Elastic Privacy Datasheet

Observability

Product Summary: Elastic Observability

Elastic Observability provides granular insights and context into the behavior of applications running in your environments, enabling quick detection and resolution of root cause events. It offers a single stack to unify your logs, infrastructure metrics, application traces, user experience data, synthetics, and universal profiling. Data is ingested directly into Elasticsearch, where it can be processed and enhanced before visualization and alerting in Kibana. Observability is designed to accelerate problem resolution with open, flexible, and unified observability powered by advanced machine learning and analytics.

- **Product type:** Software
- **Deployment model(s):**
    - On-premises (self-managed)
    - Cloud (SaaS) (Elastic Cloud Hosted, Elastic Cloud Serverless)
    - Hybrid

Data Processed in Elastic Observability

Elastic Observability processes operational telemetry data that customers ingest to monitor their systems and applications. The specific types of personal data processed depend on the customer's ingested data and configuration. Elastic provides the platform and tools for managing this data, which may include personal data depending on the nature of the logs, metrics, application traces, or user experience data.

Categories of personal data necessary to access and use Elastic Observability:

- Admin and user identifiers, credentials, permissions, session cookies and activity logs. The processing of such data is necessary for the purposes of product deployment, configuration, administration, management, secure access and auditable use.

Additional categories of personal data processed may include, but are not limited to, data contained within:

- Individual identifiers and system/network activity data from logs, metrics, and traces:
    - Examples**:** User IDs, IP addresses, hostnames, machine names, process IDs, filenames, directory paths, URLs, email addresses (if present in logs or error messages), mobile device identifiers, and location data from user experience monitoring.
- Application performance data:
    - Examples**:** Transaction details, spans, errors, metrics, and metadata collected from applications, which might contain user-specific data if not sanitized by the customer.
- User experience data:
    - Examples**:** Data related to the perceived performance of web applications by real users, potentially including browser types, operating systems, and generalized location data.

Typical Purpose(s) of Processing Customer Data Fed to Elastic Observability:

In their specific deployments, customers have sole discretion to determine the actual purpose(s) of processing of any data they choose to feed to the product. In general, Elastic Observability is designed and intended for use cases such as:

- To provide granular insights and context into application behavior.
- To unify and correlate logs, infrastructure metrics, application traces, user experience data, synthetics, and universal profiling for faster root cause detection.
- To analyze log data from various sources such as hosts, services, Kubernetes, and Apache.
- To instrument code and collect performance data and errors at runtime through Application Performance Monitoring (APM).
- To monitor system and service metrics from servers, Docker, Kubernetes, and other services and applications.
- To quantify and analyze the perceived performance of web applications through Real User Monitoring (RUM).
- To simulate user actions and requests on a site at predefined intervals in controlled environments via Synthetic Monitoring, generating consistent and repeatable data for trending and alerting.
- To build stack traces and gain visibility into system performance without requiring application source code changes or instrumentation, helping to identify expensive code lines, increase CPU efficiency, and debug performance regressions (Universal Profiling).
- To detect anomalies and spikes in time series data, and to investigate the causes of unusual changes in log rates or metric distributions by leveraging predictive analytics and machine learning.
- To alert on potential issues in environments based on defined rules and thresholds, integrating with log monitoring and APM.
- To set clear, measurable targets for service performance and monitor Service-Level Objectives (SLOs) in real-time.
- To collect and share information about observability issues by creating cases, enabling team collaboration and integration with third-party systems like ServiceNow and Jira.

Product Usage Data

Elastic Observability automatically collects certain information related to the use of the product. Such Product Usage Data is processed as described in the Elastic Product Privacy Statement. In on-premises and certain hybrid deployment models, customers may restrict or disable Product Usage Data processing in the product's configuration settings (please refer to the applicable product Documentation). However, such Product Usage Data collection and processing is indispensable to the proper functioning of the product in Cloud (SaaS) deployment models.

Access to Customer Data

- **Access by Customers:** Customers retain full control over the data they ingest into Elastic Observability. They can access and analyze their observability data through the Kibana user interface, which provides capabilities for searching, filtering, visualizing, and analyzing logs, traces, and metric data. Customers can retrieve data directly using Elasticsearch's REST API or through various Elasticsearch clients. Role-based access controls are available within Kibana to manage user access to specific content and features, ensuring adherence to least-privilege principles.

- **Access by Elastic:** For Elastic Cloud offerings, a limited number of Elastic employees are granted privileged access to the production environment, strictly for essential platform management, maintenance, and support activities. This access is provisioned based on the principles of least-privilege and need-to-know and is subject to regular review and adjustment. Elastic does not have access to data within self-managed deployments that are not Cloud connected, and will only access customer content in cloud deployments upon the customer's instruction (e.g. when providing support services at the customer's request).

    - Elastic has implemented centralized logging, encompassing proxy logs, access logs, Elasticsearch logs, and Auditbeat logs to record all access to customer data and the systems on which it resides. Elastic's internal teams actively develop and implement detections for suspicious internal account activity and unauthorized access, including file integrity monitoring and account takeover indicators.

    - Elastic's principles dictate that it will only disclose or provide access to customer data when strictly compelled by law, and it includes established protocols for challenging such requests and notifying relevant parties where legally permissible.

Processing Locations

Elastic Cloud deployments of our Observability solution are hosted on certified cloud platforms managed by industry-leading IaaS providers that include Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure. Customers can select the geographic region where they want to host their deployments.

