



CISA
CYBER+INFRASTRUCTURE



National Cybersecurity Protection System (NCPS) Cloud Interface Reference Architecture

Volume 1 - General Guidance

December 12, 2019

Version 1.0

Cybersecurity and Infrastructure Security Agency
Cybersecurity Division

Revision/Change Record

Version	Date	Revision/Change Description	Section/Pages Affected
Version 1.0	12/12/2019	Initial Release Version	All

Document Status

This document is a draft and open for public comment. The Cybersecurity and Infrastructure Security Agency is requesting feedback and comments through January 31, 2020.

EXECUTIVE SUMMARY

The National Cybersecurity Protection System (NCPS) program is evolving to ensure that security information about cloud-based traffic can be captured and analyzed and Department of Homeland Security (DHS) analysts can continue to provide situational awareness and support to the agencies. To support this goal, DHS is developing a cloud-based architecture to collect and analyze agency cloud security data. This document explains how agencies can interact with that system. It includes background about how the cloud impacts NCPS, discusses what security information needs to be captured in the cloud and how it can be captured, and provides use cases to explain how that information can be sent to DHS.

The *NCPS Cloud Interface Reference Architecture* will be released as two individual volumes. This first volume provides an overview of changes to NCPS to accommodate the collection of relevant data from agencies' cloud environments. The second volume, to be released at a later date, will provide individual use cases for how agencies can send cloud-specific data to the NCPS cloud-based architecture, with details that are specific to cloud service models and individual cloud service providers (CSPs).

A cloud-based NCPS architecture is currently in development at DHS. This *NCPS Cloud Interface Reference Architecture* is being released to Federal Civilian Agencies in advance of a deployed system in order to:

- Notify agencies about changes in the NCPS program and give them time to plan.
- Solicit feedback from agencies so that a final version of this reference architecture provides desired content and meets the needs of agencies.
- Gather requirements from agencies to ensure the cloud-based NCPS architecture can support agency use cases.

CONTENTS

1 INTRODUCTION.....	1
1.1 Document Organization	1
1.2 Purpose.....	2
1.3 Audience	2
1.4 Assumptions.....	2
1.5 Constraints	2
2 BACKGROUND	4
2.1 NCPS Overview	4
2.2 How Cloud Impacts NCPS	5
3 DHS CLOUD DATA COLLECTION	9
3.1 Cloud Security Logs	9
3.2 Cloud Log Aggregation Warehouse Overview.....	10
3.2.1 Agency Data Collection.....	10
3.2.2 Agency Data Transfer.....	11
3.2.3 Analysis of Agency Data	12
4 CONCLUSION	13
APPENDIX A: TRANSACTION & EVENT SECURITY LOGS.....	14
APPENDIX B: FLOW RECORD COLLECTION LOCATION	15

INDEX OF FIGURES

Figure 1: Current On-Premise Telemetry Configuration.....	5
Figure 2: On-Premise and Cloud Telemetry Configuration	6
Figure 3: NCPS Functions and Responsible Parties (By Deployment Model)	8
Figure 4: CLAW Pipeline (Agency Perspective).....	10
Figure 5: Responsibility for Transferring Security Data (Agency vs. DHS).....	11
Figure 6: Agency Log Ingestion (Autonomy Preserved with Log Isolation)	12
Figure 7: Network Flow Log Generation Positions for IaaS	15

DRAFT

1 INTRODUCTION

Federal civilian departments and agencies¹ are required to meet the requirements of the National Cybersecurity Protection System (NCPS).² In general, this means that the Department of Homeland Security (DHS) Cybersecurity and Infrastructure Security Agency (CISA) monitors the flow of agency network traffic and network flow logs are forwarded to DHS. DHS analysts use this data for 24/7 situational awareness, analysis, and incident response. Traditionally, network flow data has been collected by NCPS sensors located at Trusted Internet Connections (TIC) and Managed Trusted Internet Protocol Service (MTIPS) gateways, which capture security information as traffic passes between the agency and the Internet. As agencies move their information technology (IT) infrastructure to the cloud, some of their network traffic no longer traverses traditional NCPS sensors, and security information about that traffic is no longer captured by NCPS.

The NCPS program is evolving to ensure that security information about cloud-based traffic can be captured and analyzed and DHS analysts can continue to provide situational awareness and support to the agencies. To support this goal, DHS is deploying a cloud-based architecture to collect and analyze agency cloud security data. This document explains how agencies can interact with that system.

1.1 Document Organization

This document is structured to facilitate readability and ease of use. *NCPS Cloud Interface Reference Architecture: Volume 1* consists of four sections and two appendices.

- Section 1 provides a document overview, assumptions, and constraints.
- Section 2 presents an overview of NCPS and describes how the adoption of cloud computing impacts the program.
- Section 3 describes an environment that DHS is developing to collect and process NCPS-relevant data from cloud deployments of federal civilian agencies and also provides general information on how an agency can deploy a computing environment in the cloud and at the same time participate in NCPS.
- Section 4 offers summary information.
- Appendix A discusses cloud transaction and event security logs.
- Appendix B explores the various locations at which network flow information can be collected.

NCPS Cloud Interface Reference Architecture: Volume 2 is a companion document that provides vendor-specific guidance for commonly used cloud service providers (CSPs) and presents multiple use cases that can be used to inform agency implementers on best practices and considerations for different deployment scenarios. Volume 2 will be updated frequently as cloud service offerings evolve.

¹ For the purposes of this document, the term “agency” will hereinafter be used to refer to all federal civilian executive branch departments and agencies.

