOWNLOADING MODEL" and will not respond to queries. This is normal and expected. The rest of the platform — threat monitoring, compliance scoring, defense

commands — is fully operational during the model download. Alpha is the only feature that requires the download to complete.

10 Compliance Engine

HYVE Raptor includes a real-time compliance posture engine that scores your deployment against six regulatory frameworks simultaneously. The engine evaluates your actual system state — connected shields, issued defense commands, cryptographic identity management, incident response actions — and maps it to specific required controls in each framework.

Supported Frameworks

FRAMEWORK	RELEVANCE TO AI AGENT SECURITY
CMMC 2.0	AI agents processing CUI on defense contractor networks. Level 2 requires documented incident response and access control.
HIPAA	AI agents with access to electronic PHI. The zero data access architecture means PHI-adjacent signals never leave the covered entity's environment.
FedRAMP	AI tools deployed in federal cloud environments. On-premise deployment eliminates third-party data routing concerns.
CJIS	AI agents with access to criminal justice information systems. Raptor's behavioral monitoring satisfies CJIS audit and access control requirements.
NIST 800-171	Protection of Controlled Unclassified Information in non-federal systems. Directly relevant to defense contractors and research institutions.
SOC 2	Security, availability, and confidentiality controls for service organizations using AI tools in client-facing workflows.

How Controls Are Evaluated

Each required control is evaluated against live platform data. Examples:

- **Agent Identity Management:** Satisfied when raptor-shield has established a verified cryptographic identity with the command server.
- **Incident Response Plan:** Satisfied when at least one defense command has been issued — demonstrating an active response capability.

- **Cryptographic Key Management:** Satisfied when agent identity keys are present, protected, and verified during the session handshake.
- **Multi-Factor Authentication:** Evaluated based on your authentication configuration in Settings.

Reading the Compliance Tab

Each framework shows an overall pass percentage and a list of individual controls. Each control has one of three states:

- `PASS` — Control is satisfied based on current system evidence.
- `FAIL` — Control is required but not currently satisfied. Evidence and remediation guidance are shown.
- `N/A` — Control does not apply to your current deployment configuration.

> **Compliance posture is dynamic**
>
> Compliance scores update in real time as shields connect, as defense commands are issued, and as system configuration changes. A fresh deployment will show several failing controls — this is expected. As you connect shields and use the platform, controls are satisfied automatically. Use Alpha Advisor to get a prioritized remediation plan for any remaining gaps.

11 Launcher Files Reference

Your Raptor package includes seven numbered launcher files. Each one is a self-contained executable that performs a specific operation. Double-click to run. All launchers that modify system state will request UAC elevation.

FILE	WHAT IT DOES	WHEN TO USE IT
1 - Install Server.exe	First-time installation. Starts both server processes, sets up Ollama, begins the Alpha model download.	First setup, or if both server processes need to be fully re-initialized.
2 - Generate License Key.exe	Generates one new HYVE-XXXX-XXXX-XXXX license key and saves it to the key file.	Before onboarding each new client.
3 - Restart Server.exe	Stops and restarts both server processes, reloading all configuration including any new license keys.	After generating a new key. After any configuration change. Routine maintenance.
4 - Check Status.exe	Shows running process status (with PIDs), all active license keys, and the log folder path.	Verifying the server is healthy. Confirming a key loaded correctly.
5 - Stop Server.exe	Stops both server processes and any associated services.	Planned maintenance, server shutdown, or before moving the installation.
6 - Install Shield (Client).exe	GUI installer for raptor-shield on a client machine. Requires server address, license key, and shield name.	On each client machine during onboarding.
7 - Restart Shield Service.exe	Stops and restarts the RaptorShield Windows service on the client machine.	After updating the shield binary. If the shield stops reporting unexpectedly.

12 Ports & Network Reference

PORT	PROTOCOL	PURPOSE	EXPOSE EXTERNALLY?
8080	HTTP	Dashboard REST API. The Command Center connects here to fetch threat data, shields, compliance, and issue commands.	No. LAN / localhost only. Never expose to the internet.
8443	TCP / ML-KEM	Shield-to-server encrypted channel. All threat events and defense commands flow through this port.	Yes, if clients connect from outside your LAN. Open inbound in your firewall.
9443	HTTPS	License validation. Shields authenticate here on first connection and periodically thereafter.	Yes, if clients connect from outside your LAN. Open inbound in your firewall.
11434	HTTP	Ollama local AI engine. Used internally by the Alpha Advisor only.	Never. Localhost only. Exposing this port is a security risk.

