carahsoft

Synack

{I-NJI- tŜyúŜăűY ! Lπ t2ǿŜNJŜŘ tŜyŜúNJI-űƛ2y ¢Ŝăűƛy∃

/f2ăƛy∃ űĸKŜ ăŜŎdzNJƛűè Ŏ2ǾŜNJI·∃Ŝ ∃I·LJ ǿƛűK I·∃Ŝyűƛǒ ! L I·yŘ KdzY I·yπØI·fƛŘI·ûŜŘ ŜÈLJf2ƛűI·űƛ2y I·ú ăǑI·fŜ

Thank you for your interest in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider**® supporting a broad portfolio of industry-leading technologies through GSA, NASA SEWP V, E&I Carahsoft Cloud Solutions & Services Distributor Contract and a wide range of other contract vehicles.

As the **Master Government Aggregator**®, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with Synack, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.

To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com

Explore More Resources:
carah.io/synackresources

Join Events & Webinars:
carah.io/synackevents

Discover Technology Solutions:
carah.io/synack

Learn About Procurement:
carah.io/synackcontracts

Connect With Our Team:
Synack@carahsoft.com
(703) 871-8585

Synack

DATA SHEET

Sara Pentest: AI-Powered Penetration Testing

Closing the security coverage gap with agentic AI and human-validated exploitation at scale

According to new research from Omdia and Synack, while 95% of organizations prioritize penetration testing, they are currently testing only 32% of their total attack surface. This leaves 68% of your enterprise environment exposed to blind spots—a risk that is compounding as AI-enabled adversaries move faster than ever. The threat landscape is shifting as new AI models such as Mythos uncover hidden, high-risk attack paths at scale shrinking blind spots for adversaries.

The Sara (Synack Autonomous Red Agent) Pentest is powered by agentic AI and built on Synack's Penetration Testing as a Service (PTaaS) platform. It safely runs real-world attacks to find and prioritize exploitable vulnerabilities, which are then validated by human experts. This happens at a scale that traditional pentesting can't achieve, integrating seamlessly into your broader security program.

As your security team evolves to meet the coverage gaps and AI-driven threats, we have created two AI-led penetration tests: Sara Pentest and Sara Pentest+. Both are streamlined, fixed-scope penetration tests that provide testing coverage of both hosts and web applications. Sara Pentest is designed for smaller and low-complexity scopes. Sara Pentest+ is ideal for larger or more complex assets.

Key Benefits of Sara Pentest

REDUCE THE WINDOW OF EXPOSURE

Launch on-demand tests to respond quickly to product updates, code deployments, or emergent threats. In the era of shrinking Time-to-Exploit (TTE), this on-demand capability is critical for validating vulnerabilities before they are exploited.[1]

ACCELERATED TIME-TO-VALUE

Receive AI-powered - human-validated results in 2-3 days. Sara Pentest executes the pentest in hours, with all exploitable vulnerabilities then validated by Synack security experts.

SCALABLE & COST-EFFECTIVE COVERAGE

Expand test coverage to assets typically out-of-scope for pentesting. With AI-powered attacks now a top risk and already impacting 75% of organizations, this scalability is essential for matching AI scale and speed.[2]

CONSOLIDATE ALL PENTEST DATA IN ONE PLATFORM

Consolidate Sara Pentest results with other Synack human-led testing data. Get a unified view of your attack surface to improve overall security posture and eliminate recurring risk.

1 https://beelzebub.ai/blog/how-advanced-malware-self-update-systems-enable-exploitation-before-patches -can-be-applied/

2 https://lsvp.com/cyber60-2025-2026/

CUSTOMER CASE STUDY HIGHLIGHT

Sara Agentic AI Pentest

ENGAGEMENT SCOPE

Secondary Web Assets

VULNERABILITY FINDINGS

During the 48-hour autonomous pentest engagement, the Sara agent successfully identified and validated the following vulnerabilities in a previously untested secondary application:

- Information Disclosure
- Server/Application Misconfiguration
- Authentication/Session Management
- Cross-Site Scripting (XSS)
- SQL Injection (SQLi)

Key Performance Metrics

ASSESSMENT DURATION

48 hours

AVERAGE CVSS SCORE

7 (High)

MAX CVSS SCORE

9.1 (Critical)

FALSE POSITIVE RATE

0%

COST EFFICIENCY

At roughly 25% of the cost of a human-led test, Sara Pentest delivered significant value at speed.