² <https://www.dhs.gov/cisa/national-cybersecurity-protection-system-ncps>

1.2 Purpose

A reference architecture is an authoritative source of information about a specific subject area that guides and constrains the instantiations of multiple architectures and solutions. The purpose of this reference architecture is to explain what information agencies need to capture in the cloud for NCPS, how that information can be captured, and how it can be sent to DHS. This reference architecture is divided into two volumes:

1. Volume 1 of the *NCPS Cloud Interface Reference Architecture* provides general guidance for agencies on subscribing to NCPS in the cloud. The discussion in Volume 1 is vendor-agnostic and not specific to any given cloud service provider.
2. Volume 2 of the *NCPS Cloud Interface Reference Architecture* contains reference solutions for how agencies can participate in NCPS in the cloud under different cloud service models. These solutions will include vendor-specific guidance unique to individual CSPs (e.g., Amazon Web Services (AWS), Microsoft Azure, etc.).

1.3 Audience

This document is designed primarily for the federal civilian agencies, contractors, and vendors that are required to comply with the NCPS program. This document can also be leveraged by stakeholders ranging from policy, acquisition, technical, and cybersecurity personnel to agency information technology leadership (e.g., Chief Information Officers (CIOs) and/or Chief Information Security Officers (CISOs)). Non-federal organizations may also derive value from this document as programs, strategies, and approaches are considered to address cloud security needs.

1.4 Assumptions

The following assumptions were used in the development of this reference architecture:

1. DHS will expand NCPS to include cloud data sources (rather than develop a new program to accommodate this new deployment model).
2. Agencies will continue to seek DHS assistance in securing their data by participating in NCPS.
3. Cloud computing products and services will continue to evolve and expand and their adoption by Federal Civilian Executive Branch agencies will increase.
4. Federal cybersecurity policy will permit agency security data hosted on cloud services to be accessed directly by DHS (rather than through agency on-premise infrastructure).
5. Agencies are expanding the use of encryption for all types of data and encryption is expected to become increasingly common in the future.
6. DHS's initial telemetry requirements can be satisfied with packet header-level details and do not require payload decryption.

1.5 Constraints

The following constraints were used in the development of this reference architecture:

1. Agencies remain as data owners for all cloud telemetry and are merely sharing a copy of that data with DHS.
2. DHS makes efforts to reduce costs to agencies for sending cloud telemetry to DHS. However, agencies may still incur financial expense to fully participate in NCPS in the cloud.
3. DHS and agencies will have a written and signed Memorandum of Understanding (MOU) which governs the information sharing and handling relationship between both parties.
4. DHS information collection and use shall comply with public Privacy Impact Assessments (PIA) for the NCPS program.
5. Richness of telemetry shared with DHS is bound by the agency's encryption policy. If the agency does not perform encryption "break and inspect" functions, the agency and DHS will both be unable to observe traffic payload details.

DRAFT

2 BACKGROUND

NCPS is an integrated system-of-systems that delivers a range of capabilities, including intrusion detection, analytics, intrusion prevention, and information sharing capabilities that defend the civilian federal government's information technology infrastructure from cyber threats. The NCPS capabilities, operationally known as EINSTEIN, are one of a number of tools and capabilities that assist in federal network defense.

NCPS sensors are integrated into TIC access points. As such, agencies have traditionally been able to fulfill NCPS requirements simply by complying with the TIC program. However, in 2019, Office of Management and Budget (OMB) issued an updated TIC policy, OMB Memorandum M-19-26³, which does not require TIC access points to be embedded in all TIC use cases. Many of these new TIC use cases describe cloud services. In these use cases, network traffic between an agency and a CSP does not pass through an NCPS sensor.

As agencies adopt cloud environments and conform to the new TIC use cases, they may still need to share telemetry with DHS. This document provides guidance on how agencies can share telemetry with DHS and fulfill the requirements of NCPS when operating in cloud environments. It furthers the NCPS objective to support “cyber” information sharing between DHS and federal agencies in order to enable a shared situation awareness between DHS and federal networks. Under this platform, DHS and agencies gain increased security visibility and enhance existing incident response capabilities needed to tackle modern cyber threats on U.S. networks.

2.1 NCPS Overview

Traditionally, TIC access points (either MTIPS gateways⁴ or agency-managed TIC Access Points⁵) contain EINSTEIN⁶ sensors, so when an agency participates in the TIC program, they also automatically utilize the capabilities of the NCPS program. EINSTEIN 1 (E1) monitors the flow of network traffic (i.e., network flow records) to and from a Federal civilian executive branch agency's on-premise networks. EINSTEIN 2 (E2) is an intrusion detection service that identifies potentially malicious network activity in Federal government network traffic based on specific known signatures.⁷

Today, both E1 and E2 are deployed and screen all network traffic that is routed from an agency through TICs, MTIPS, and the EINSTEIN 3A NEST⁸ locations. For E1 and E2, the agency's telemetry, in the form of network traffic, is forwarded to DHS, and DHS analysts use this data for 24/7 situational awareness, analysis, and incident response. Hence, participation in TIC and ensuring all agency traffic from “inside” their networks to “outside” systems is traversing a TIC access point were all that was required to be in full compliance with NCPS demands for E1 and E2. This E1 and E2 traditional telemetry pattern is depicted by the blue data flow path from an agency to DHS in Figure 1.