13 Data & Privacy Reference

DATA TYPE	WHERE IT LIVES	WHO CAN ACCESS IT
Threat event database	`%APPDATA%\HyveRaptor\` on your server	You — the operator. HYVE has no access.
License keys	`%APPDATA%\HyveRaptor\license-keys.txt`	You only. File is restricted to the running user account.
Shield identity keys	`%ProgramData%\HyveRaptor\vault\` on the client machine	The RaptorShield service only. Protected by Windows DPAPI machine-scope encryption.
Alpha Advisor conversations	In-memory only — never written to disk	You — the operator. Never transmitted anywhere.
Client threat metadata	Your raptor-command database	You — the operator. Architectural separation prevents you from reading client plaintext data.

DATA TYPE	WHERE IT LIVES	WHO CAN ACCESS IT
Logs	`%APPDATA%\HyveRaptor\logs\` (server) and `%ProgramData%\HyveRaptor\raptor-shield.log` (client)	Operator and client administrator respectively.

14 Troubleshooting — Every Known Failure Path

This section enumerates every failure mode observed or anticipated in HYVE Raptor deployment, with a resolution path for each. It is organized by which component the symptom presents on. Begin at the symptom you are seeing.

Diagnostic command you will use repeatedly

Every launcher writes to `%APPDATA%\HyveRaptor\logs\`. The two most informative log files are `raptor-command-err.log` (server) and `hyve-backend-err.log` (license authority). On the client side, the shield service logs to `%ProgramData%\HyveRaptor\raptor-shield.log`. When in doubt, open these first.

14.1 — Installation Failures

Q — "1 - Install Server.exe" fails with "access is denied" or "UAC denied".

Cause: The installer requires administrator privileges to write to `%ProgramData%`, register Windows Firewall rules, and (on shield installs) register a Windows service.

Resolution: Right-click `1 - Install Server.exe` → "Run as administrator". When the UAC dialog appears, click **Yes**. If UAC is disabled on the machine, enable it temporarily via the Control Panel → User Accounts, or run PowerShell as Administrator and execute `Deploy-HyveRaptor.ps1` directly.

Q — "PowerShell execution policy prevents this script from running."

Cause: Group Policy or local hardening has set the execution policy to Restricted or AllSigned. The bundled launchers set execution policy Bypass inline, but some domain-managed machines override this.

Resolution: Open an elevated PowerShell session and run:

```
Set-ExecutionPolicy -Scope Process -ExecutionPolicy Bypass -Force
cd "C:\path\to\Raptor"
.\Deploy-HyveRaptor.ps1
```

This sets the policy only for the current process, bypasses the launcher, and produces identical behavior.

Q — Installer reports "hyve-backend.exe not found in bin/".

Cause: The distribution ZIP was not fully extracted before running the installer, or the `bin/` folder was moved/renamed.

Resolution: Re-extract `hyve-raptor-package-v1.0.0.zip` to a fresh folder. Do not run the launchers from inside a ZIP preview — Windows does not extract bundled folders in that mode. Confirm the folder structure matches:

```
Raptor/
├── 1 - Install Server.exe
├── bin/
│   ├── hyve-backend.exe
│   ├── raptor-command.exe
│   └── raptor-shield.exe
├── alpha-models/
└── ...
```

Q — Windows Defender / SmartScreen blocks the installer ("Windows protected your PC").

Cause: The bundled EXEs are not signed with an extended-validation code-signing certificate (common for unsigned or freshly-signed software until sufficient reputation builds).

Resolution: On the SmartScreen dialog, click **More info** → **Run anyway**. If your organization has disabled this option via Group Policy, your IT team must whitelist the installer EXEs. Confirm integrity before whitelisting by verifying the SHA-256 hashes of the binaries against the hashes documented in your delivery email.

Q — Antivirus flags raptor-shield.exe or hyve-backend.exe and quarantines them.

Cause: Heuristic scanners often flag binaries that open raw network sockets, modify firewall rules, or terminate processes — all of which the shield and defense modules legitimately do. Raptor's defense commands use documented Windows APIs (netsh, TerminateProcess, sc stop).

