splunk> turn data into doing

Privacy and Security in the Splunk Cloud Platform

September 2022

Customer FAQ

Splunk is committed to providing our customers with straight-forward, transparent explanations about how we process and secure Personal Data in the Splunk Cloud Platform. As part of that commitment, we created this FAQ to provide helpful information about the privacy and security we offer and to serve as a resource for you when you are completing data protection impact assessments. For more detailed information about Splunk's privacy and security programs, please visit Splunk Protects.

General Description: Splunk Cloud Platform

What does the Splunk Cloud Platform do?

The Splunk Cloud Platform ("Platform") is a search engine for machine-generated data that provides customers with real time insights into the performance and security of their IT infrastructure.

What are log files?

Every time a computing device takes an action, it notes it in a file. For example, when a user logs into a network server, the server may record "on X date, at Y time, user jsmith logged in and typed the correct password." These notations ("log entries") are appended into a file ("log file") that is ingested into the Splunk Cloud Platform. Log files are created on all types of computing devices, e.g., network servers, routers, firewalls, laptops, mobile phones, iWatches, etc. In short, any system or machine that processes data.

Log files provide important information about how devices or IT systems within them are performing and how they are secured. For example, log files from network file servers may indicate that a network hard drive is performing erratically and needs to be replaced before it fails; hundreds of unsuccessful logins from a single IP address may suggest that a brute force attack has occurred on a user account.

Personal Data Processed in the Platform

What categories of Personal Data are processed in the Platform?

Log files typically only contain minor elements of Personal Data, such as a user ID or IP address. In almost all cases, identifying an individual from log file data can be difficult without corroborating, external information that is not normally ingested into the Platform. Splunk does not curate or have any visibility into customer data that is ingested into the Platform. However, customers can self-manage their data (including Personal Data) within the Platform by applying hashing, redaction or suppression techniques prior to or after the data is ingested into the Platform. For more, see guidance provided in the Splunk Cloud Service Description.

Does the Personal Data processed in the Platform include special categories of Personal Data?

Log files almost never contain any special categories of Personal Data as defined under Article 9 of the GDPR.
Splunk Customers are responsible for ensuring that submission of any special categories of Personal Data complies with applicable laws.

Who can access Personal Data processed in Splunk Cloud Platform?

A list of Splunk's sub-processors can be found on our Website Page for Sub-processors. Splunk offers data hosting globally in select AWS and GCP regions. Customers can choose the region where their data is hosted. Data hosting can be limited to within the EEA. Processing (as defined by the GDPR) is performed in the United States and other locations as set forth in the link above to provide global cloud operations and support. Customers may sign up to receive notifications of any changes to Splunk's sub-processors through its Notification Portal.

International data transfers

Splunk may transfer data from the EEA to countries not deemed "adequate" by the European Commission and from the UK to countries not deemed "adequate" by the UK Information Commissioner's Office. These locations are set forth in the sub-processor section referenced above. In such cases, Splunk relies on the standard contractual clauses pursuant to European Commission Decision 2021/914/EU ("EU Clauses") including the modules for controller-processor with its customers and the modules for processor-processor with its sub-processors.

Splunk has reviewed the EDPB guidance regarding supplementary measures for international data transfers and the specific questions raised by NOYB in their questionnaire regarding international data transfers post-Schrems II. Splunk has published a whitepaper and issued responses to the NOYB which are available on its website and which detail how Splunk meets the requirements specified in both. Please see A Risk Assessment of EU Cross-Border Data Transfers and Splunk's Responses to the European Center for Digital Rights (nyob) Questions. Splunk provides these materials in compliance with its obligations under Article 28(f) of the GDPR to provide customers who are controllers with assistance in completing a transfer impact assessment.

What's Splunk's policy regarding law enforcement requests?

Splunk's practices for responding to requests by government agencies and other third parties for customer data are detailed in the Splunk Data Request Guidelines

Protecting Personal Data Processed in the Platform

Splunk's security and privacy programs meet the highest standards in the industry and are further set forth in the Splunk Cloud Platform Security Addendum. This includes such things as:

- GDPR required breach notification
- Policies, Practices and Training
- Access and user management
- Governance and audit management
- Password management and authentication controls
- Encryption at rest and in transit
- Threat and vulnerability management
- Logging and monitoring
- Secure software development
- Network Security, Physical Security, Disaster Recovery Plans
- Asset Management and Disposal
- Human Resources Security
- Splunk Vendor Security
- Annual third-party audits for ISO 27001, Soc2, Type2, HIPPA/PCI-DSS, FedRAMP, IL5

Further, if you elect AWS as your hosting provider, Splunk leverages the Amazon Web Service Key Management Service (AWS KMS) to create and maintain a primary encryption key used to secure data on your Splunk Cloud deployment. KMS is a fully managed service, backed by Federal Information Processing Standards (FIPS)-140 hardware security modules, that is also supported on PCI and HIPAA deployments. With this model, Splunk is responsible for the management of the keys, including all creation, rotation, and revocation operations.

