SPLUNK INC.

INDEPENDENT PRACTITIONER'S TRUST SERVICES REPORT FOR THE SPLUNK CLOUD SYSTEM

FOR THE PERIOD OF OCTOBER 1, 2016, TO AUGUST 31, 2017

Attestation and Compliance Services

INDEPENDENT PRACTITIONER'S TRUST SERVICES REPORT

To the Management of Splunk Inc.:

We have examined management's assertion that during the period October 1, 2016, to August 31, 2017, Splunk Inc. ("Splunk") maintained effective controls over the Splunk Cloud system (the "system"), for the security, availability, and confidentiality principles set forth in the 2016 TSP section 100A, *Trust Services Principles and Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy (AICPA, Trust Services Principles and Criteria)* (applicable trust services criteria), to provide reasonable assurance that

- the system was protected against unauthorized access, use, or modification to meet the entity's commitments and system requirements;
- the system was available for operation and use to meet the entity's commitments and system requirements;
- information designated as confidential is protected to meet the entity's commitments and system requirements.

As indicated in the description, Splunk uses Amazon Web Services, Inc. ("AWS") for all of Splunk's cloud hosting services. The description indicates that certain applicable trust services criteria can be met only if certain types of controls that management expects to be implemented at AWS are suitably designed and operating effectively. The description presents Splunk's system; its controls relevant to the applicable trust services criteria; and the types of controls that Splunk expects to be implemented, suitably designed, and operating effectively at the AWS to meet certain applicable trust services criteria. The description does not include any of the controls expected to be implemented at AWS. Our examination did not extend to the services provided by AWS, and we have not evaluated whether the controls management expects to be implemented at AWS have been implemented or whether such controls were suitably designed and operating effectively throughout the period October 1, 2016, to August 31, 2017.

Splunk's management is responsible for the attached assertion. Our responsibility is to express an opinion based on our examination. Management's description of the aspects of the Splunk Cloud system covered by its assertion is attached. We did not examine this description, and accordingly, we do not express an opinion on it.

Our examination was conducted in accordance with attestation standards established by the American Institute of Certified Public Accountants and, accordingly, included (1) obtaining an understanding of Splunk's relevant controls over the security, availability, and confidentiality of the Splunk Cloud system; (2) testing and evaluating the operating effectiveness of the controls; and (3) performing such other procedures as we considered necessary in the circumstances. We believe that our examination provides a reasonable basis for our opinion.

Because of the nature and inherent limitations of controls, Splunk's ability to meet the aforementioned criteria may be affected. For example, controls may not prevent or detect and correct error or fraud, unauthorized access to systems and information, or failure to comply with internal and external policies or requirements. Also, the projection of any conclusions based on our findings to future periods is subject to the risk that changes may alter the validity of such conclusions.

In our opinion, management's assertion referred to above is fairly stated, in all material respects, based on the AICPA and CPA Canada applicable trust services criteria.

Schellman & Company, LLC

Tampa, Florida
October 4, 2017

MANAGEMENT'S ASSERTION

October 4, 2017

During the period October 1, 2016, through August 31, 2017, Splunk Inc. ("Splunk") maintained effective controls over the Splunk Cloud system (the "system"), for the security, availability, and confidentiality principles set forth in the 2016 TSP section 100A, *Trust Services Principles and Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy (AICPA, Trust Services Principles and Criteria)* (applicable trust services criteria), to provide reasonable assurance that:

- the system was protected against unauthorized access, use, or modification to meet the entity's commitments and system requirements;
- the system was available for operation and use to meet the entity's commitments and system requirements;
- information designated as confidential is protected to meet the entity's commitments and system requirements.

The attached system description identifies the aspects of the Splunk Cloud system covered by the assertion.

Timothy Emanuelson
Vice President, Controller

System Description of the Splunk Cloud System

Company Background

Splunk Inc. was founded in 2003 to pursue a disruptive new vision: make machine data accessible, usable and valuable to everyone. Splunk offers software solutions that enable organizations to search, analyze, and visualize, machine-generated Big Data. Machine data is one of the fastest growing and most valuable parts of big data - generated by every component of information technology (IT) infrastructures, applications, mobile devices, website clickstreams, social data, sensors and more.

Splunk is a software platform for machine data that enables customers to gain real-time Operational Intelligence. Splunk's mission is to address the challenges and opportunities of managing massive streams of machine-generated big data. Splunk customers use the software to harness the power of their machine data for application management, IT operations, security, web intelligence, customer and business analytics and more.

