



SECURITY AND COMPLIANCE

Traceforce Security and Compliance Overview

How the platform handles data, who can reach it, and the independent assurance behind it.

Version 1.2
October 2026

PORT1

Prepared by PORT1, Traceforce's channel partner

Shareable before an NDA. This overview covers what can be shared up front. The full documentation package is available once an NDA or partner agreement is in place.

What's here, and what's available next

This overview explains how Traceforce handles data, who can reach it, and the independent assurance behind the platform. It covers what can be shared before an NDA or partner agreement is in place. More is available once one is signed, including Traceforce's SOC 2 Type II report, its system architecture documentation, and its HIPAA assessment.

Compliance at a glance

ISO/IEC 27001:2022

Certified

Glocert International Certifications
(UK) Limited
Certificate 26ISMS-1227
Valid through August 6, 2029

SOC 2 TYPE II

Attested

Security trust services criteria
February 1 to April 30, 2026
Ezzy & Associates, LLC, unmodified
opinion

HIPAA

Assessed

Internal assessment, August 2026
OCR HIPAA Audit Protocol
BAA signed directly by Traceforce

The documentation package

DOCUMENT	WHAT IT COVERS	ACCESS
Security and Compliance Overview	This document	AVAILABLE NOW
Security Q&A	Answers to common security and compliance questions, with stable question IDs and a version history	AVAILABLE NOW
ISO 27001 certificate	Certification scope, issuing body, and validity	AVAILABLE NOW
ISO 27001 Statement of Applicability	All 93 Annex A controls and how each one applies	AVAILABLE NOW
SOC 2 Type II report	The auditor's opinion, system description, controls, and test results	UNDER NDA
System and network architecture	Detailed data flow and infrastructure, including subprocessors	UNDER NDA
HIPAA compliance assessment	Assessment against the HIPAA Security Rule and Breach Notification Rule	UNDER NDA
Business Associate Agreement	Signed directly with Traceforce by healthcare customers that need one	WHEN NEEDED

Under NDA means once an NDA or partner agreement is in place with PORT1. The last page explains how to request the full package.

What Traceforce does

Traceforce runs as a lightweight agent on each company computer, paired with a browser extension for Chrome and Edge. It sees the AI running on the device and in the browser, including desktop apps, coding assistants, agents, and the connectors they use.

Shadow AI AI tools in use that haven't been approved or reviewed.	Personal accounts Personal AI accounts used for work instead of company ones.	Agents and assistants AI agents, coding assistants, IDE extensions, and AI browsers.
Connected workflows MCP servers, connectors, and integrations that may create risk.	Sensitive-data indicators Credentials and personal data appearing in AI activity, on supported tools.	Users and devices The people and company computers behind all of it.

Why it works on the device

The context that decides whether AI activity is risky lives on the device: which connector was invoked, what permissions it holds, and whether a person or an agent started the action. Conventional security tools see network events, processes, and file writes. Traceforce observes where that context exists, and every finding ties back to a specific tool, on a specific computer, used by a specific person.

How risk is scored and controlled

Traceforce scores the AI tools and connectors it finds, drawing on its Atlas registry, which rates hundreds of publicly available AI agents and MCP servers. Policies set a requirement for each control, and when what's observed doesn't meet it, an issue is raised against a specific device and user.

Low 0-39	Medium 40-59	High 60-79	Critical 80 and above
--------------------	------------------------	----------------------	---------------------------------

Controls act on the specific operation. For coding agents, a risky delete through one connector can be blocked while the rest of the tool keeps working. On supported tools, sign-ins that don't use a company email domain can be blocked, and sensitive data in prompts can trigger a warning or a block.

Traceforce is additive
It runs alongside endpoint detection and response, data loss prevention, and cloud access controls, and covers the AI activity on each computer that those tools can't see.

Where your data lives

Two categories of data are involved, and they're handled differently. Most security reviews start with this question, so here's the precise answer.

Analysis happens on the device

The agent analyzes AI use locally and sends only metadata and findings to the platform. In the default configuration, no prompt or tool-call content is stored, and credentials found on a device are never uploaded.

