Service Description:

Human on the Loop Penetration Testing

Overview

eSentire's Human on the Loop Penetration Testing (**"Service")** is a continuous, automated penetration testing solution powered by the Atlas Platform. The Service enables organizations to proactively identify, validate, and remediate security exposures across their external, cloud, and web application environments on a scheduled cadence aligned to their risk posture. Specific elements of the Service are described in this Service Description. The Service is scoped using the total number of users. "**Users**" are defined as registered human identities registered in the Client's Identity Platform (IdP). For the avoidance of doubt, this does not include machine, automation or agent identities. The User count will map to test scope entitlements detailed in this Service Description (see the section "Entitlements and Usage Policies" below). The number of Users included in the service are detailed on the Order Form and subject to overage fees, if applicable.

Service Definitions

In addition to the below, capitalized terms have the meaning given to such terms in this Service Description. Capitalized terms used but not defined in this Service Description have the meaning given to such terms in the Master Security Services Agreement between eSentire and Client that governs this Service Description.

"Atlas AI" means the Atlas agentic AI framework, designed to enrich and investigate Findings generated by the Service.

"Atlas Platform" means the eSentire security operations platform which consolidates all data and drives workflows for all eSentire services, including those in this Service. All data on the Atlas Platform related to delivery of the Service for Client is retained for the Term and deleted upon Service expiration.

"Atlas UI" means the user interface of the Atlas Platform where interaction with and configuration of the Atlas Platform and penetration testing services is performed.

"Client" means the entity ordering the Service described in this Service Description.

"Client Environment" means the Client systems, infrastructure, cloud services, and applications in scope for penetration testing, as agreed upon during onboarding and specified in the Order Form.

"Engagement" means an autonomous test, review, or assessment performed by Atlas AI. Available Engagement types are described in the table in the "Service Elements" section below.

"Engagement Entitlement" means the agreed-upon number of scheduled Engagements per year — as specified in the Order Form.

"Finding" means a confirmed security vulnerability or exposure identified during a penetration test, delivered via the Atlas UI and automatically routed to the Client's ITSM platform.

"Identity Platform (IdP)" means a system that creates, maintains, and manages digital identities for users, computers, or devices.

"Order Form" means an ordering document that specifies the eSentire services ordered by Client, including the selected Engagement Entitlement, user tier and Scope Entitlement.

"Scope Entitlement" means the number of external IPs, web applications, and cloud accounts included in the Service for a given user tier, as specified in the applicable Order Form.

"Validate Fix" means a targeted, on-demand re-test of a specific Finding, performed to validate that a specific remediation or mitigation has been successfully applied.

Service Elements

The Service consists of the following elements:

- Scheduled Human-assisted Engagements at the agreed Engagement Entitlement ;
- Self-Scheduled Automated Tests at the agreed Engagement Entitlement ;
- Validate Fixes up to a 100 per year entitlement;
- Automated Finding delivery to the Atlas UI;

The Service consists of the following engagement types:

Engagement Type	Description
Internal Vulnerability Assessment	Human-assisted assessment targeting Client internal network infrastructure across endpoints, servers, and network devices. A qualified security analyst validates, contextualises, and prioritises identified exposures — distinguishing true vulnerabilities from false positives and providing risk-ranked remediation guidance aligned to the Client environment.
Internal Penetration Test	Human-assisted simulated insider attack targeting Client internal network and Active Directory/Hybrid environments. A certified penetration tester actively attempts credential attacks, lateral movement, and privilege escalation — validating real-world exploitability and documenting the confirmed path to Domain Admin with evidence-backed findings.
PCI Vulnerability Assessment	Human-assisted vulnerability scan of Client systems within the cardholder data environment (CDE), conducted in alignment with PCI DSS 4.0 requirements. A qualified security analyst reviews and validates scan results, removes false positives, and produces audit-ready evidence and attestation documentation for quarterly ASV compliance reporting.
Mobile App Vulnerability Assessment	Human-assisted assessment targeting Client iOS and Android applications. A certified mobile security tester conducts static and dynamic analysis to identify vulnerabilities including insecure data storage, weak authentication, improper session handling, API misconfigurations, and OWASP Mobile Top 10 exposures — with manual validation of all findings prior to delivery.
External Vulnerability & Penetration Test	Automated test that targets internet-facing assets including IP addresses, hostnames, and exposed services. Simulates an unauthenticated external attacker across the full reconnaissance-to-exploitation chain.
Cloud Penetration Test	Automated test that targets Client cloud infrastructure across AWS, Azure, and GCP. Assesses IAM misconfigurations, overprivileged roles, publicly exposed storage, cross-service lateral movement, and cloud-native attack vectors.
Cloud Security Best Practices Review	Automated review of Client cloud configuration against CIS benchmarks and provider hardening guides. Identifies deviations from security best practices across identity, network, encryption, and logging controls.
Web Application Vulnerability	Automated assessment that targets Client web applications. Identifies vulnerabilities including OWASP Top 10, authentication weaknesses, injection flaws, API misconfigurations, and business logic errors.