³ <https://www.whitehouse.gov/wp-content/uploads/2019/09/M-19-26.pdf>

⁴ <https://www.dhs.gov/cisa/managed-trusted-internet-protocol-services>

⁵ <https://www.dhs.gov/cisa/trusted-internet-connections>

⁶ <https://www.dhs.gov/einstein>.

⁷ <https://www.dhs.gov/cisa/national-cybersecurity-protection-system-ncps>

⁸ <https://www.gao.gov/assets/680/674829.pdf> (page 48)

Figure 1: Current On-Premise Telemetry Configuration



Security insights are concrete intelligence data formulated for the timely identification and prevention of imminent cyber threats. Security insights may include security rules provided in a rule-based language (e.g., Snort⁹ rules, Yara¹⁰ rules, etc.), attack signatures (e.g., malware hash, malicious macros, etc.), and indicators (e.g., blacklisted IPs, Email Header Indicators, etc.). Security insights are furnished by DHS and delivered to agencies to enable them to mitigate and counter cyber-attacks. Security insights may trigger internal processes and incident response within the agency's network to enact needed security reinforcements. Under the NCPS program, security insights can also be provided in the form of a DHS security alert to an agency concerning detected suspicious activity on the agency's network. This DHS alert may include a mitigation recommendation from DHS analysts, which will trigger an agency workflow to remediate the security threat. The flow of security insights from DHS to a TIC access point is shown by the blue data flow path from DHS to an agency in Figure 1.

2.2 How Cloud Impacts NCPS

As part of their IT modernization efforts, many agencies are utilizing commercial cloud products and adopting cloud email, collaboration, and software tools. Many agencies are using multiple CSPs in order to meet their mission needs and are utilizing all three cloud service models: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS).¹¹ When an agency creates a tenancy within a CSP, traffic between that CSP and the agency may no longer pass through a TIC access point or an NCPS sensor.

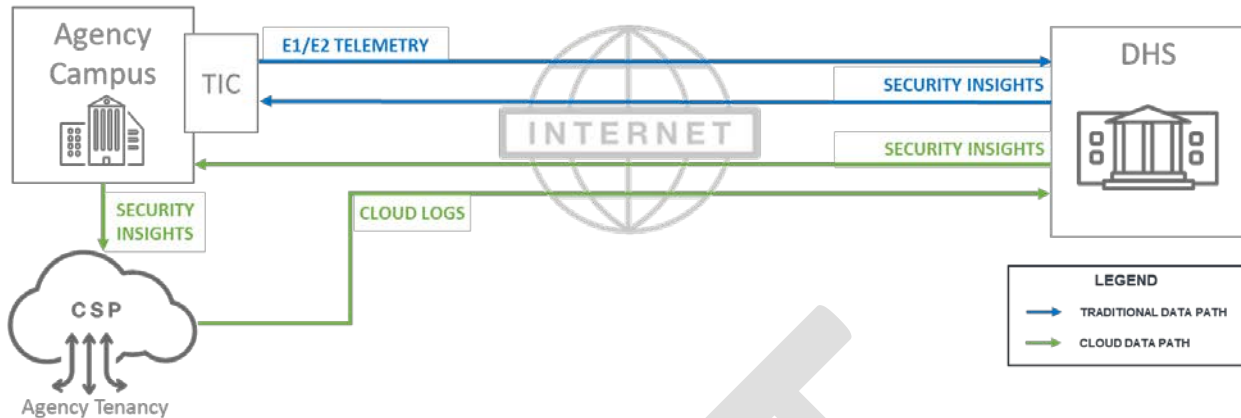
This new telemetry pattern is depicted in Figure 2. In this diagram, an agency still has some of its network traffic traversing the traditional TIC access point, but network traffic to one or more CSPs does not pass through the TIC access point. The blue data flow paths show the traditional flow of E1/E2 telemetry and security insights to and from an agency to DHS. The green data flow paths show the new data flows between the agency, the CSP, and DHS. Detailed use cases for data flows and telemetry collection and sharing are included in *NCPS Cloud Interface Reference Architecture: Volume 2*.

⁹ <https://www.snort.org>

¹⁰ <https://yara.readthedocs.io/en/latest/>

¹¹ Email as a Service (EaaS) is a sub-type of SaaS.

Figure 2: On-Premise and Cloud Telemetry Configuration



There are a wide range of CSPs and tenant-controlled security tools, and thus, there will be new data formats for telemetry (other than traditional network flows) and potential new formats for security insights for NCPS in the cloud. Data formats are discussed in Section 3 and vendor-specific details are discussed in *NCPS Cloud Interface Reference Architecture: Volume 2*.

Benefits of Sharing Cloud Security Data With DHS

There are several benefits associated with sharing cloud security data with DHS:

1. NCPS in the cloud provides DHS with the ability to gain situational awareness of threats and threat actors across the .gov domain, including on federal agencies' cloud communications. As a result, DHS can proactively respond to and mitigate cloud-based attacks against federal networks.
2. NCPS in the cloud extends DHS' security visibility and protection perimeter to include cloud-hosted software interactions and third-party services. Because threats to government systems are rarely targeted at a single agency, this increased visibility informs and enhances incident response capabilities and federal cloud security posture. All agencies and DHS benefit from that extended visibility.
3. NCPS in the cloud provides DHS with the ability to aggregate and correlate threat data generated and consumed in the cloud to aid in the timely discovery of security vulnerabilities and attack campaigns facing federal network cloud infrastructure.
4. Data gathered from the cloud network flow and cloud security logs provide DHS with additional intelligence and information to predict the changing security landscapes of both on-premise and cloud infrastructure, as well as to accurately plan, execute, and manage security countermeasures on the federal scale.
5. NCPS in the cloud provides a centralized model for log aggregation and analysis of a broad data set from federal cloud deployments, which result in a greater risk reduction for individual agencies as well as better availability of indicators of compromise to federal government information resources.

