Elastic Privacy Datasheet

AutoOps in Cloud Connect

Product Summary: AutoOps in Cloud Connect

The evolution of distributed search architectures has necessitated a corresponding advancement in operational management and diagnostic capabilities. Elastic AutoOps in Could Connect represents a cloud-connected utility designed to provide real-time issue detection, automated root cause analysis, and actionable performance recommendations for Elasticsearch clusters, allowing organizations to gain expert-level insights without the requirement of maintaining a comprehensive local monitoring infrastructure.

The architecture of Elastic AutoOps is founded on a cloud-connected model, which is particularly relevant for self-managed enterprise users running deployments on Elastic Cloud Enterprise (ECE), Elastic Cloud on Kubernetes (ECK), or standalone on-premises hardware. This model ensures that while diagnostic telemetry is analyzed in Elastic Cloud, the actual content residing in the customer's indices never leaves the local environment. By strictly limiting the processing of data to operational metadata and providing robust transparency controls, Elastic ensures that organizations can optimize their most complex search workloads without compromising the privacy of their data. To use AutoOps with your ECE, ECK, or self-managed Elasticsearch cluster, you first need to create an Elastic Cloud account or log in to your existing account. An installation wizard then guides you through the steps of installing Elastic Agent to send metrics from your cluster to AutoOps in Elastic Cloud.

- **Product type:** Software
- **Deployment model(s):**
    - Cloud (SaaS)

Data Processed in AutoOps

The processing of data within Elastic AutoOps is strictly bifurcated between administrative data required to operate the service and operational metadata required to diagnose cluster health. It is critical to note that AutoOps is designed to be a read-only solution regarding customer content; it does not access, fetch, or store the documents or source customer data managed within the cluster.

Categories of personal data necessary to access and use Elasticsearch:

To manage the service, certain categories of personal data are processed. This data is essential for maintaining secure access, managing subscriptions, and ensuring an auditable log of system interactions.

- **Admin and User Identifiers:** This includes names, email addresses, and phone numbers provided during account registration, support interactions, and event notification configurations.
- **Credentials and Permissions:** Usernames, passwords (encrypted), session cookies, and activity logs are processed to manage secure access and enforce role-based access control (RBAC).

Audit Logs: The system records administrative actions, such as when a cluster is registered or when notification settings are modified, to ensure accountability and security auditing.

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

- Operational Metadata and Diagnostic Metrics
    - Examples: Detailed shard information, comprehensive node statistics, cluster-wide configurations, overall health summary, index and composable template lists, active cluster tasks.
        - The collection of these metrics allows AutoOps to perform complex analyses, such as identifying shard imbalances, detecting high CPU usage, and pinpointing slow search and indexing queries. While some of these fields, such as index names or node hostnames, might indirectly identify internal infrastructure as configured by the customer, they do not contain the actual business data or end-user records stored within the indices.

Typical Purpose(s) of Processing Customer Data Fed to the Product:

The processing of metadata by Elastic AutoOps is governed by the objective of maximizing the operational efficiency of the Elasticsearch platform and is intended for use cases such as:

- Real-Time Issue Detection and Root Cause Analysis
    - AutoOps correlates hundreds of metrics to provide a unified picture of cluster health. For example, the system can identify that a transition to a "Red" cluster status was preceded by a buildup of high disk watermark events over several days, allowing administrators to address the underlying storage issue before an outage occurs. This automated root cause analysis (RCA) removes the need for manual analysis of charts and indices, which is traditionally a time-consuming process for DevOps teams.
- Resource and Cost Optimization
    - The processing of resource utilization metrics, such as CPU pressure, memory utilization, and disk I/O, enables AutoOps to highlight underutilized nodes or inefficiently configured indices. By identifying ingestion bottlenecks or data structure misconfigurations, the utility provides clear resolution paths to improve resource efficiency and lower hardware costs. For instance, the Template Optimizer view analyzes mappings and settings to suggest optimizations that reduce storage overhead.
- Support Integration and Collaborative Diagnosis
    - A significant insight into the AutoOps privacy model is the "Better Support" feature. Customers may grant Elastic Support engineers read-only access to their AutoOps diagnostics. This shared context allows engineers to see exactly what the customer sees, facilitating faster and more precise resolutions of support tickets without the need for manual diagnostic bundle uploads. This mechanism facilitates "least-privilege" support delivery, where access to diagnostic data is temporary and explicitly authorized by the customer.