Resolution: Whitelist the Raptor installation directory in your AV. For enterprise AV (CrowdStrike, SentinelOne, Defender for Endpoint), add exclusions for:

- `%ProgramFiles%\HyveRaptor\` (shield install location)
- `%APPDATA%\HyveRaptor\` (server runtime data)
- `%ProgramData%\HyveRaptor\` (shield vault + forensic)

Add the binaries by name to the process-allowlist: `raptor-shield.exe`, `raptor-command.exe`, `hyve-backend.exe`.

Q — Installer completes but no confirmation dialog appears.

Cause: PowerShell ISE execution, or the processes exited before the post-start probe fires.

Resolution: Run `4 - Check Status.exe`. If both processes show as running, the installation succeeded and the dialog was suppressed. If either shows as Stopped, open `%APPDATA%\HyveRaptor\logs\raptor-command-err.log` for the exit reason.

14.2 — Server Runtime Failures

Q — Server starts, then crashes within seconds. Error log shows "listen tcp :8443: bind: address already in use".

Cause: Another process is already bound to port 8443, 8080, or 9443.

Resolution: Identify the conflicting process:

```
netstat -ano | findstr ":8443 :8080 :9443"
```

The last column is the PID. Look up the process name:

```
tasklist /FI "PID eq <PID>"
```

Stop it (`taskkill /PID <PID> /F`) or change the Raptor listen ports via environment variables: `RAPTOR_LISTEN`, `RAPTOR_API_LISTEN`, and the hyve-backend flag `--listen :<new-port>`.

Q — Server crashes with "database is locked" or "disk I/O error".

Cause: SQLite database file is corrupted, or two raptor-command processes are trying to open the same DB file simultaneously, or Windows antivirus is holding the file open during scanning.

Resolution:

1. Run `5 - Stop Server.exe` and wait 10 seconds.
2. Check for orphaned processes: `tasklist | findstr raptor-command`. Kill any that remain: `taskkill /IM raptor-command.exe /F`.
3. If the database file is corrupted (rare), rename it: `%APPDATA%\HyveRaptor\raptor-command.db` → `.db.corrupted`. The next server start creates a fresh DB. Historical events will be lost but license keys and shield identities are preserved (they live in separate files).

4. Restart: `3 - Restart Server.exe`.

Q — Server runs but memory usage grows steadily until the machine swaps.

Cause: High-volume threat event ingestion with a large number of connected shields can cause the in-memory event buffer to grow faster than disk writes flush it. Normal for a server with >50 shields producing continuous events.

Resolution: Restart the server nightly via a Windows Scheduled Task invoking `3 - Restart Server.exe`. For Sentinel and Command deployments, consider increasing server RAM to 32 GB+ or adding the `RAPTOR_EVENT_BUFFER_SIZE` environment variable (default 10000, reduce to 2000 for tighter memory pressure).

Q — Server log shows repeated "ML-KEM decapsulation failed" errors.

Cause: A shield is presenting a mismatched session key. This happens when a shield's vault is corrupted, or when a clone of a shield identity is attempting to connect (two machines using the same keypair).

Resolution: Identify the offending shield ID in the log. On that client machine, stop the RaptorShield service, delete `%ProgramData%\HyveRaptor\vault\`, and restart the service — this forces fresh key generation. If the issue persists, re-run `6 - Install Shield (Client).exe` on that machine to reset its identity completely.

14.3 — Shield Installation & Runtime

Q — "6 - Install Shield (Client).exe" completes but the shield service does not start.

Resolution: Open `services.msc`, locate **RaptorShield**. If it is listed but not running, double-click → **Start**. If it refuses to start, check the service error in Event Viewer → Windows Logs → Application.

Common service-start errors:

- **Error 1067** ("process terminated unexpectedly") — check `%ProgramData%\HyveRaptor\raptor-shield.log` for the underlying exception.
- **Error 5** ("access denied") — the LocalSystem account lacks access to `%ProgramData%\HyveRaptor\`. Reset folder permissions with `icacls "%ProgramData%\HyveRaptor" /reset /T /C`.
- **Error 1053** ("did not respond to start or control request in a timely fashion") — the shield binary has an initialization bug or is blocked by AV. Check the log; whitelist in AV if applicable.

Q — Shield service starts but logs "identity load failed: DPAPI decrypt".

Cause: The shield's cryptographic identity vault was encrypted on a different Windows machine (e.g., cloned VM image). DPAPI machine-scope encryption is deliberately tied to the specific machine.

Resolution: Delete `%ProgramData%\HyveRaptor\vault\` and restart the RaptorShield service. The shield will generate a fresh keypair and re-register with the command server on its next heartbeat.