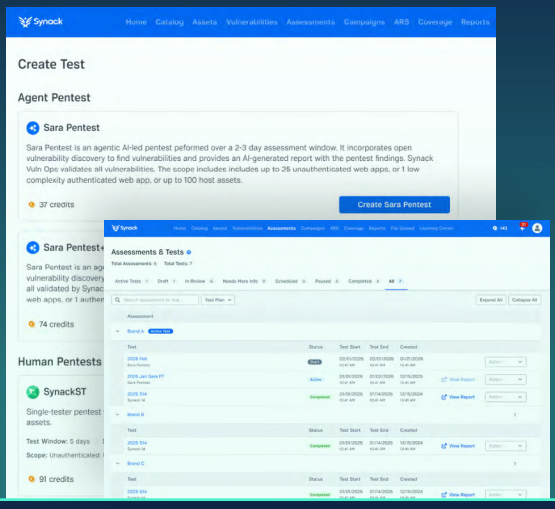

COMMON USE-CASES FOR SARA PENTEST

- Expanding pentest coverage to previously unmanaged or untested assets
- Quickly test or retest web and host assets for exploitable vulnerabilities
- Pentesting for internal compliance frameworks (i.e. CMMC Levels 1 or 2, OWASP vulns)

Sara Pentest Process Overview

Sara employs a multi-agent methodology for efficient and accurate testing. The following steps detail how Sara executes this comprehensive pentest process:

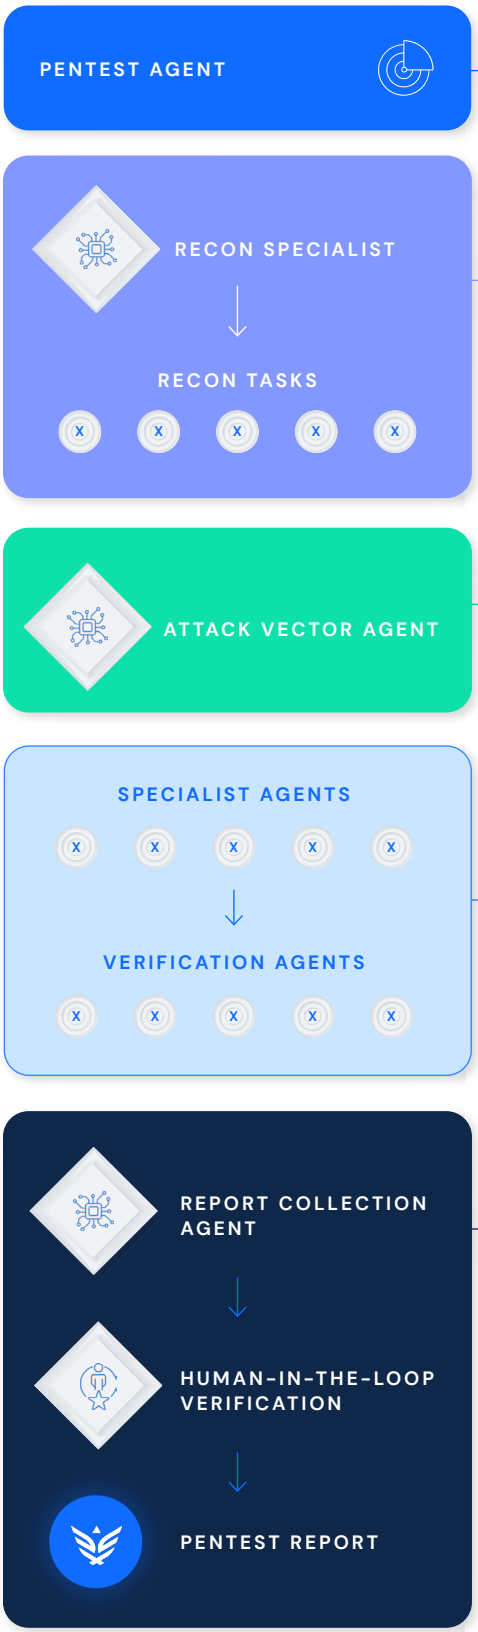

PHASE 1

DISCOVERY

Sara begins by performing a scan against the defined scope to accurately map the live attack surface, ensuring testing targets actual assets.