Description of Services Provided

The Splunk Cloud provides operational intelligence capabilities to organizations around the world. It provides an easy and fast way to analyze the massive streams of machine-generated data coming from IT systems and applications, wherever they are deployed (on-premises, in virtualized environments, or in the cloud).

Splunk Cloud delivers the Splunk Enterprise software in a cloud environment, where the setup and management of the support infrastructure is performed by Splunk Cloud Operations (Cloud Ops).

The Splunk Cloud provides the following:

- Splunk Enterprise delivered as a service – no waiting on hardware or personnel resources
- Troubleshoot, secure, monitor, and analyze IT infrastructure and applications wherever they are deployed: on–premises or in the cloud
- Gain visibility into customer transactions, user experience, and user behavior across your on-premises and cloud-based applications and infrastructure

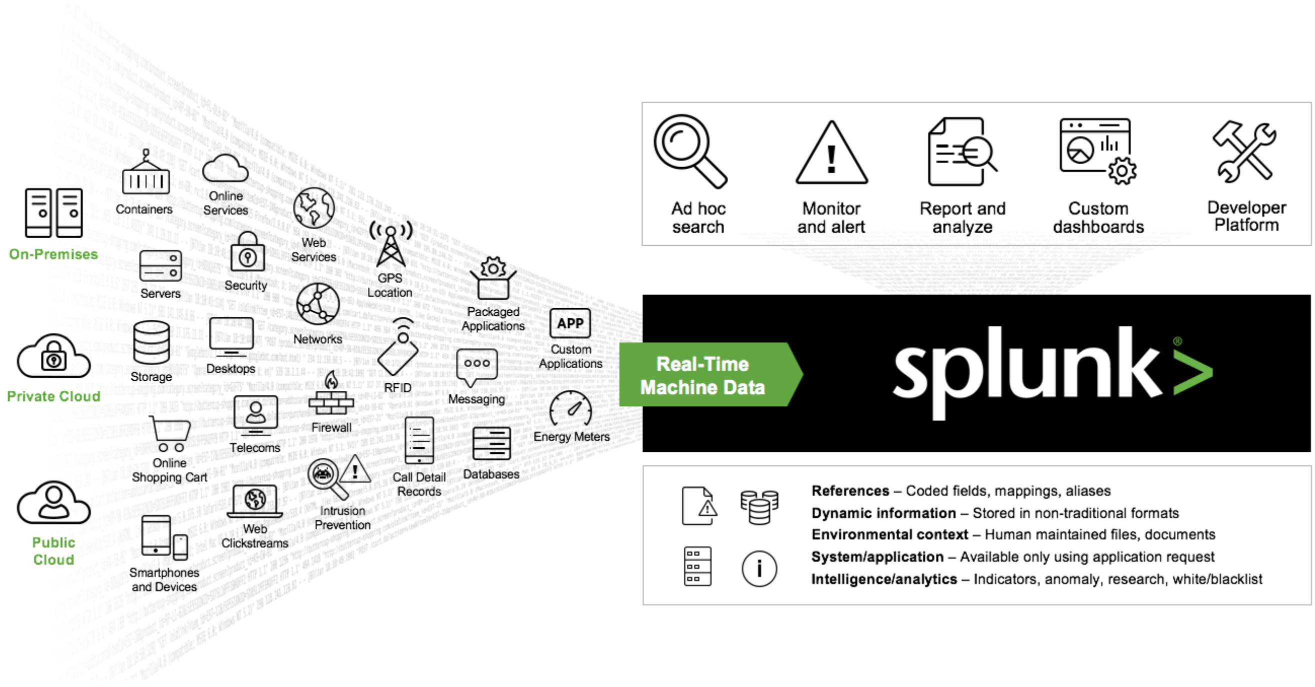

References – Coded fields, mappings, aliases
Dynamic information – Stored in non-traditional formats
Environmental context – Human maintained files, documents
System/application – Available only using application request
Intelligence/analytics – Indicators, anomaly, research, white/blacklist

The scope of this review is limited to the Splunk Cloud enterprise customer environments provisioned by the Provisioning software. Customer environments provisioned on the Microservices provisioning platform are covered under a separate report.