Q — Shield appears in the Command Center for a few minutes, then drops off permanently.

Resolution: Three likely causes:

1. **License key revoked.** The shield authenticates on every heartbeat. If the key was removed from `%APPDATA%\HyveRaptor\license-keys.txt` on the server, the shield is disconnected on next heartbeat. Verify the key is present and the server has been restarted.
2. **Network firewall timeout.** Stateful firewalls sometimes drop idle long-poll connections. Configure the firewall to allow connections idle up to 90 seconds. Or, add `RAPTOR_POLL_TIMEOUT=30` to the client environment to force shorter polls.
3. **Shield service crashed.** Check `%ProgramData%\HyveRaptor\raptor-shield.log` on the client for panic traces. Restart via `7 - Restart Shield Service.exe`.

Q — Shield log shows "handshake failed: invalid license".

Cause: The license key entered during shield install does not match any key on the server's `license-keys.txt`, or the server was not restarted after the key was added.

Resolution: On the server, open `%APPDATA%\HyveRaptor\license-keys.txt` and confirm the key is present exactly as delivered (format `HYVE-XXXX-XXXX-XXXX`, no trailing whitespace). Run `3 - Restart Server.exe`. Run `7 - Restart Shield Service.exe` on the client.

Q — Shield connects but never emits threat events, even with AI agents running on the machine.

Cause: The shield's recon discovery did not classify any running process as an AI agent. Discovery looks for framework markers (langchain, autogen, crewai, openai, ollama, llamaindex, vllm, plus explicit HYVE_AGENT_ID tagging) in the process name or command line, OR an explicit `HYVE_AGENT_ID` tag in `%ProgramData%\HyveRaptor\agent-tags.txt`.

Resolution: To explicitly tag a process as a monitored agent, create `%ProgramData%\HyveRaptor\agent-tags.txt` with one line per tagged process in the format:

```
12345   my-production-agent
67890   another-agent
```

Where 12345/67890 are PIDs and the second column is the agent ID. Restart the shield with `7 - Restart Shield Service.exe`.

14.4 — Command Center Desktop

Q — Raptor Command Center.exe shows a blank white window on launch.

Cause: The Electron renderer could not reach the raptor-command REST API on port 8080 at startup. Most commonly, the server was not running when the Command Center opened.

Resolution: Close the Command Center. Run `1 - Install Server.exe` first (or `3 - Restart Server.exe` if already installed). Wait for the "HYVE Raptor is running" dialog. Then launch Raptor Command Center.exe.

Q — Command Center shows "raptor-command exited with code 1".

Cause: The Command Center attempted to auto-launch a bundled raptor-command instance, but port 8080 is already in use by an independently-running server (from `1 - Install Server.exe`).

Resolution: This is cosmetic — the independently-running server is actually fine. Close the error dialog and re-open Command Center. It will now detect the existing server and connect to it instead of trying to start its own.

Q — Dashboard displays, but metrics are all zero and no shields are visible.

Cause: Either no shields have connected yet, or the Command Center is pointed at the wrong server address.

Resolution: Open the Settings tab and verify the Server Address field shows `http://127.0.0.1:8080` (local) or your remote server IP:8080. If remote, verify network reachability from your workstation to that IP on port 8080.

Q — Issuing a defense command shows success but the target shield doesn't show the effect.

Cause: Commands are delivered asynchronously via long-poll. Delivery can lag up to 30 seconds if the shield is mid-poll-cycle. Also, defense actions require the shield service to have privileges for the specific OS action (firewall modification, process termination).

Resolution: Wait up to 30 seconds for delivery confirmation in the Threat Feed tab. Check the shield log (`%ProgramData%\HyveRaptor\raptor-shield.log`) for the command execution trace. If the log shows "defense action FAILED" with an error, the RaptorShield service may lack the required privilege — ensure it runs as LocalSystem (default) and not a limited user account.

14.5 — Alpha Advisor

Q — Alpha shows "OFFLINE" status.

Cause: Ollama is not running on the server machine, or not reachable at `http://localhost:11434`.

Resolution: Open an elevated Command Prompt and run:

```
ollama serve
```

Leave that window open; the Alpha tab should turn ONLINE within 5 seconds (click REFRESH). For a permanent fix, set Ollama to auto-start by running its installer or adding `ollama serve` to a Windows Scheduled Task at logon.