PHASE 2

RECONNAISSANCE

Specialized Recon agents analyze the scan results, decomposing them into targeted tasks (e.g., web application analysis, database enumeration, network service fingerprinting). These tasks leverage built-in domain knowledge and execute concurrently across different asset types for maximum speed and depth.

PHASE 3

ATTACK PLANNING & PRIORITIZATION

An Attack Vector Specialist agent synthesizes all reconnaissance findings. It correlates potential weaknesses and prioritizes the most probable attack paths based on known adversary techniques (TTPs) and relevant CVE intelligence, strategically directing the exploitation phase.

PHASE 4

CONTROLLED EXPLOITATION & VERIFICATION

Targeted Protocol-Specific agents execute safe, controlled exploits against prioritized vulnerabilities (e.g., attempting SQLi, testing for XSS, assessing SSH/FTP configurations). Every potential finding is automatically submitted to a separate Verification agent which independently re-tests and confirms exploitability, minimizing false positives.

PHASE 5

CONSOLIDATED REPORTING

A Report Collection Specialist agent aggregates all findings, and correlates related vulnerabilities to identify potential attack chains. Findings are then sent to Synack security experts for final human-in-the-loop review. Finalized findings are then presented in a single, consolidated report within the Synack platform, detailing proof-of-work, clear reproduction steps, and prioritized guidance for remediation.

Methodology

The testing methodology for this penetration test aims to provide compliance evidence and practical remediation guidance for both web applications and underlying host/infrastructure as scope allows. Testing is designed to surface exploitable weaknesses, validate security controls, and cover OWASP Top 10 categories of vulnerabilities.

Reporting

Sara Pentest delivers detailed, high-signal findings directly to the Synack PTaaS platform. All exploitable vulnerabilities surfaced by the agent are triaged by a human expert to minimize false positives.

EACH SARA PENTEST REPORT INCLUDES:

- Executive Summary
- Engagement Details
- Vulnerability List, Categorization/Distribution
- Detailed Findings for each verified exploitable vulnerability including:
  - Description, severity, & impact analysis
  - Specific recommendations of steps to remediate / fix
  - Steps to reproduce / proof of work
- Metrics for suspected vulnerabilities
- Explanation of Synack testing methodology

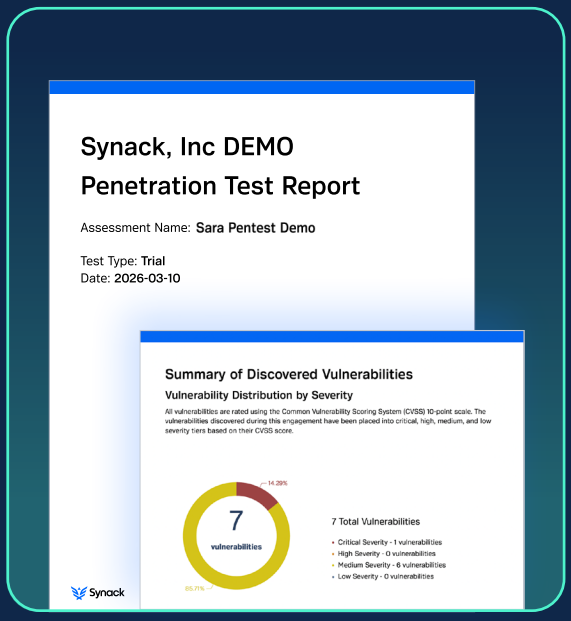

The platform also includes on-demand patch verification to confirm that fixes are successful.

All discovered vulnerabilities can be exported via Synack's integrations (e.g., Jira, Splunk, ServiceNow, Nucleus), or leveraging Synack's API, to streamline your management workflow.

Sara Pentest in Context

The Synack PTaaS platform provides a range of pentesting solutions, from agent-driven to human-led. The table below outlines how Sara Pentest is designed to work alongside Synack's human-led offerings, giving organizations the flexibility to align the right testing methodology with the specific business risk.