Product Usage Data

For Elastic Cloud deployments utilizing AutoOps, Elastic 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.

Access to Customer Data

- **Access by Customers:**
    - Customers retain full control and ownership of the data they ingest, utilizing integrated Role-Based Access Control (RBAC) to define granular permissions and enforce the principle of least privilege. Through a unified cloud-based interface, users can access high-level metrics, real-time health events, and detailed performance visualizations for their nodes, indices, and shards.
    - To provide customers running self-managed clusters with full visibility into the data being shared, the AutoOps Elastic Agent includes a debug configuration file. By restarting the agent with this configuration, administrators can view a sample of the data gathered from the cluster locally before it is sent to Elastic Cloud. This allows for a proactive assessment of the information being disclosed and supports internal compliance reviews.
    - Customers can decide on the type of data Elastic agent can pull - Read here.

- **Access by Elastic:**
    - For Elastic Cloud offerings, a limited number of Elastic employees have privileged access to the production environment solely for platform management, maintenance, and support. This access strictly adheres to the principles of least-privilege and need-to-know and is regularly reviewed and modified as necessary. Elastic does not have access to data within self-managed deployments, 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).
    - The processing of data within Elastic AutoOps is strictly bifurcated between administrative data required to operate the service and operational metadata required to diagnose cluster health. It is critical to note that AutoOps is designed to be a read-only solution regarding customer content; it does not access, fetch, or store the documents or source event data managed within the cluster, facilitating faster and more precise troubleshooting within a shared context while ensuring the actual business data remains within the customer's local environment.
    - 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.

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 Datasheet. Our Trust Center serves as a comprehensive resource for information on our 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 retain full control over the data they ingest into their own clusters. For the metadata processed by AutoOps, users can access high-level metrics via the Overview and Nodes views. While the specific telemetry stored in the cloud is managed by Elastic, the original metrics remain available in the customer's self-managed cluster for extraction via standard Elasticsearch APIs.

Retention and Deletion

Elastic AutoOps implements strict data retention policies to align with the principle of data minimization and the transient nature of operational telemetry. AutoOps Operational Metrics are retained for a period of 10 days within Elastic's Internal Cloud Infrastructure. The 10-day retention period for AutoOps metrics ensures that recent performance trends are available for diagnosis while older data, which typically loses its operational relevance in dynamic environments, is automatically purged.

Security

Elastic implements a defense-in-depth security model designed to protect customer data throughout its entire lifecycle. Access to the systems processing AutoOps data is governed by strict organizational and technical measures. The security of your Elastic Cloud data also relies on you keeping your Elasticsearch cluster configured securely and maintaining the confidentiality of your Elastic Cloud login credentials (please see our Security FAQ for more information).

- **Auto-Ops Architecture**: The architecture of Elastic AutoOps is particularly relevant for self-managed enterprise users running deployments on Elastic Cloud Enterprise (ECE), Elastic Cloud on Kubernetes

(ECK), or standalone on-premises hardware. This model ensures that while diagnostic telemetry is analyzed in the cloud, the underlying business data, the actual content residing in the customer's indices, never leaves the local environment.

- **Communication Protocols**: Elastic AutoOps utilizes two primary communication protocols to ensure secure telemetry ingestion. Both channels operate over standard Port 443, ensuring compatibility with common firewall configurations while maintaining high encryption standards.
    - Registration Channel (HTTP/HTTPS): This channel is used to register the local cluster with Elastic Cloud. It utilizes an Elastic Cloud API key that is strictly limited for use with the Cloud Connect framework.
    - Telemetry Channel (OTLP over HTTP): The OpenTelemetry Protocol (OTLP) is used to gather operational data. This channel is authenticated using an AutoOps token, which is functionally equivalent to an API key and ensures that only authorized agents can stream data to the specific customer organization.
- **Encryption:** Customer data is encrypted at rest using AES-256 and in transit via TLS 1.2 or higher. Elastic implements strict encryption key management procedures to ensure that access to encrypted data is limited to authorized processes and users. For hosted cloud deployments, customer data is also protected through robust key management integrated into the cloud service provider (CSP) infrastructure.
- **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, including file integrity monitoring and account takeover indicators.
- **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 and 27701 certifications, as part of its third-party risk management program.
- **Penetration Testing:** Independent third parties conduct annual application and network penetration tests against Elastic Cloud, at a minimum. 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.