Q — Alpha shows "NO MODELS" even after first-run setup completed.

Cause: The tier-aware installer's `ollama pull` was interrupted — power loss, network disconnect, or manual cancellation during download.

Resolution: Identify your tier's base model (Scout = llama3.2:3b, Guardian = llama3.1:8b, Sentinel = llama3.3:70b, Command = llama3.1:405b) and resume the pull:

```
ollama pull llama3.2:3b
```

Then build the HYVE Alpha profile:

```
cd "C:\path\to\Raptor\alpha-models"
ollama create hyve-alpha-3 -f Modelfile.alpha-3
```

Click REFRESH in the Alpha tab.

Q — Alpha-70 or Alpha-405 fails to load with "out of memory" error.

Cause: The selected Alpha profile exceeds the hardware capacity of this machine. See Section 9 and the product page for the real hardware requirements per tier.

Resolution: Fall back to a smaller profile. For a Sentinel or Command customer running on a commodity server, use Alpha-8 (16 GB RAM is enough):

```
ollama create hyve-alpha-8 -f alpha-models\Modelfile.alpha-8
ollama pull llama3.1:8b
```

Then set the Command Center's Alpha model selection to `hyve-alpha-8` in the Settings tab. The smaller profile is still fully functional — only the reasoning depth on complex queries is reduced.

Q — Alpha responds but answers reference shields or events that don't exist.

Cause: LLM hallucination. The model is constructing plausible-looking identifiers from the general pattern rather than from the live context.

Resolution: This is a known limitation of any open-weight LLM. Cross-check Alpha's answers against the live Threat Feed and Shields tabs. Ask Alpha to "cite the specific event ID" — it will either produce a real one or acknowledge uncertainty. Do not treat Alpha output as gospel for compliance or incident-report documentation; use it as an analyst draft.

Q — Alpha responses take over 2 minutes, even on a fast machine.

Cause: Ollama is running on CPU-only inference, or the selected profile is too large for the available hardware and is spilling to disk (swapping).

Resolution: Check GPU acceleration with:

```
ollama ps
```

Look for "GPU" in the output. If only "CPU" is shown, install the NVIDIA CUDA toolkit (or AMD ROCm on supported platforms) and restart Ollama. Alternatively, downgrade to a smaller Alpha profile whose weights fit entirely in system RAM without swapping.

14.6 — License Keys

Q — Newly generated key is not recognized by shields.

Always run `3 - Restart Server.exe` after generating a key. The server reads the key file only at startup.

Verify key formatting — `%APPDATA%\HyveRaptor\license-keys.txt` must contain comma-separated keys on a single line with no quotes or extraneous whitespace.

Q — Key file was corrupted or deleted.

Re-run `2 - Generate License Key.exe` for each client you need to re-issue. Maintain a separate out-of-band record of which key belongs to which client — the Raptor platform does not store this association.

Q — I need to revoke a shield's access immediately.

Open `%APPDATA%\HyveRaptor\license-keys.txt`, delete the key, save the file, run `3 - Restart Server.exe`. This terminates all active sessions — the revoked shield will fail re-authentication on its next heartbeat (within 30 seconds).

14.7 — Compliance & MFA

Q — Compliance score dropped unexpectedly overnight.

Cause: Compliance scores reflect live system state. Common drops:

- A shield went offline → HYC-001 (Agent Identity) drops proportionally
- No threat events received in 24+ hours → HYC-005 (Continuous Threat Monitoring) evidence ages
- No defense commands issued recently → HYC-006 (Incident Response Plan) may show as unverified
- Audit log file not written in 24h → HYC-004 fails

Check the Compliance tab for the specific failing control's Evidence line — it tells you exactly which condition is not met.

Q — I enabled MFA but can't log in — code is rejected.

Cause: Most common — clock drift between the server and the authenticator app. TOTP codes are valid for a ~30-second window.

Resolution: Sync your server's clock via `w32tm /resync` (Windows). Try the code at the start of the next 30-second interval. If still failing, your enrollment may not have saved correctly — delete `%APPDATA%\HyveRaptor\mfa.json` and re-enroll via `/api/mfa/enroll`. This will force a fresh QR code.

Q — I lost my phone / authenticator and can't access the Command Center.

Break-glass procedure: On the server machine (where MFA enrollments are stored), open `%APPDATA%\HyveRaptor\mfa.json` in Notepad. Delete your operator entry from the JSON array. Save. Restart the server. MFA gate will now allow enrollment again — re-enroll with a fresh authenticator.