NCPS Roles and Responsibilities

Because the transition to the cloud introduces new roles, actors, and procedures (e.g., an autonomous CSP, absence of TIC, third-party cloud monitoring tools, etc.), the existing system for NCPS security insights transmission needs to be adapted. Specifically, in existing NCPS on-premise deployments, security insights in E2 are forwarded from DHS to the TIC access point (as shown in Figure 1). However, when an agency utilizes a CSP, E2 security insights continue to be transmitted to DHS via the TIC access points (as seen in the blue data flow), but agencies also need to “pull” E2 security insights from DHS and transmit those security insights to their agency tenant protections hosted by CSP(s) (as seen in the green data flow path in Figure 2).

A more detailed analysis of the responsible parties, as depicted in Figure 3 below, shows the shifting relationship for NCPS capability implementation. Traditional NCPS, shown in the leftmost column, was almost entirely implemented by DHS, with the agency only playing a role in provisioning a network tap for DHS observation and use. Under IaaS, PaaS, and SaaS deployment models the agency and CSP responsibilities greatly increase as they utilize CSP offerings and their own supplemental services to satisfy NCPS capabilities. Capabilities which must be implemented or coordinated by more than one party are designated as “shared”, with the participants identified. For example, shared capabilities for Traditional NCPS include determination of the Sensor Position, where the agency provisions a traffic mirror and DHS deploys the sensor with the associated ingestion capacity, media type, and protocol support. For cloud services, the shared responsibility may include one party (the CSP) responsible for implementation of the platform and another party (the agency) responsible for populating that platform with specific configuration values.

Figure 3: NCPS Functions and Responsible Parties (By Deployment Model)

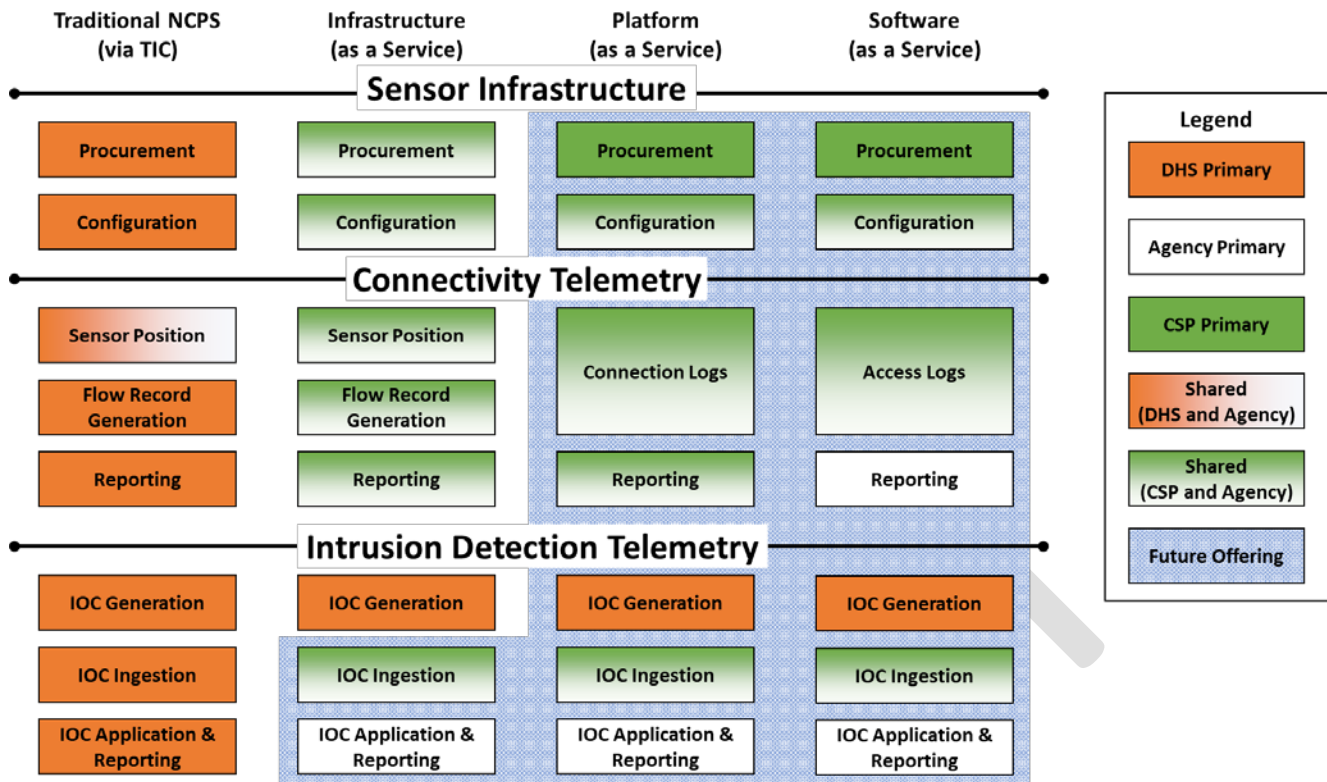


Figure 3 groups NCPS Functions into three categories: Sensor Infrastructure, Connectivity Telemetry, and Intrusion Detection Telemetry. Sensor Infrastructure functions are focused on the technologies, systems, and infrastructure required to perform policy enforcement. The Connectivity Telemetry group of functions enumerate security data collection for traffic traversing the sensor infrastructure. Similarly, the Intrusion Detection Telemetry focuses on functions related to threat discovery information sharing. An Indicator of Compromise (IOC), or forensic artifact to identify suspicious activity, can be used for rapid threat discovery in pattern matching countermeasures. In this case, the IOCs represent DHS threat intelligence for implementation in Intrusion Detection System (IDS) protections. In addition, Figure 3 also illustrates how cloud information sharing with DHS is not yet fully functional, with some functions designated as Future Offerings.