Infrastructure and Software

The Splunk Cloud service encompasses two primary applications sitting within the Amazon Web Services infrastructure:

- Splunk Cloud (customer stack based on the most current cloud version of Splunk Enterprise)
- Splunk Provisioner (StackMakr provisioning software)

Splunk Cloud consists of multiple components which may be installed on an individual host or many hosts:

- Search Head(s) – Server that manages search requests the individual stacks and provides the customer facing user interface (UI)
- UI – Console in AWS where users can manage their instance of Splunk
- Indexer – Stores and indexes data to facilitate the search process
- License Manager – Stores customer license related configurations

Splunk provisioning consists of an aggregation of processes and tools used to deploy a customer stack:

- GitHub – File sharing repository used to store configuration files for the client
- Jenkins – Component of provisioning used to gather and store scripts
- Chef – Automation platform which stores instructions (cookbooks) on directions for deploying and maintaining individual Splunk Cloud instances
- Whisper – Source code base developed by Splunk Cloud Engineering that defines how Splunk runs in the Cloud
- 1SOT (Artifactory) – Directory storing the Splunk Enterprise binary file and secure shell (SSH) keys
- Nagios – Log monitoring solution
- Ansible – Agentless tool used to run tasks on server machines
- AWS Console – Internal operations console for managing hardware infrastructure
- StackMakr – Provisioning application to build Splunk Cloud stacks in AWS for customer environments
- Salesforce – Frontend system that manages the customer relationship lifecycle

The in-scope infrastructure consists of multiple applications, operating system platforms and databases, as shown in the table below:

Primary Infrastructure			
Production Application	**Business Function Description**	**Operating System Platform**	**Physical Location**
StackMakr	Front-end Provisioning software application utilized to build Splunk Cloud stack(s) within AWS.	Amazon Linux - Amazon Machine Image (AMI) MySQL	AWS
Okta	Access to the production environment performed through multi-factor authentication.	Cloud Delivered	Cloud Delivered

Primary Infrastructure			
Production Application	**Business Function Description**	**Operating System Platform**	**Physical Location**
LastPass	Access to the production environment performed through multi-factor authentication.	Cloud Delivered	Cloud Delivered
Virtual Private Network (VPN) Appliance	Access to the production environment requires external users to authenticate through a VPN.	Palo Alto Global Protect VPN	San Jose, CA
AWS	Virtualized infrastructure (servers/storage) to host Splunk for supporting strategies (monitoring, security testing, etc.).	Cloud Delivered	Cloud Delivered
Amazon Simple Storage Services (S3)	Virtualized infrastructure (servers/storage) to host Splunk- dedicated infrastructure for customers.	Amazon AWS	AWS
Windows Active Directory (AD)	AD domain utilized to control access to the corporate and production networks.	Windows 2008 R2 Standard and Windows 2012 R2 Standard	Equinix

People

Personnel directly involved in the operation and use of the system are the Cloud Team. The Cloud Team is composed of the Cloud Operations (Cloud Ops), Cloud Engineering and Cloud Support team. The Cloud Team functions have operationalized an effective framework based on the Security and Confidentiality principles and criteria. Leadership regularly reviews and updates respective policies, training materials, compliance updates, and conducts internal and external risk assessments of the control environment. Product, Human resources and Legal provide supporting roles.

Cloud Ops

The Cloud Ops team is heavily involved in daily activities pertaining to the Splunk Cloud environment.

Cloud Engineering Team

The Cloud Engineering team is responsible for delivering the ability to provision and run Splunk in Cloud. They produce the underlying product that then Cloud Ops runs as a service.

- UI Engineer – Software engineer primarily focused on building internal UI tools for operations and customer stack management.
- Dev Engineer – Software engineer primarily focused on building the StackMakr scripts used to provision and manage customer stacks.
- Test Engineer – Quality engineer primary focus on building the test framework and test functionality for operations.
- Product Manager – Primarily focused on defining requirements and product strategy for the system.
- Project Manager – Primarily focused on the overall project planning and resourcing.

Cloud Support

The Cloud Support team provides on-going support and remediation efforts for customers during Splunk Cloud fixes, updates, and upgrades. Additionally, the following groups support all lines of business within Splunk and were not directly responsible for the operations of the Splunk Cloud.

Security

The Security team ensures Splunk Cloud deploys best practices to protect assets and third party service providers adhere to Splunk Cloud policies and procedures. They develop and implement security awareness campaign and training program, threat and vulnerability management as well as incident response. Also enable the business to securely operate within necessary technologies.