14.8 — Network & Firewall

Q — Client machines outside the LAN can't reach the server on port 8443.

Resolution: Open inbound port 8443 (and 9443 if clients validate licenses directly) on every firewall between the client and the server — Windows Firewall on the server, corporate edge firewall, ISP router, cloud security group. Never open port 8080 externally — that's the operator dashboard API.

Verify end-to-end connectivity from the client:

```
Test-NetConnection -ComputerName your.server.com -Port 8443
Test-NetConnection -ComputerName your.server.com -Port 9443
```

Q — Shield connects via LAN but not from remote office / VPN.

Cause: VPN split-tunneling or DNS mismatch — the server hostname resolves differently inside vs. outside the VPN.

Resolution: Use the server's public IP rather than a hostname in the shield configuration. Or, publish a consistent DNS record accessible from both inside and outside the VPN.

Q — Air-gapped network — no internet access — Alpha won't set up.

Resolution: On a connected machine, install Ollama and pull the desired base model (`ollama pull llama3.1:8b`). Copy the resulting model files from `%USERPROFILE%\.ollama\models\` to the same path on the air-gapped server. Also copy the Ollama binary to the air-gapped machine (from `%LOCALAPPDATA%\Programs\Ollama\`). Start the server, run `ollama create hyve-alpha-X -f Modelfile.alpha-X` against the local Modelfile. Alpha now works fully offline with no external connections.

14.9 — Emergency Recovery

Q — Server machine died — need to restore Raptor on a new server.

Procedure:

1. Install Raptor on the new server via `1 - Install Server.exe`.
2. Stop the new server immediately (`5 - Stop Server.exe`).
3. Restore the `%APPDATA%\HyveRaptor\` folder from your backup of the old server. Critical files: `signing-secret.txt`, `license-keys.txt`, `raptor-command.db`.
4. Update each client shield's server address to point to the new server's IP (re-run `6 - Install Shield (Client).exe` on each client, or edit the shield's config file directly).

5. Start the new server (`3 - Restart Server.exe`). Shields reconnect within minutes.

Q — All compliance evidence missing after a restore.

The compliance evidence is computed live from the event database. If your `raptor-command.db` is fresh, the evidence will rebuild as new shields reconnect and new events arrive. For audit continuity, restore the old DB file per the procedure above.

14.10 — Still Stuck

Q — I've tried everything and still can't get it working.

Gather this information and email it to majixx@vibesoftwaresolutions.com:

1. The specific error message you see (screenshot or exact copy)
2. Output of `4 - Check Status.exe`
3. Contents of `%APPDATA%\HyveRaptor\logs\raptor-command-err.log` (last 200 lines)
4. Contents of `%ProgramData%\HyveRaptor\raptor-shield.log` (last 200 lines, if client-side)
5. Output of `ollama ps` (if an Alpha Advisor issue)
6. Your Windows version, server specs (CPU, RAM), and license tier

Response target: within 48 hours during business hours. Critical incidents (a live threat in progress on a Command-tier deployment) are prioritized — clearly mark the email subject **[CRITICAL]**.

15 Frequently Asked Questions

Installation & Setup

Q — Do I need an internet connection to run HYVE Raptor?

Only during the initial installation, to download the Ollama engine and the Alpha AI model (~2 GB combined). Once installation is complete, HYVE Raptor operates with zero internet connectivity. There is no telemetry to HYVE, no license check over the internet, no update server to reach. The platform can run indefinitely in an air-gapped or classified environment. Alpha Advisor continues to work fully offline once the model is downloaded.

Q — Can I install the server on Linux instead of Windows?

The raptor-command and hyve-backend binaries are cross-compiled for both Windows and Linux. The launcher EXE files are Windows-only convenience wrappers. On a Linux server, start the processes directly from the `bin/` directory. Contact support for Linux-specific deployment documentation. The client-side raptor-shield service is Windows-only in the current release.

Q — Can I move my HYVE Raptor installation to a different server?

Yes. Stop the server with `5 - Stop Server.exe`. Copy the entire Raptor package folder and the contents of `%APPDATA%\HyveRaptor\` to the new machine. Run `1 - Install Server.exe` on the new machine. Your signing secret, license keys, and event database will be carried over. Shields will need to be updated with the new server's IP address — run `6 - Install Shield (Client).exe` on each client with the new address, or update the configuration on each client machine directly.