3 DHS CLOUD DATA COLLECTION

Replicating NCPS capabilities in the cloud for agencies demands an effective approach to facilitate the collection and correlation of an agency's cloud network data by DHS. Current cloud-based approaches are limited by existing data and log reporting capabilities provided by CSPs in the cloud. Because agency traffic may not traverse TIC access points when going to the cloud, the feasibility of having situational awareness on agency data and network activities in the cloud depends on both DHS and agency ability to effectively augment available traffic monitoring and reporting technologies to generate, collect, and aggregate E1/E2-equivalent data from cloud traffic generated by agencies. This section reviews supported cloud security log types and specifies a reference approach for their subsequent collection, aggregation, analysis, and storage.

3.1 Cloud Security Logs

The success of the NCPS program is directly impacted by the type of data or logs on which it operates. The selection of the cloud log types used as E1/E2-equivalent cloud telemetry will impact whether DHS is able to attain efficient and high-fidelity threat correlation. Different types of security logs are available for different cloud service models (IaaS, PaaS, or SaaS).

In order to satisfy NCPS in the cloud, network flow logs will initially be considered as the primary source of data. At a later time, additional types of cloud logs (such as those discussed in Appendix A) may be considered as a data source to NCPS in the cloud.

Connectivity Telemetry - Network Flow Logs

IP network traffic statistics describe the communication which takes place between endpoints and enables modern network management and security. Network flow log protocols specify how those statistics are to be generated and formatted. Participating network devices implement these protocols by generating, compiling, and organizing network flow records as traffic traverses them. When collected and sent to DHS, network flow logs will enable DHS to have situational awareness of an agency's cloud activities.

Cloud activities deployed in multiple service models (IaaS and PaaS specifically) can generate network flow logs. For example, when a node in the cloud initiates communication with another node, an intermediate sensor device with network flow record generation capability creates a flow record for that communication. Subsequent packets with the same network flow attributes update previously created flow records, which are continuously monitored and updated until the communication concludes. When communication is over, flow records are sent to a collector, where data logs are stored and further analyzed.

The specific collection location for network flow records is determined by the agency's requirements and those of DHS. In all cases, the following guiding principles help scope the generation of network flow records to be sent to DHS and those retained by the agency:

1. Network flow records are generated at the demarcation point where agency cloud tenancies interact with systems or components beyond the agency visibility, control, and administration.

2. Network flow record collection is enabled for all data sensitivity designations of agency information hosted in the cloud (i.e., Low and High sensitivity data will both require observation).
3. When agencies have more robust information collection needs for their internal purposes, as well as for post-collection processing or filtering of logs, network flow records can be used to align data sharing with agency and DHS MOU requirements (described in Volume 2).

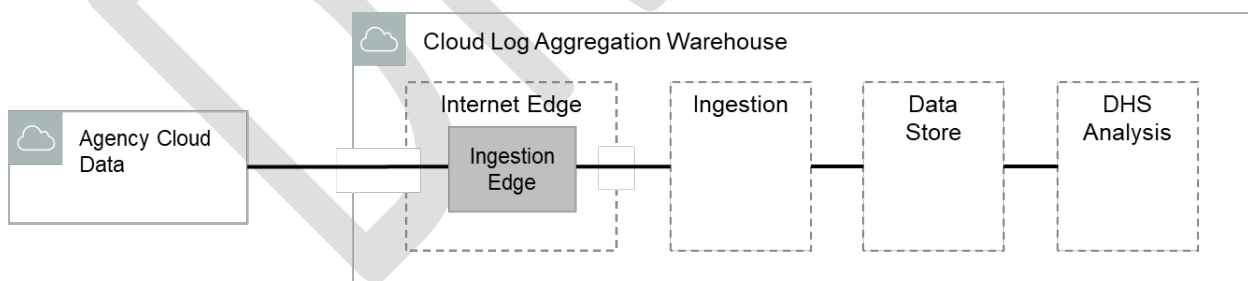
Appendix B discusses potential network flow data collection locations in more detail. *NCPS Cloud Interface Reference Architecture: Volume 2* details specific use cases for sensor deployment and discusses network flow log protocols and formats.

3.2 Cloud Log Aggregation Warehouse Overview

The Cloud Log Aggregation Warehouse (CLAW) is a DHS-deployed architecture for the collection and aggregation of NCPS data from agencies using commercial CSP services. While agency NCPS data is currently aggregated on-premise at DHS, the CLAW is deployed in the cloud to aggregate agency security logs and EINSTEIN sensor data that originates in the cloud. The CLAW presents a functional, module-based architecture to ingest, store, and analyze security logs and sensor data from agencies. It is geared towards enabling secure and efficient methods to process cloud data in a manner that offers DHS a similar level of situational awareness provided by current EINSTEIN on-premise deployments. The CLAW architecture supports log aggregation at multiple locations (optimized for performance, cost, and efficiency) utilizing centralized threat discovery with distributed analytics.