Synack Testing Features	Sara PT/PT+	Synack ST/ST+	Synack14	Synack90	Synack365
Assessment Window	2-3 Days	5-10 Days	14 Days	90 Days	365 Days
Testers	Agent Driven	Assigned Researcher	Pool of Researchers	Rotating Pools of Researchers	Rotating Pools of Researchers
Test Methodology	Open Vulnerability Discovery	Guided Vulnerability Discovery	Open Vulnerability Discovery	Open Vulnerability Discovery	Open Vulnerability Discovery
Test Type	Point-in-Time	Point-in-Time	Point-in-Time	Continuous	Continuous
Primary Mission	Scaled Coverage & Validation, Risk Assessment	Compliance, Risk Assess-ment	Risk Reduction, Compliance	Risk Reduction, Compliance	Risk Reduction, Compliance
Test Scope	Web or Host	Web or Host	Web, Host, Cloud, Mobile, API, AI/LLM	Web, Host, Cloud, Mobile, API, AI/LLM	Web, Host, Cloud, Mobile, API, AI/LLM
Vulnerability Management	Yes	Yes	Yes	Yes	Yes

Open Vulnerability Discovery

SARA PT AND SYNACK14, 90, AND 365 DELIVER HIGH-IMPACT COVERAGE BY ACTIVATING AGENTIC AI-TESTING OR A CURATED POOL OF RESEARCHERS ON YOUR AGREED-UPON SCOPE.

Each researcher brings their own tools, tactics, and tradecraft and is incentivized to uncover exploitable vulnerabilities. This model provides fast, comprehensive coverage with high-confidence results that drive meaningful risk reduction. Sara PT/PT+ agentic-AI testing utilizes diverse specialized testing agents, vulnerability chaining, and orchestration delivering broad tools, tactics, and techniques to find exploitable vulnerabilities on in-scope assets.

Checklist-Based Assessments

CHECKLIST-BASED ASSESSMENTS PROVIDE A CLEAR, STRUCTURED PATH TO MEETING COMPLIANCE AND CONTROL ASSURANCE REQUIREMENTS.

A dedicated researcher is assigned to methodically validate security controls across your in-scope assets for detailed evidence of coverage and configuration. Rather than broad vulnerability hunting, this approach ensures consistent evaluation of your foundational security controls—helping you demonstrate compliance, identify gaps, and confidently prepare for audits.

Sara PT and Sara PT+ for Web Applications

Web App Scope Details

Use the criteria below to determine whether Sara PT or Sara PT+ is the best fit for a test of your web application. At this time Sara, only tests external web assets.

Sara PT	
Unauth Web Apps	**Auth Web App**
Up to 25 Web Apps	Small / basic web app
Recommended for testing unmanaged or untested assets including a single-purpose web applications, eg.: • Straightforward functionality • Minimal input surfaces • Limited business logic	

Sara PT+	
Unauth Web Apps	**Auth Web App**
Up to 50 Web Apps	Large / full-featured web app
Recommended for testing unmanaged or untested assets with increased coverage for full-featured web applications, eg.: • Form-dense • Support multiple user roles • Incorporate expanded business logic	

Host Scope Details

Use the criteria below to determine whether Sara PT or Sara PT+ is the best fit for a test of your host. At this time, Sara only tests external hosts.

Sara PT
Hosts
Up to 100 Host Assets

Sara PT+
Hosts
Up to 250 Host Assets

Web App Assessment Details

WHAT'S INCLUDED IN A WEB APP ASSESSMENT?

Our web application penetration testing methodology for Sara and Sara Pentest+ covers the following key categories. The specific techniques applicable to and executed in each assessment depend on the outcome of the agent's recon activities and always adhere to the guardrails. The agent is non deterministic and, similar to human-led engagements, does not check for every vulnerability listed here in every assessment.

CROSS-SITE SCRIPTING (XSS)

- ☐ Reflected XSS
- ☐ Stored/Persistent XSS
- ☐ DOM-Based XSS
- ☐ Blind XSS
- ☐ Context-specific exploitation (HTML, JavaScript, attribute injection)
- ☐ Polyglot payloads

SQL INJECTION (SQLi)

- ☐ Error-based SQL injection
- ☐ UNION-based SQL injection
- ☐ Boolean-based blind SQL injection
- ☐ Time-based blind SQL injection
- ☐ Out-of-band SQL injection
- ☐ Second-order SQL injection
- ☐ NoSQL injection (MongoDB, CouchDB, Redis)
- ☐ ORM injection (Hibernate, Sequelize, SQLAlchemy)

INSECURE DIRECT OBJECT REFERENCE (IDOR)