Splunk also offers EMEK as an optional capability for encrypting data at rest, which allows you to provide your own primary encryption key. By leveraging this capability, Splunk Cloud Platform administrators can grant and subsequently rotate, revoke or disable access to your complete data set while maintaining the same degree of real-time data encryption and decryption operations that comes with the managed-service model. EMEK gives you the flexibility of managing the encryption key yourself, which ensures you maintain complete control of your Splunk Cloud Platform deployment.

If you enable EMEK for your deployment, note that you are the sole controller of the master encryption key and, by design, you cannot unlock your data without this key. This means that disabling or deleting the EMEK key will result in permanent loss of data access. Please see Secure Data with Enterprise Managed Encryption Keys - Splunk Documentation) for more information.

Splunk employs robust security by encrypting customer data in transit, including data flowing across transatlantic cables, and at rest. Customers can also purchase additional encryption at rest at the application layer to prevent unauthorized access by third parties. Splunk encryption standards currently include:

- TLS 1.2+ (in transit) and AES 256 (at rest)
- Industry standard encryption tools vetted to meet Splunk's security standards
- A requirement that Splunk sub-processors encrypt data in transit
- Regular rotation and monitoring of keys

Splunk Monitors Access to Your Data

Splunk continuously monitors your instance to detect and investigate suspicious activity. Splunk employs a Host-based Intrusion Detection, which logs and monitors access attempts and uses automatic alerts to trigger investigation and incident management procedures in certain cases.

Splunk Cloud Platform provides frameworks that prevent unauthorized access to the platform and the data that you store in it these frameworks include, but are not limited to:

- Role-based access control (RBAC) to help you manage access to your Splunk Cloud Platform instance
- Security of configurations, data ingestion points, data storage and internal/external communications using various certificates and encryption schemes
- Obfuscation of log-in credentials

Splunk has a range of both policy-based and technical controls in place to prevent and detect unauthorized access. Splunk has a deny-by-default policy with regards to Splunk employee access to customer data. This includes documented approval systems to request access, scoped/time-limited access based on the business requirement, 24/7 logging and monitoring (including for insider threat scenarios), external and internal audit and contractual agreements.

If you request support which would require Splunk to access your data, you must explicitly authorize us to do so through a customer-initiated support ticket. Access is managed under the principle of "least privilege", using scoped and/or ephemeral access tokens to ensure that access is granted only insofar as is required to resolve your specific issue and only through a secure Virtual Desktop Infrastructure (VDI), which has Data Loss Prevention (DLP) controls built in, such as disabled USB, limited internet ingress/egress, etc. Splunk support users will appear in your audit index, and alerts can be set up for this activity. For more details, please see Secure Data with Enterprise Managed Encryption Keys - Splunk Documentation).

Privacy by Design

Splunk Cloud Platform has numerous built-in features which integrate privacy by design principles.

- **Data Collection**. Customers can restrict data collection from only allowed IP addresses by using the Administrative Configuration Service (ACS).

- **Data Anonymization**. Splunk supports advanced anonymization to remove confidential data from the data that you index into Splunk Cloud Platform. You can anonymize parts of confidential fields to protect privacy while providing enough remaining data for use in event tracking. For more details, please see

Getting Data In.

During support, a customer may generate a diagnostic file (diag file) and send it to their support representative to help diagnose the problem. Diag files give Splunk support insight into how a Splunk instance is configured and operating. If created according to Splunk's documented instructions, diag files do not contain customer data, including Personal Data; for additional protection, customers also can redact or anonymize data within the diag file through the functionality of the Splunk Cloud Platform service prior to sending to Splunk.

- **Data Integrity**. The Splunk data integrity control feature allows you to verify the integrity of indexed data. For more details, please see Manage Data Integrity - Splunk Documentation).

  Indexed data can be hashed to ensure fidelity over time, giving you confidence that your data hasn't been altered. Individual events and streams of events can be signed. In addition, message integrity measures show whether an event has been inserted or deleted from the original data stream.

- **Data Deletion, Return and Portability**. Functionalities to retain, delete or export (return) data are built into the Platform. Customers choose how long to retain / delete their data, including Personal Data, within the Platform. Customers can tailor retention options by country and for any duration required (additional fees may apply). Splunk's DPA reflects the self-help design of the Platform for GDPR purposes.

- **Data Subject Requests**. The Platform is equipped with self-help capabilities. The customer can respond to Data Subject Requests via the data deletion, return and portability functions set forth above

- **Data Segregation**. While the Platform service is multi-tenant, it enforces logical separation of Customer data.

Data Hosting Location

As of June 2022, Splunk offers the following data hosting locations. Customers have the option of limiting data hosting to the EEA or the UK. For more details and a current list of data hosting locations, please see Splunk Cloud Platform Service Details.

AWS regions:
- US (Oregon, Virginia, GovCloud-West, GovCloud-East)
- UK (London)
- Europe (Dublin, Frankfurt, Paris, Stockholm)
- Asia Pacific (Seoul, Singapore, Sydney, Tokyo, Mumbai)
- Canada (Central)

Google Cloud regions:
- US (Iowa)
- UK (London)
- Europe (Frankfurt)
- Asia Pacific (Singapore, Sydney)
- Canada (Montreal)