The CLAW pipeline is depicted in Figure 4, which shows how it supports the collection and ingestion of cloud data. Enabling NCPS capabilities in the cloud involves: (1) data collection and aggregation from agencies, (2) data transfer from agency cloud environments into the CLAW architecture, and (3) data analysis by DHS. These three steps will be discussed in the next three subsections. Specific details for unique use cases will be developed in detail in Volume 2.

Figure 4: CLAW Pipeline (Agency Perspective)



3.2.1 Agency Data Collection

Cloud data collection is the first step towards achieving data-driven situational awareness in the cloud. Data collection in the cloud impacts sensor positioning and placement, virtual network layout and boundary partitioning, specific log data types to be filtered and collected, and cloud-based data protection mechanisms. To be secure, agency data collection from data sources/taps to data sinks must enforce data protection properties like data confidentiality (prevents unauthorized disclosure), integrity (tamper-proof),

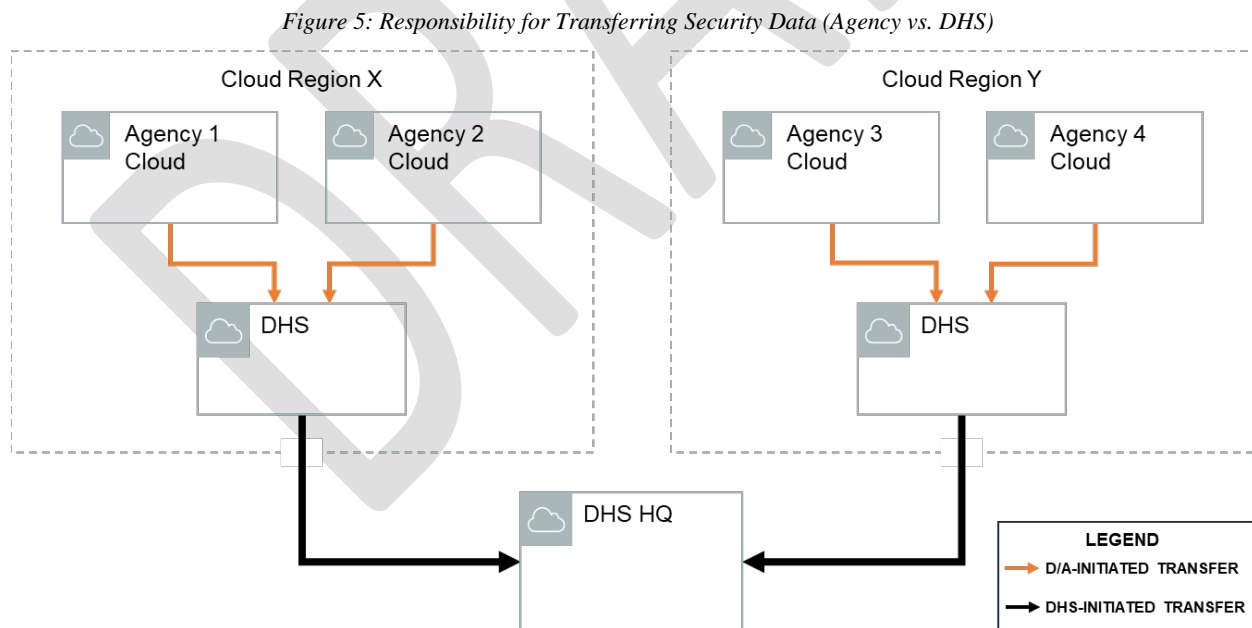
and authorization (roles and access control). To achieve these security enforcements and to prevent leakage/corruption, deployed collector modules implement a secure channel (e.g., via encryption) and access control mechanisms between data taps/interfaces and interim collection points. The collected data will be securely aggregated, pre-processed (if necessary), and transferred to DHS for further analyses and threat correlation.

The agency is responsible for provisioning the cloud setup to enable the collection of network flow logs and providing the logs to DHS. Specifically, this implies that every agency is responsible for placing sensors and/or network taps at the cloud virtual network boundary to collect network flow records. Appendix B discusses potential network flow data collection locations in more detail.

3.2.2 Agency Data Transfer

Data transfer involves the mechanism by which agency cloud data is transferred to DHS after the data is collected. The frequency or timeliness of the transfer can be based on time differentials (e.g., polling every twenty minutes) or based on storage utilization (e.g., after every 20MB of new data) of accumulated new data.

The data transfer method supports a push mechanism. Data push specifies the mechanism by which the agency initiates and transfers their collected log data to DHS. As such, the agencies themselves are responsible for moving this data, as well as providing the infrastructure to support the transfer. Figure 5 depicts how agencies in different cloud regions can push their data to DHS.