Sub-Processing: Elastic utilizes certain infrastructure and customer support sub-processors to deliver its services. The specific sub-processors involved depend on the data center location, services, and functionalities chosen by the customer. These sub-processors are contractually bound to provide a level of data protection at least equivalent to that described in the Elastic Information Security Addendum. Elastic maintains full transparency regarding its sub-processors (see internal and external lists) and is liable for their actions and omissions to the same extent as if Elastic performed the services itself.

Compliance with Privacy Regulations

Elastic captures, processes, stores, and protects Customer Personal Data in accordance with the applicable customer Data Processing Addendum, and the commitments outlined in this Privacy Data Sheet. Our Trust Center serves as a comprehensive resource for information on its privacy practices and compliance efforts.

- **Data Ownership:** Customer data remains the property of the customer. Elastic only processes this data for the purposes specified in the customer's agreement, and never sells customer data to third parties.
- **GDPR Compliance:** Elastic Cloud features are designed to support compliance with the General Data Protection Regulation and other global data protection and privacy laws. This includes prioritizing the security of personal data through effective technical and organizational measures, offering a GDPR-compliant Data Processing Addendum. Our dedicated privacy team at Elastic oversees GDPR compliance.
- **Data Privacy Framework (DPF):** Elasticsearch, Inc. is certified under the EU-U.S. Data Privacy Framework, UK Extension to the EU-U.S. Data Privacy Framework, and the Swiss-U.S. Data Privacy Framework as set forth by the U.S. Department of Commerce.
- **Cross-Border Data Transfer:** Elastic legalizes transfers of personal data across the jurisdictions where we operate by relying on current Standard Contractual Clauses, and through our participation in the Data Privacy Framework for transfers to the U.S. and from there onwards. Additionally, we implement robust supplementary measures to protect data during transfers, such as encryption in transit and at rest, protocols for challenging public authority requests, and providing customers with the option to select regional servers for hosting.

Data Portability

Customers can easily manage the portability of their data from Elastic Observability. You can retrieve data stored in Elasticsearch using its REST API and through various Elasticsearch clients for common programming languages. This capability supports customers in fulfilling data subject access requests. Furthermore, Elastic provides options to export Kibana dashboards and visualizations to formats like PDF, PNG, or CSV files.

Retention and Deletion

Elastic provides tools and capabilities that enable customers to manage their data retention and deletion policies effectively.

- Customers have the ability to define and enforce retention policies for the personal data they collect and process within the Elastic environment.
- Elastic can help identify unused data and informs data retention practices by allowing customers to categorize data into storage tiers and review access logs.
- Logstash, an Elastic component, can be used to perform data transformation, including anonymization and pseudonymization, which helps achieve data minimization goals and reduces data security risks.
- Elastic's platform enables organizations to more closely analyze their actual use of retained personal data, allowing them to tailor data retention periods and policies more effectively.
- For data subject deletion requests, Elastic provides capabilities to tag data for retention under an exception, or to permanently delete data using permissible deletion amd de-identification techniques such as anonymization, to help customers remain compliant within short response timeframes.

Security

The security of your data within Elastic Cloud relies on you keeping your deployments configured securely and maintaining the confidentiality of your Elastic Cloud login credentials (please see our Security FAQ for more information). In addition, Elastic supports a comprehensive defense-in-depth security model designed to protect customer data throughout its entire lifecycle: in transit, at rest, and in memory, as well as through robust key management procedures.

- **Encryption:** Customer data is encrypted at rest using AES-256 and in transit via TLS 1.2. Elastic implements strict encryption key management procedures.
- **Access Controls:** Elastic maintains technical, logical, and administrative controls to restrict data access solely to authorized users. These controls include measures such as multi-factor authentication, strong password strength standards, and the use of Virtual Private Networks (VPNs) for administrative access. Additionally, Role-based Access Controls (RBAC) are integrated into Elastic deployments and the Elastic Cloud management platform, allowing granular control over user permissions.
- **Logging and Monitoring:** Elastic has implemented centralized logging, encompassing proxy logs, access logs, Elasticsearch logs, and Auditbeat logs to record all access to customer data and the systems on which it resides. Elastic's Threat Detection and Response team utilizes automated workflows to detect and investigate suspicious internal account activity and unauthorized access.
- **System Updates and Patches:** Elasticsearch instances are regularly updated and deployed based on the latest operating system kernels. Appropriate patches are applied promptly whenever a Common Vulnerability and Exposure (CVE) is identified in any component software.
- **Incident Detection and Response:** Elastic has implemented and continuously updates detection rules for suspicious activity and unauthorized access. These detections are part of automated workflows that alert the Threat Detection and Response team, triggering analyst investigations. Elastic uses Elastic Security internally which includes Endpoint Detection and Response capabilities and integrates threat intelligence feeds to support intrusion detection. It supports automated responses through integrations with third-party SOAR platforms, reducing the likelihood of a breach and speeding response times.
- **Secure Software Development Framework (SSDF):** Elastic maintains a secure software development framework based on NIST 800-218. This framework guides the process to securely design, develop, deploy, track, and maintain all Elastic software, ensuring a "secure by design" and "secure by default" approach.
- **Third-Party Vendor Review:** Elastic partners with major IaaS providers (AWS, GCP, Azure) and conducts rigorous reviews of their security and compliance standards, including SOC 2 audits and ISO 27001 certifications, as part of its third-party risk management program.
- **Penetration Testing:** Independent third parties conduct annual application and network penetration tests for Elastic Cloud. Elastic also hosts a public Bug Bounty Program for continuous researcher testing.
- **Employee Training:** All Elastic employees are required to complete comprehensive information security and data protection/privacy training upon hire and at least annually thereafter.