The two categories, side by side

CATEGORY	WHAT IT CONTAINS	WHERE IT LIVES
Usage metadata	AI application and web usage data, derived findings, and device and user attribution	Encrypted storage inside the Traceforce platform
Audit logs and finding snippets	Prompt-level activity and short snippets of sensitive-data findings, only if you turn on audit logging or customer-owned storage	Storage you own. Endpoints write directly to your Amazon S3, Google Cloud Storage, or Azure Blob Storage, and none of it is stored on Traceforce infrastructure.

Retention and deletion

Usage metadata is kept for 12 months by default, and customers can request deletion. Anything in your own storage follows the retention you set there. Device records clear about 1 month after a device stops reporting.

Hosting and residency

Core platform infrastructure runs in AWS us-east-2 (Ohio, United States), and more hosting regions are being added. Data you choose to log stays in your own cloud account, in the region you choose.


Personal data the platform holds

Contact information	Email address
Technical data	Device name, device ID, user name, and usage data
Transactional data	Service usage

Collection is deliberately minimized. Where the browser extension records web activity, URL parameters are stripped and HTTP headers aren't uploaded. When an admin opens a finding, the console reads it from your storage using a signed link that expires after 15 minutes.

System and data flow

Your environment sits outside the platform's security boundary. Components are described here by function. The detailed architecture document is available under NDA.



Solid lines carry data. Dashed lines are optional. Every flow crossing the security boundary uses TLS 1.2 or better, and traffic between the agent and the platform uses TLS 1.3.

The flow, step by step

- 1** Your administrators sign in to the console through your identity provider using SAML 2.0.
- 2** Administrators download the agent package and deploy it through your RMM or MDM.
- 3** Administrators set policies and notification settings in the console. Integrations use the API.
- 4** The agent periodically uploads AI application and web usage metadata to encrypted storage in the platform.
- 5** Operational logs from the agent go to the platform's monitoring service.
- 6 7** New data queues an analysis job, and processing jobs analyze it. Other jobs run on events or a schedule.
- 8** Results are stored in the platform database and shown in the console.
- 9** Notifications are sent according to your organization's settings.
- 10** Optional. If configured, endpoints write audit logs and finding snippets directly to storage you own, and your admins view that content through the console.

Identity, isolation, and encryption

Console access runs through your own identity provider, and tenant separation is enforced in the database itself.

Single sign-on	SAML 2.0 with your identity provider. Documented setups for Microsoft Entra ID, Okta, and OneLogin, including single logout.
MFA and conditional access	Applied by your identity provider. Because console sign-in runs through SAML, your MFA and conditional access policies apply to every sign-in.
Sign-in scope	Limited to your own email domain, plus admins in the managing MSP's parent tenant. Guest and external accounts can't sign in.
Admin access	A client's own admins see only their own organization. Admins in the managing MSP's parent tenant manage that MSP's client organizations.
Directory sync	Users and groups sync from your identity provider over SCIM, including Microsoft Entra ID and Okta. Membership is managed in your identity provider, and synced groups can be assigned to policies.
Tenant isolation	Enforced in the database with row-level security and unique identifiers for every organization and user. Endpoint data is held in dedicated stores with additional access controls.
Encryption in transit	TLS 1.3 between the agent and the platform. The console uses TLS 1.3 and supports TLS 1.2 for compatibility.
Encryption at rest	AES-256 across platform databases and file storage, with keys managed and rotated through a managed key service.
Secrets	On macOS, the agent stores its credentials in the System Keychain. API client credentials can be viewed by admins with console access to that organization, so treat console access as access to those credentials.
API authentication	OAuth 2.0 client credentials flow issuing short-lived bearer tokens.
Hosting and perimeter	Core infrastructure in AWS us-east-2, behind a web application firewall. The database is reachable only from inside the security boundary.
Traceforce staff access	Least privilege, granted by documented request and approval, with unique named accounts, MFA, and monthly access reviews.

How Traceforce runs security

Traceforce's security program is certified under ISO 27001 and covered by its SOC 2 Type II report. These are the areas a vendor review usually asks about.

Governance

A designated executive owns information security and compliance. An IT leadership committee meets monthly, and executive management reviews risk and controls at least annually.

Risk management

A formal risk assessment at least annually, with a treatment plan for each risk and remediation tracked to resolution.