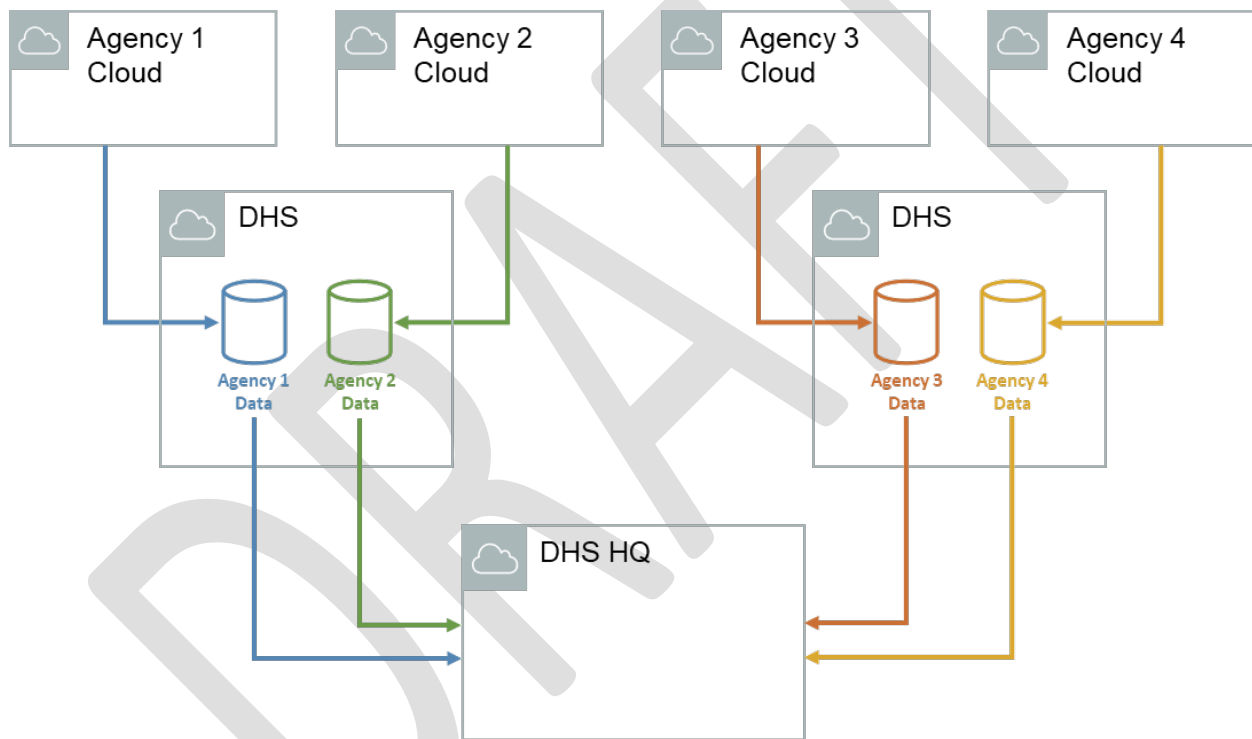
The agency is able to push their data to a “local” DHS CLAW location, reducing data transfer costs, technical complexity, and transmission latency. The CLAW architecture supports collocating CLAW aggregation points with agency tenants on major CSPs. Any additional data aggregation or consolidation required will occur within DHS’ purview.

3.2.3 Analysis of Agency Data

DHS analysis focuses on providing the environment and tools to correlate and discover threats from application and network data that has been collected and aggregated from the agency cloud tenants. Current analytics approaches involve signature-based (pattern recognition) and non-signature-based (heuristic and statistical) analytics for identification of IOCs and for identification of anomalous activities. Analysis will also bring in enrichment data to enhance the analysis results.

Figure 6 shows how cloud data from individual agency cloud tenancies is collected and analyzed at DHS cloud sites. Each agency has a separate analysis compartment to prevent data comingling and corruption. Analysis results obtained will subsequently be sent to DHS HQ for processing and assimilation (i.e., for threat detection and correlation, and synthesis of security indicators).

Figure 6: Agency Log Ingestion (Autonomy Preserved with Log Isolation)



Analysis will be performed from a central location with standardized tools. Those centralized tools will be able to interact with the sensor data distributed across CLAW locations. The data will be ingested and processed at a “local” CLAW location; in other words, the agency data will not be backhauled to a central repository. This will provide analysts with global situational awareness without requiring a corresponding centralized data store or requiring multiple copies of the same tools at each of the distributed data stores.

4 CONCLUSION

As agencies move more of their applications and services to CSPs, the NCPS program is evolving to ensure that security information about cloud-based traffic can be captured and analyzed and that DHS analysts can continue to provide situational awareness and support to the agencies. This document provides an overview of how NCPS will provide capabilities to agencies using CSPs and what cloud logs will be collected and transferred to the “CLAW.” The companion document (*NCPS Cloud Interface Reference Architecture: Volume 2*) contains reference solutions that consider different cloud service models and specific guidance unique to individual CSPs (e.g., AWS, Microsoft, etc.). Together, these two documents provide guidance for how an agency can adapt their cloud environments to allow for security data to be sent to NCPS.

DRAFT

APPENDIX A: TRANSACTION & EVENT SECURITY LOGS

Cloud transaction and event security logs are actively generated and stored to collect security and statistical information on observable cloud activities. The types of logs that are available are dependent on the type of cloud service model (IaaS, PaaS, or SaaS) and the specific CSP.

When transaction records and event security logs are collected and stored, they may be sent to DHS under certain circumstances (as defined in Volume 2) to enable DHS to have situational awareness of additional agency cloud activities (beyond network flow). The applications that are deployed at the IaaS, PaaS, and SaaS domains in the cloud actively generate transaction and event security logs for analysis by DHS and the individual agency. Popular cloud applications associated with transaction and event security logs include:

- Web (HTTP)
- Email (SMTP)
- Naming (DNS)
- Identity and Authentication services (Active Directory and Certificates)

Typically, these logs are first generated by the application server when clients request server resources. Subsequent interaction with the same attribute would update previously created transaction records, which are continuously monitored and updated until the communication ends. When the communication is over, the transaction records are sent to a collector, where data logs are stored and further analyzed.