Q — Does the server installation affect any other software running on my machine?

No existing software is modified. The installer creates a data directory in `%APPDATA%\HyveRaptor\`, starts two background processes, installs Ollama (if not already present), and adds firewall rules. It does not touch your system registry beyond what is required for the processes to run, does not

modify any existing application, and does not change any environment variables that other applications depend on.

Operations & Management

Q – How many shields can be connected at once?

This is determined by your license plan — Scout (10 shields), Guardian (50 shields), Sentinel (250 shields), Command (unlimited). The platform enforces no hard technical limit on simultaneous connections beyond your plan entitlement. If you need to expand, contact Vibe Software Solutions to upgrade your plan.

Q – What happens if my server goes down while shields are still running on client machines?

Each raptor-shield deployment operates independently. If the connection to your server is lost, the shield continues monitoring AI agents on the client machine. Threat events are queued locally and transmitted when the connection is restored. The shield does not stop protecting the client — it just cannot report back or receive defense commands until connectivity is restored. The client will not see any disruption or error on their machine.

Q – Can I have multiple operators using the Command Center at the same time?

Yes. Multiple operators can open the Command Center simultaneously from different machines, all connecting to the same raptor-command server on port 8080. They will see the same live data. Defense commands issued by any operator are visible to all others in the threat feed audit trail. Each operator's machine needs network access to port 8080 on your server — ensure your LAN routing allows this.

Q – How long does HYVE Raptor retain threat events?

Event retention is configured by your plan. Events are stored in a local database on your server. The database uses write-ahead logging for reliability — events are never lost to process crashes or unexpected shutdowns. If disk space becomes a concern, older events can be archived or pruned. Contact support for guidance on event database management for high-volume deployments.

Q – Can I export threat event data for use in a SIEM or other tool?

Yes, via the REST API on port 8080. The API exposes all threat events, shield data, and compliance state in structured JSON format. Sentinel-tier and above plans include webhook delivery — events

are pushed to your configured endpoint in real time as they arrive. Contact support for the API reference documentation.

Q – What AI agents does raptor-shield monitor?

raptor-shield is designed to monitor AI agent processes broadly — it profiles any process that exhibits AI agent behavior characteristics (autonomous API calls, capability declarations, model interactions, external service communication). It is not limited to a specific framework or tool. Agents built on OpenAI, Anthropic, local models, LangChain, AutoGen, and custom stacks are all within scope. Contact support if you have a specific agent runtime you want to confirm coverage for.

Security & Privacy

Q – Does HYVE Raptor send any data to Vibe Software Solutions?

No. Zero. There is no telemetry, no crash reporting, no usage analytics, no license phone-home, and no automatic update check that transmits data to HYVE. The only outbound internet traffic during installation is the download of Ollama and the Alpha model from ollama.com. After installation, the platform generates no outbound traffic to any external server. You can verify this by monitoring your network traffic — you will find no connections to HYVE infrastructure.

Q – Can my operator see my clients' data through the platform?

No. This is an architectural property, not a policy statement. The protocol between raptor-shield and raptor-command transmits threat metadata — behavioral signals, anomaly scores, attack vector classifications, event timestamps — not the underlying data that the AI agents were processing. A healthcare client's PHI, a law firm's case files, a government agency's classified documents — none of this can reach the operator's server because the shield is designed to never transmit it. This is not enforced by a terms of service. It is enforced by how the system is built.

Q – How is the shield-to-server channel protected?

The channel between each raptor-shield and your raptor-command server uses ML-KEM-768, a NIST-standardized post-quantum key encapsulation mechanism (FIPS 203). This is applied at the application layer — not just at the transport layer. This means that even if standard TLS encryption is broken by future advances in quantum computing, the threat intelligence in transit remains protected. This is the "harvest now, decrypt later" threat model — adversaries recording your traffic today cannot decrypt it in the future with quantum hardware.

Q – Where is the shield's identity key stored, and who can access it?

Each shield generates a unique cryptographic identity key the first time it runs. That key is stored at `%ProgramData%\HyveRaptor\vault\` on the client machine, encrypted using Windows DPAPI with machine-scope protection. Machine-scope DPAPI means the key can only be decrypted by processes running on that specific machine — it cannot be copied to another machine and decrypted, and it cannot be accessed by user-level processes without the appropriate service context. Only the RaptorShield service can read its own identity key.

