THE PITSTOP

Beta Tester Brief

AI Agent Security Scanner

Free Beta Access

EXECUTIVE SUMMARY

AI agents are being deployed at scale with little to no security auditing. The Pitstop is the first tool purpose-built to scan autonomous AI agents across 20 security checks -- covering identity, memory, backups, infrastructure, and supply chain. Run a scan in seconds, get a letter grade and specific fixes. No installs, no modifications, fully transparent.

The Pitstop scans your AI agent's security posture in under 10 seconds. You get a scored performance card (A+ to F) with a professional PDF report and specific, actionable tune-up recommendations. It's free during beta -- we just want your feedback.

24 Security Checks | 6 Categories | ~10 Seconds

One command. Full security audit. No installs. No modifications.

curl -sL https://thepitstop.ai/scan.sh | bash

thepitstop.ai

Version 1.2.0 | March 2026

Created by Nicholas Lynchini & Beeglie

What Gets Scanned

The Pitstop evaluates 20 security checks across 5 categories. Each check is scored individually, contributing to your overall grade.

1. Identity & Auth	API key exposure, file permissions, .env protection, SSH key hygiene
2. Memory & Encryption	Memory directory health, encryption status, permissions, size hygiene
3. Backup & Recovery	Backup scripts, recent backups, cloud backup tools, recovery documentation
4. Infrastructure	Firewall, non-root execution, container isolation, auto-updates
5. Supply Chain & Skills	Skill count, .gitignore protection, audit logging, workspace permissions
6. Dependency Chain	NPM vulnerability audit, postinstall script detection, dependency count, lock file integrity

Grade Scale

Grade	Score	Status	Description
A+	95-100	**Pole Position**	Hardened and race-ready
A	90-94	**Race-Ready**	Minor optimizations available
B	75-89	**Running Strong**	Some tune-ups recommended
C	60-74	**Needs Service**	Action needed
D	45-59	**Off Pace**	Significant tune-ups required
F	<45	**Off Track**	Critical issues, immediate attention

How to Run Your Scan

Prerequisites

Linux/macOS with bash, curl or wget. Your AI agent's workspace accessible on the machine.

Step 1: Open a terminal on the machine where your agent runs.

Step 2: Run one of these commands:

```
curl -sL https://thepitstop.ai/scan.sh | bash
```

or

```
wget -qO- https://thepitstop.ai/scan.sh | bash
```

Step 3: Review your results in the terminal. A PDF report link is generated automatically -- click it to download your professional Performance Card.

What happens:

The script checks 24 security categories across 6 areas, scores each one, calculates your overall grade, and saves a detailed JSON report locally. An anonymous summary (pass/fail per category only -- no agent name, paths, or system details) is submitted to generate a professional PDF Performance Card. A download link appears in your terminal automatically.

Duration:	~5-10 seconds
Installs anything?	No
Modifies your system?	No
Sends personal data?	No -- see Privacy section on next page

Terms of Use -- Beta Program

1. The Pitstop Beta Scanner ("the Tool") is provided as-is for evaluation purposes.
2. The Tool performs read-only analysis of your agent's security configuration. It does not modify any files, install any software, or make any system changes.
3. Beta testers agree to provide feedback on their experience. Feedback may be used to improve the product.
4. The Tool is free during the beta period. No payment information is collected.
5. You may discontinue use at any time by simply not running the scanner again.

Privacy Policy

We DO NOT collect:

- x Your agent's name or identity
- x File paths, directory names, or system hostnames
- x Contents of any files (memory, config, credentials, etc.)
- x API keys, tokens, passwords, or any secrets
- x IP addresses tied to scan results
- x Any personally identifiable information

We DO collect (anonymous only):

- > Pass/fail/warn status for each of 24 security checks
- > Point scores per check
- > Overall grade and total score
- > A randomly generated scan ID (not linked to your identity)

Purpose: Anonymous aggregate data helps us understand common security gaps across the AI agent ecosystem. This data is used solely to improve The Pitstop and publish aggregate security research (e.g., "72% of agents lack memory encryption").

Data retention: Anonymous scan summaries are retained indefinitely for research. You can verify exactly what is sent by reading the scan.sh source code -- it's fully transparent.

Source code: The complete scanner source is available at thepitstop.ai/scan.sh -- read it before you run it. We believe in full transparency.

Contact
nick@lynchini.com | thepitstop.ai