- ☐ Sequential ID enumeration
- ☐ GUID/UUID manipulation
- ☐ Parameter tampering (user IDs, account numbers, document IDs)
- ☐ Horizontal privilege escalation
- ☐ Vertical privilege escalation
- ☐ Mass assignment vulnerabilities
- ☐ File path manipulation
- ☐ API endpoint enumeration

SERVER-SIDE REQUEST FORGERY (SSRF)

- ☐ Internal network access
- ☐ Cloud metadata exploitation (AWS, Azure, GCP)
- ☐ Protocol smuggling (file://, gopher://, dict://)
- ☐ DNS rebinding attacks
- ☐ URL parser confusion
- ☐ Blind SSRF detection
- ☐ Port scanning via SSRF
- ☐ SSRF to RCE chains

CROSS-SITE REQUEST FORGERY (CSRF)

- ☐ CSRF token validation bypass
- ☐ SameSite cookie attribute testing
- ☐ GET based state changes
- ☐ Origin and Referer header bypass
- ☐ JSON/XML CSRF
- ☐ Multi-step transaction CSRF
- ☐ Subdomain CSRF
- ☐ WebSocket CSRF

XML EXTERNAL ENTITY (XXE)

- ☐ Classic XXE (in-band file disclosure)
- ☐ Blind XXE (out-of-band data exfiltration)
- ☐ XXE based SSRF
- ☐ Parameter entity injection
- ☐ XInclude attacks
- ☐ Parser-specific vulnerabilities
- ☐ Document processing XXE (DOCX, SVG, XLSX)

SERVER-SIDE TEMPLATE INJECTION (SSTI)

- ☐ Template engine identification (Jinja2, Twig, Freemarker, Velocity, ERB)
- ☐ Remote Code Execution via templates
- ☐ Sandbox escape techniques
- ☐ File system access via template injection
- ☐ Expression language injection (JSP EL, Spring EL, OGNL)
- ☐ Blind SSTI detection

LOCAL FILE INCLUSION (LFI)

- ☐ Path traversal attacks
- ☐ Null byte injection
- ☐ Wrapper exploitation (php://, data://, zip://)
- ☐ Log poisoning
- ☐ LFI to RCE escalation
- ☐ Session file inclusion
- ☐ Filter bypass techniques (encoding, double encoding)

REMOTE FILE INCLUSION (RFI)

- ☐ Remote code inclusion
- ☐ URL wrapper exploitation (http://, ftp://)
- ☐ RFI to RCE
- ☐ SMB/UNC path inclusion
- ☐ Configuration-based RFI testing
- ☐ Filter evasion techniques

COMMAND INJECTION

- ☐ Direct command injection
- ☐ Blind command injection
- ☐ Time-based detection
- ☐ Command chaining
- ☐ Expression evaluation
- ☐ Shell metacharacter injection
- ☐ Context-specific injection (Python, PHP, Node.js)

BROKEN AUTHENTICATION

- ☐ Session fixation and hijacking
- ☐ Weak password policies
- ☐ MFA bypass
- ☐ JWT vulnerabilities
- ☐ OAuth/SAML flaws
- ☐ Authorization bypass and privilege escalation
- ☐ Password recovery vulnerabilities
- ☐ Timing attacks and authentication oracles

JWT (JSON WEB TOKEN) VULNERABILITIES

- ☐ Algorithm confusion (none algorithm, RS256 to HS256)
- ☐ Weak signing secrets
- ☐ Missing signature validation
- ☐ Token payload manipulation
- ☐ Key confusion attacks
- ☐ Token expiration issues
- ☐ JKU/JWK header injection
- ☐ Token storage vulnerabilities

BROKEN ACCESS CONTROL

- ☐ Horizontal privilege escalation
- ☐ Vertical privilege escalation
- ☐ Missing function-level access control
- ☐ IDOR vulnerabilities
- ☐ Multi-tenant isolation bypass
- ☐ CORS misconfiguration
- ☐ Path traversal in authorization
- ☐ RBAC bypass

Host Assessment Details

WHAT TYPES OF VULNERABILITIES DOES SARA PENTEST FIND?

Our host penetration testing methodology for Sara and Sara Pentest+ covers the following key categories. The specific techniques applicable to and executed in each assessment depend on the outcome of the agent's recon activities and always adhere to the guardrails. The agent is non deterministic and, similar to human-led engagements, does not check for every vulnerability listed here in every assessment.