Human Resource (HR)

The HR team is responsible for recruitment, resource management, and organizational development and training to support the ability for Splunk Cloud to hire and retain personnel with appropriate skill. The HR team also supports employee development, employee welfare, and performance management.

Leadership ensures all personnel adhere to HR procedure documentation set forth. The HR procedure documentation outlines the pre-employment, on-boarding, performance review, training, backup personnel, and off-boarding policies and procedures.

Legal / General Counsel (GC)

Legal personnel work closely with the chief information officer (CIO) & chief information security officer (CISO) in formulating relevant policies. The Legal team is tasked with reviewing all company policies and procedures including but not limited to IT Security Policy, Acceptable Use Policy, Privacy Policy, and Information Management policy. Annually, the team will review and policies to address growing needs and concerns of the company.

Procedures

Personnel directly involved in the operation and use of the system are the Cloud Team. The Cloud Team is composed of the Cloud Operations (Cloud Ops), Cloud Engineering and Cloud Support team. The Cloud Team functions have operationalized an effective framework based on the Security and Confidentiality principles and criteria. Leadership regularly reviews and updates respective policies, training materials, compliance updates, and conducts internal and external risk assessments of the control environment. Product, Human resources and Legal provide supporting roles.

Data

Customer data is held in accordance with the relevant security and availability policies, data protection and other regulations, and contractual requirements. All processing is carried out on the Splunk Cloud platform product. All customer data is treated as confidential and subject to Splunk’s Terms of Service, Acceptable Use Policy, Privacy Policy, IT Security Policy, and Information Management Policy.

Data that is residing in the Indexer is unencrypted by default, but is in a compressed proprietary format. There are logical access controls to restrict direct access to the data. As a premium service enhancement, customers may elect to have Indexer data encrypted at rest.

Data stored by the Splunk Cloud SH is held on AWS Elastic Block Store (EBS) volumes. These are network attached block devices where AWS manages the encryption keys via AWS Key Management Services (KMS). By this method, AWS encrypts data at rest by default and guarantees the security of encryption keys. Splunk relies on the built-in encryption that EBS offers to provide data confidentiality with no access to the raw encryption key material. As for the root volumes that have regular Operating System files like binaries for Splunk and other systems, they do not store any customer data and thus are not encrypted. The encryption occurs on the servers that host elastic compute cloud (EC2) instances, providing encryption of data-in-transit from EC2 instances to EBS storage.

Splunk Cloud production is logically and physically separated from Splunk Corporate and requires additional layers of authentication. Customer stacks are logically separated from each other using AWS Security Groups.

The following table describes the information used and supported by the system.

Data Used and Supported by the System		
Data Description	**Data Reporting**	**Classification**
The primary types of data handled by Splunk are network, system, and application log files. These files can come in a variety of forms and the Splunk engine collects, processes, and analyzes the data.	Customer data is stored at AWS environments. Output data is available to customers via the customer Splunk Cloud environment.	Confidential

Significant Changes During the Review Period

No relevant changes to the Splunk Cloud system occurred during the review period.

System Boundaries

As outlined in the 2016 TSP section 100A, *Trust Services Principles, Criteria, and Illustrations for Security, Availability, Processing Integrity, Confidentiality, and Privacy,* a system is designed, implemented, and operated to achieve specific business objectives (for example, delivery of services, production of goods) in accordance with management-specified requirements. The purpose of the system description is to delineate the boundaries of the system, which includes the services outlined above and the five components described below: infrastructure, software, people, procedures and data.

Subservice Organizations

The cloud hosting services provided by AWS were not included within the scope of this examination. The following table presents the applicable Trust Services criteria that are intended to be met by controls at AWS, alone or in combination with controls at Splunk, and the types of controls expected to be implemented at AWS to meet those criteria.

Control Activity Expected to be Implemented by AWS	Applicable Trust Services Criteria
AWS is responsible for managing logical access to the underlying network, virtualization management, and storage devices for its cloud hosting services where Splunk systems reside.	CC5.1 – CC5.4, CC5.6
AWS is responsible for implementing controls for restricting physical access to data center facilities, backup media, and other system components including firewalls, routers, and servers.	CC5.5
AWS is responsible for implementing controls for the transmission, movement, and removal of the underlying storage devices for its cloud hosting services where Splunk systems reside.	CC5.7