The format and fields of the transaction and event security logs are defined by the underlying application and not necessarily by cloud providers. For example, a Web application transaction and event logs will include fields such as HTTP headers, Client User Agent, Content type, number of bytes, TCP Port numbers, TLS session information, and timestamp.

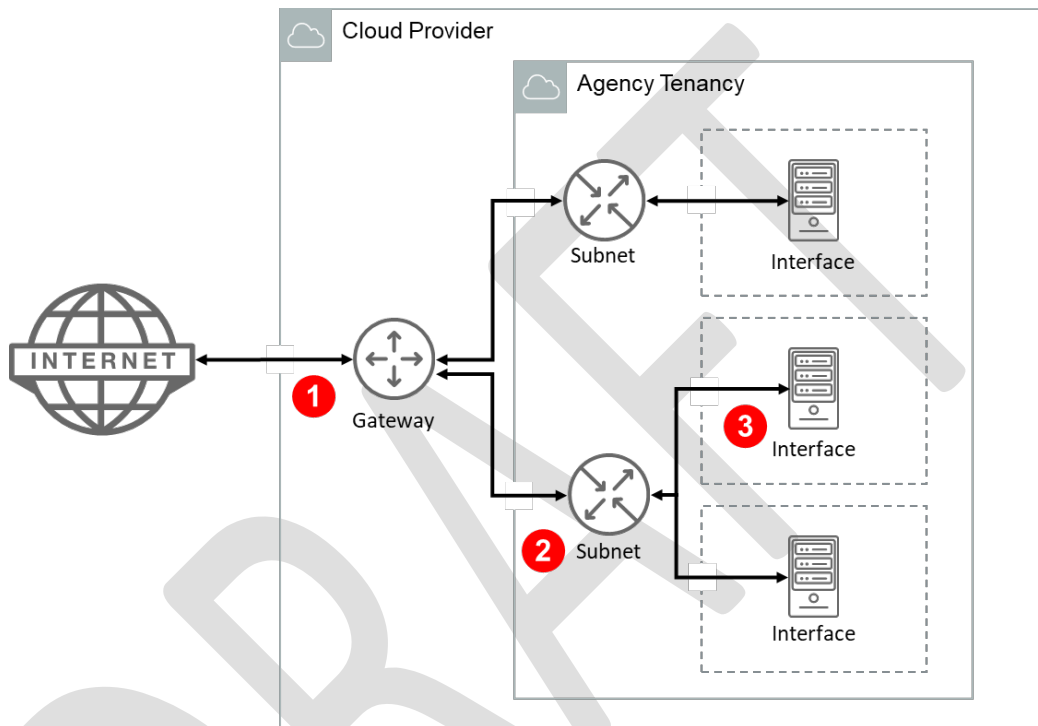
If stored transactions and event security logs are collected and sent to DHS, these logs will enable DHS to have visibility and situational awareness into an agency's cloud application activities. Moreover, these records can be used during post-event analysis, incident response, and potential root cause analysis on known and perceived threats.

More details about transaction and event security logs will be discussed in Volume 2.

APPENDIX B: FLOW RECORD COLLECTION LOCATION

These flow record collection guidelines apply differently depending on where in the agency's cloud system the demarcation point occurs. Three typical deployment locations for network flow generation in IaaS deployments are shown below in Figure 7. In addition, this appendix enumerates some conditions under which each is suitable.

Figure 7: Network Flow Log Generation Positions for IaaS



There are three potential collection locations for Agency Tenancy network flow records. Each of these collection locations has unique visibility scope and detail.

1 Gateway

The first potential collection location is the Gateway at the Internet to Agency Cloud Tenancy interface(s). Collection of network flow records at the Gateway allows monitoring of all traffic to and from all agency cloud resources. When NAT (Network Address Translation) is used, the agency must ensure that the records gathered reports the public IP addressing. The traffic monitored at this location may include agency "private/internal" sources not typically monitored by the NCPS sensors. The records gathered here may require processing at the post-collection phase to exclude those records prior to being sent to DHS. An example would be a publicly accessible web site hosting publicly available information, where the agency is not cohosting any additional resources on same cloud tenancy.

2 Subnet

The second potential collection location(s) are the subnet(s) utilized by the agency tenancy to provide cloud server access. Collection of network flow records at the subnet level allows for monitoring of all traffic to and from cloud server(s) on each individual subnet. The “private/internal” and “public” data flows can be separated, thereby constraining the sharing of data flow information with DHS to only the “public” resources, reducing post-collection processing requirements. An example would be a publicly accessible web site hosting publicly available information and internal human resources applications in the cloud with “public” and “private/internal” data flows (respectively) destined for resources on independent subnets. The subnet with the publicly available information is provisioned to share network flow records with DHS.

3 Interface

The third potential collection location(s) are the interface(s) utilized by the agency tenancy to provide access to cloud virtual servers. Collection of network flow records at the interface level allows for monitoring of all traffic to and from the individual interfaces on each of the cloud server(s) that has been properly configured to provide this capability. The “private/internal” and “public” data flows are separated by the individual virtual interfaces. This type of collection location provides the finest granularity for the network flow records, minimizes the post-collection processing requirements, and permits greater insights for event correlation and analysis. An example would be a publicly accessible web site hosting publicly available information in the cloud with the public data flows destined for resources on independent interfaces. The interfaces with the publicly available information are provisioned to share network flow records with DHS.