People

Background checks before onboarding, security awareness training at hire and annually, and signed confidentiality agreements.

Secure development

Formal change control, an additional approval before changes reach production, separate development and test environments, and no unauthorized changes directly in production.

Vulnerability management

Annual penetration testing, annual internal and external vulnerability assessments, and ongoing scanning, with findings tracked to remediation.

Monitoring

Centralized logging and managed threat detection with alerting, plus monthly access reviews.

Incident response

A documented process covering identification, communication, investigation, root cause analysis, remediation, and post-incident review.

Continuity and backups

Daily database backups in the production region, periodic restore tests, and business continuity and disaster recovery plans reviewed annually.

Vendor management

A risk assessment before any new vendor is onboarded, including a review of its assurance reports, and an annual review of subprocessors.

Subprocessors

Traceforce uses a small number of subprocessors, including for cloud hosting, the web console, the platform database, and directory sync, with AWS as the primary hosting provider. The hosting, console, and database providers are named, with each one's role, in the SOC 2 Type II report and architecture documentation, available under NDA.

Traceforce's team works fully remotely and runs no data centers of its own, so physical security is handled by its cloud providers. To report a vulnerability in Traceforce, see the Vulnerability Disclosure Policy at traceforce.ai/legal.

Who handles what

Security works best when each side knows its part. Where an MSP manages Traceforce for you, the MSP may handle some of your responsibilities on your behalf.

AREA	TRACEFORCE	YOU, OR YOUR MSP ON YOUR BEHALF
Identity and access	SAML sign-in, sign-in limited to your email domain, and tenant isolation	Enforce strong passwords and MFA in your identity provider, add and remove users promptly, and review access periodically
Deployment	A signed agent and browser extension, and encrypted communication with the platform	Deploy through your RMM or MDM and keep devices enrolled
Data you choose to log	Writes only to the storage you configure, with audit logs redacted by default	Set encryption, access, and retention on that storage
Policies	Detection, scoring, and enforcement on supported tools	Decide which controls to enforce, and for whom
Incidents	Detect, investigate, and communicate incidents affecting the platform	Report suspected incidents promptly and keep your own incident response plan
Network	Web application firewall, threat detection, and encrypted transport	Configure firewalls and remote access for your environment, and monitor your network

The identity, incident, and network rows reflect the customer controls listed in Traceforce's SOC 2 Type II report.

For healthcare organizations

Business Associate Agreement

Traceforce signs a BAA directly with healthcare customers that need one, scoped to how the platform could come into contact with ePHI.

When ePHI could be involved

By default the platform holds only usage metadata. ePHI could be involved only if you turn on audit logging or customer-owned storage and ePHI appears in monitored AI activity. That content is written to your own storage, which the platform reads when your admins review findings.

HIPAA assessment

An internal assessment completed in August 2026 against the HIPAA Security Rule and Breach Notification Rule, using the OCR HIPAA Audit Protocol. Available under NDA.

Deploying, evaluating, and getting the full package

Platforms

Windows 10 and 11, macOS 15 and 26, Chrome and Edge through the browser extension, and Linux dev containers on a host running the agent. Support for iOS and Android is on the roadmap for early 2027.

Deployment

The agent installs silently through the RMM or MDM you already run. On macOS, MDM profiles grant the permissions it needs. An initial inventory usually appears within about an hour.

Integrations

A REST API exposes the same data and controls as the console, so findings can feed a SIEM, SOAR, ticketing system, or internal tools. Policies can apply to specific user groups.

Evaluation

During an AI Risk Assessment, controls report to the admin only. Nothing is blocked and no workflow changes. Enforcement is a separate decision made afterward, control by control.

Getting the full documentation package

1

Sign an NDA or partner agreement with PORT1.

2

Receive the SOC 2 Type II report, architecture documentation, and HIPAA assessment.

3

If you're a healthcare organization, sign a BAA directly with Traceforce.

Questions, or ready for the full package?

Reach out to your Partner Success Manager at PORT1. If you don't have a direct contact, fill out the contact form at port1.io.

PORT1

Provided for information as of the version date. Capabilities vary by AI tool and platform. Contractual commitments are set out in the applicable agreements.