Q – Does my client need to know the server address is mine? Can I deploy this under my own brand?

HYVE Raptor is sold as a white-labelable operator platform. You are the authority — you generate keys, you run the server, you own the relationship with your clients. The shield installer can be configured with your organization's branding. Contact Vibe Software Solutions for white-label configuration options for your deployment.

Alpha Advisor

Q – I asked Alpha a question and it gave me wrong or outdated information. Why?

Alpha reads live system state for operational questions (threats, shields, compliance). For questions about external topics — regulatory frameworks, security concepts, specific CVEs — Alpha uses the knowledge baked into its language model, which has a training cutoff date. For current events or recently published standards, cross-reference Alpha's answers with authoritative sources. Alpha is most reliable and most powerful when you ask it about your specific deployment, not about the external world.

Q – Can I train Alpha on my organization's specific documents or procedures?

In the current release, Alpha uses the llama3.2 base model without custom fine-tuning. Customization through document ingestion (retrieval-augmented generation) and custom model fine-tuning are planned for a future release. Contact support to register your interest in this capability.

Q – Can I use a different AI model with Alpha?

Alpha is designed to work with models served by Ollama. If you have specific model requirements — a larger model for higher quality, a smaller model for faster inference on constrained hardware, or a domain-specific model — contact support for guidance on supported model configurations.

Swapping models is technically possible but the integration has been tested and validated against llama3.2 specifically.

Q — Alpha gave me a defense recommendation that seems aggressive. Should I always follow it?

Alpha's recommendations are advisory, not automatic. The platform will never issue a defense command without your deliberate action. Alpha surfaces its reasoning — you can see what evidence it drew on and why it made a given recommendation. Use Alpha's analysis as one input alongside your own judgment, your knowledge of the client's environment, and the operational context. For NEUTRALIZE and LOCKDOWN recommendations especially, treat Alpha as a well-informed analyst whose advice you weigh, not an authority whose instructions you execute without thought.

Compliance

Q — Can I use the compliance reports from HYVE Raptor in an actual audit?

The compliance engine provides real-time posture scoring based on your live deployment state. The evidence and control mapping are based on published framework requirements. Whether this data is accepted by a specific auditor as audit evidence depends on the auditor and the audit framework. The reports are accurate representations of your platform state at the time of export. We recommend presenting them as part of a broader evidence package alongside your policies, procedures, and other technical controls. Alpha Advisor can help you prepare audit narratives based on your current compliance state.

Q — My compliance score dropped overnight without me changing anything. Why?

Compliance scores reflect live system state. If a shield went offline, a required control that depended on connected shield coverage may have lapsed. If no defense commands have been issued recently, incident response controls may have aged out. Check the Compliance tab for the specific failing controls and review the evidence column — it will show you exactly what condition is not being met. Alpha Advisor can give you a fast path to restoring the score.

Billing & Licensing

Q — What happens if I exceed my shield limit?

The platform will accept connections up to your plan's shield limit. Shields that attempt to connect beyond that limit will fail authentication. They will not damage the platform or affect existing

connected shields. To add capacity, upgrade your plan and contact Vibe Software Solutions. The upgrade takes effect immediately — no reinstallation required.

Q — If I cancel my subscription, what happens to my deployment?

HYVE Raptor is software you run on your infrastructure. Cancellation ends your entitlement to updates, new feature access, and support. The currently deployed binaries continue to run. Nothing in the platform phones home to validate your subscription status at runtime — there is no kill switch. Your existing deployment remains functional. However, continuing to run a version that no longer receives security updates is not recommended for production environments.

Q — Is there a trial or evaluation period?

Yes. Contact Vibe Software Solutions at **majixx@vibesoftwaresolutions.com** or book a demo at **vibesoftwaresolutions.com/hyve-raptor** to request an evaluation package. Trial deployments are fully functional with a time-limited license key and a reduced shield limit. All features — including Alpha Advisor and the full compliance engine — are available during the trial.

SUPPORT & CONTACT

If your question is not answered in this manual, contact Vibe Software Solutions directly:

Email	majixx@vibesoftwaresolutions.com
Website	vibesoftwaresolutions.com/hyve-raptor
Book a demo	vibesoftwaresolutions.com/consult

When contacting support, please include: your plan tier, the version of the Raptor package you are running (visible in the Settings tab), and the contents of the relevant log files from `%APPDATA%\HyveRaptor\logs\`. This allows us to assist you without unnecessary back-and-forth.