Engagement Type	Description
Assessment (WAVA)	

Human on the Loop Delivery Model

The test is conducted primarily through semi-automated tooling, with a qualified human leading the engagement. The human reviews all automated output, removes false positives, validates high-severity findings, and contextualises results against the Client environment. Where automated tooling cannot confirm exploitability or complete an attack chain, the human performs targeted manual testing to fill the gap. A final review is conducted before delivery to ensure all findings are accurate, evidence-backed, and accompanied by actionable remediation guidance.

Findings and Reporting

Every Finding generated by the Service is delivered as a Finding in the Atlas UI. Each Finding includes:

- a description of the vulnerability and affected assets;
- severity classification (Critical, High, Medium, Low, Informational);
- contextual evidence including proof-of-concept details where applicable;
- assigned ownership based on asset attribution; and
- recommended remediation guidance.

Validate Fixes

Validate Fixes enable targeted, on-demand validation of specific Findings outside scheduled tests. Validate Fixes are appropriate when:

- a remediation or mitigation has been deployed, and the Client wishes to confirm the vulnerability has been resolved;
- a new vendor, technology, or code release introduces a specific exposure that warrants immediate testing; or

Validate Fixes are scoped to the specific Finding being validated and consume one Validate Fix entitlement from the allocation. Validate Fixes are user-initiated and tracked via the Atlas UI.

Deployment

Following receipt of a fully executed Order Form, Client will be provided with access to the Atlas UI which provides visibility over all subscribed eSentire services. The Atlas UI acts as a central hub for managing and monitoring the Service, allowing Client to configure, integrate, and view Service elements.

Entitlements and Usage Policies

The Service is offered in three tiers. All Human on the loop and autonomous Engagement types (outlined above) are available in each tier. The tier determines Engagement entitlement as shown below.

	Small	Medium	Large
Engagements/ year	2	8	24

Each Engagement is scoped by User count as set out in the applicable Order Form to entitle Client to a defined maximum number of hosts, external IP blocks (256 IPs), web applications, and cloud accounts as shown below.

User count	Hosts	Internal IP Blocks (24)	External IP Blocks (256 IPs)	Web applications	Cloud accounts
50–100	180	1	10	2	1
101–250	450	2	24	4	2
251–500	900	4	45	12	3
501–1,000	2000	6	90	18	5
1,001–1,500	3750	10	160	25	8
1,501–2,500	6250	16	250	40	12
2,500–5,000	15000	28	450	60	20

Fair Use policies will be in effect to enforce reasonable usage of the Service. The User count scoped for the Service will be mapped to expected volumes of in-scope hosts, IP blocks, cloud accounts, and web applications. "Reasonable use" means that the volume of in-scope assets will fall within 20% for each entitlement. In the event that scope requirements exceed expected bounds, eSentire will work with Client to investigate and recommend additional configurations or scope.

The number of Validate Fixes is limited to 100 per year and is not subject to the Fair Use policy above.

Add-Ons

The following add-ons may be purchased.

SKU	Unit
Additional Engagement	Per Engagement
Additional Hosts	Per 200 Host / yr

SKU	Unit
Additional Internal IP Block (24 IPs)	Per IP Block/ yr
Additional External IP Block (256 IPs)	Per IP Block / yr
Additional Web application	Per Web application / yr
Additional Cloud account	Per Cloud account / yr