SSH (SECURE SHELL) PORT 22

- ☐ Authentication bypass
- ☐ User enumeration
- ☐ Weak ciphers and key exchange
- ☐ Credential testing
- ☐ Configuration weaknesses (PermitRootLogin, PasswordAuthentication)
- ☐ Version-specific vulnerabilities

FTP (FILE TRANSFER PROTOCOL) PORT 21

- ☐ Anonymous FTP access
- ☐ Default credentials
- ☐ Directory traversal
- ☐ TLS/SSL misconfiguration
- ☐ Version-specific vulnerabilities (vsftpd, ProFTPD, Pure-FTPd)
- ☐ File permission issues

SMTP (SIMPLE MAIL TRANSFER PROTOCOL) PORT 25

- ☐ Open relay testing
- ☐ User enumeration (VRFY, EXPN, RCPT TO)
- ☐ Authentication bypass
- ☐ TLS/STARTTLS downgrade
- ☐ Information disclosure
- ☐ Version-specific vulnerabilities (Postfix, Sendmail, Exim)

CVE BASED TESTING

- ☐ CVE identification and mapping
- ☐ Exploit availability assessment
- ☐ Vulnerability validation
- ☐ Zero-day assessment
- ☐ Vendor advisory tracking

SMB (SERVER MESSAGE BLOCK) PORTS 139/445

- ☐ EternalBlue (MS17-010)
- ☐ Null session enumeration
- ☐ Share access testing
- ☐ Version downgrade (SMBv1)
- ☐ Default credentials
- ☐ PsExec exploitation

RDP (REMOTE DESKTOP PROTOCOL) PORT 3389

- ☐ BlueKeep (CVE-2019-0708)
- ☐ Default credentials
- ☐ Network Level Authentication (NLA) bypass
- ☐ User enumeration

NFS (NETWORK FILE SYSTEM) PORT 2049

- ☐ Export enumeration
- ☐ Unauthorized mount access
- ☐ Version specific vulnerabilities (NFSv2, NFSv3, NFSv4)
- ☐ Kerberos authentication bypass
- ☐ File permission issues
- ☐ Information disclosure

REMOTE CODE EXECUTION (RCE)

- ☐ Command injection
- ☐ Deserialization attacks (Java, Python, PHP, .NET)
- ☐ Logic flaws leading to code execution
- ☐ File upload exploitation
- ☐ Template injection to RCE
- ☐ Expression language injection
- ☐ API abuse

Adopting AI Pentesting with Safety, Control, and Confidence

Adversaries are using AI to increase speed and scale of their attacks, and 75% of organizations have already experienced AI-related security incidents. This increased risk requires full visibility into exploitable vulnerabilities and faster response from security teams.

To meet this challenge, organizations are beginning to adopt autonomous and agentic AI AI-driven pentesting tools. However, adopting these tools requires careful evaluation. Trusting an agent with privileged access based on its training alone is insufficient. The evaluation framework in this guide, along with the two downloadable attachments, provides the tools to assess the technical guardrails of a vendor's agentic AI safety guardrails.

BY REQUIRING VERIFIABLE, HARD-CODED GUARDRAILS, SECURITY LEADERS CAN CONFIDENTLY ADOPT AGENTIC AI PENTESTING AS A TRUSTED AND EFFECTIVE SOLUTION AS PART OF THEIR SECURITY PROGRAM.

See How Synack's Agentic AI Work in Practice

Ready to adopt AI-driven pentesting?

Schedule a demo of Sara Pentest today to see how our agentic AI is bound by a decade of proven, hard-coded safety controls.

SCHEDULE A DEMO

Use AI as a Force Multiplier for Penetration Testing

Synack is the leader in human-led and AI-powered Penetration Testing as a Service (PTaaS), transforming offensive security to help organizations proactively reduce risk, stay compliant, and defend against evolving cyber threats. We are committed to making the world more secure by harnessing agentic AI innovations and a talented, vetted community of security researchers to deliver continuous penetration testing and autonomous vulnerability management. Founded by former NSA operatives, Synack has enabled nearly 10 million hours of expert testing to protect critical assets, from global financial systems to U.S. Defense Department networks.

To learn more about Synack, visit www.synack.com.